White Paper

Consolidating x86 on System p on Power Systems

Better availability, performance, and security–comparable TCO and lower energy consumption

November 2010
Josh Krischer

Contents

Consolidating x86 on System p on Power Systems

Josh Krischer & Associates GmbH
Enterprise Servers, Storage and Business Continuity

Josh Krischer is an expert IT advisor with 41 years of experience in high-end computing, storage, disaster recovery, and data center consolidation. Currently working as an independent analyst at Josh Krischer & Associates GmbH, he was formerly a Research Vice President at Gartner, covering enterprise servers and storage and disaster recovery techniques from 1998 until 2007. He authored many papers and spoke on these topics and others at a multitude of worldwide IT events, including Gartner conferences and symposia, industry and educational conferences, SNW, as well as major vendor events.

Introduction

In the 90-ies many organizations decentralize their IT infrastructure to relatively-inexpensive x86-servers that could meet the processing needs of departments and branch locations. The addition of a single application could typically add three to five servers to the data center for such things as production, development, testing, and backup. This trend caused server sprawl which created huge, difficult to manage, and low-utilization infrastructures. It was not uncommon to see large data centers requiring huge floor space, consuming lots of power, and emitting plenty of carbon-dioxide. Many organizations could not even quote the number of the servers in their data centers. The situation started to change with the introduction of several server virtualization techniques for UNIX and x86 platforms.

Beginning in the early 2000s, a trend reversal emerged in which organizations began consolidating and deploying a centralized IT-structure. This development was enabled by emerging virtualization techniques, and driven by the goal to lower costs, increase efficiencies, as well as provide higher Service Levels Agreement (SLAs) and better resiliency requirements.

By consolidating servers into more powerful, energy-efficient virtual-partitions/machine-hosts, organizations can replace outdated, power-hungry hardware, increase server utilization, and thus, achieve significant savings in space, power, and cooling requirements. In addition to energy cost savings, such consolidation also results in environmental benefits. However, such consolidation increases the risk and impact of malfunction, which requires careful selection of the target hardware, hypervisor, and software. The new technology has to provide better availability, advanced disaster recovery techniques, and easier to use management and monitoring tools. To ensure benefits, the consolidation deployments should be seen as an ongoing IT-organization strategy, not a one-time project.

> "*However, such consolidation increases the risk and impact of malfunction, which requires careful selection of the target hardware, hypervisor and software.*"

This white paper provides an overview of IBM POWER7 server technology as a powerful server platform supremely suitable for enterprise-level server-consolidations, and analyses two case studies of customers who deployed a migration from x86 to Linux and

Consolidating x86 on System p on Power Systems

AIX on System p. There are also cases of users who migrated from x86 to System i on Power platforms, but these can be the subject of another paper.

About POWER7 and System p

On February 9th, 2010 IBM announced POWER7 technology, with technology breakthroughs resulting in increased computing power and state-of-the-art virtualization. POWER7 servers deliver up to twice the performance and four times the virtualization partitions for the same price, with the same energy usage as similar POWER6 systems. In addition to performance and energy efficiency, the new POWER7 servers technology supports Active Memory Expansion, which allows the effective memory capacity of the system to be much larger than the installed physical memory, as well as PowerVM virtualization, which allows organizations to deploy server consolidation with optimal utilization of system resources.

On August, 17th 2010 IBM enhanced its offerings by adding the new high-end IBM Power-795 system, four entry-level POWER7 processor-based servers designed specifically for mid-market clients, a POWER7-processor-based workload-optimized Smart Analytics System, and a new revision of AIX (IBM's UNIX version) operating system – the AIX 7.

Processors and Servers

The POWER7 chip includes 4, 6 or 8 cores per socket with up to 4 threads per core and 3.0 to 4.14 GHz clock cycle. All models have an integrated eDRAM L3 Cache with fast access time (see Figure 2)
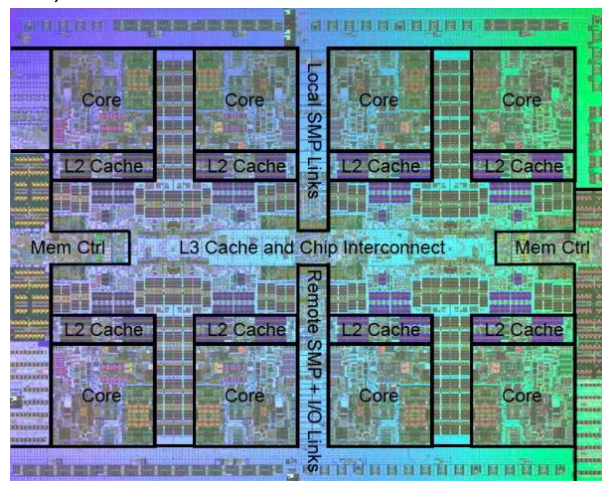

Figure 1 POWER7 Processor Chip

Consolidating x86 on System p on Power Systems
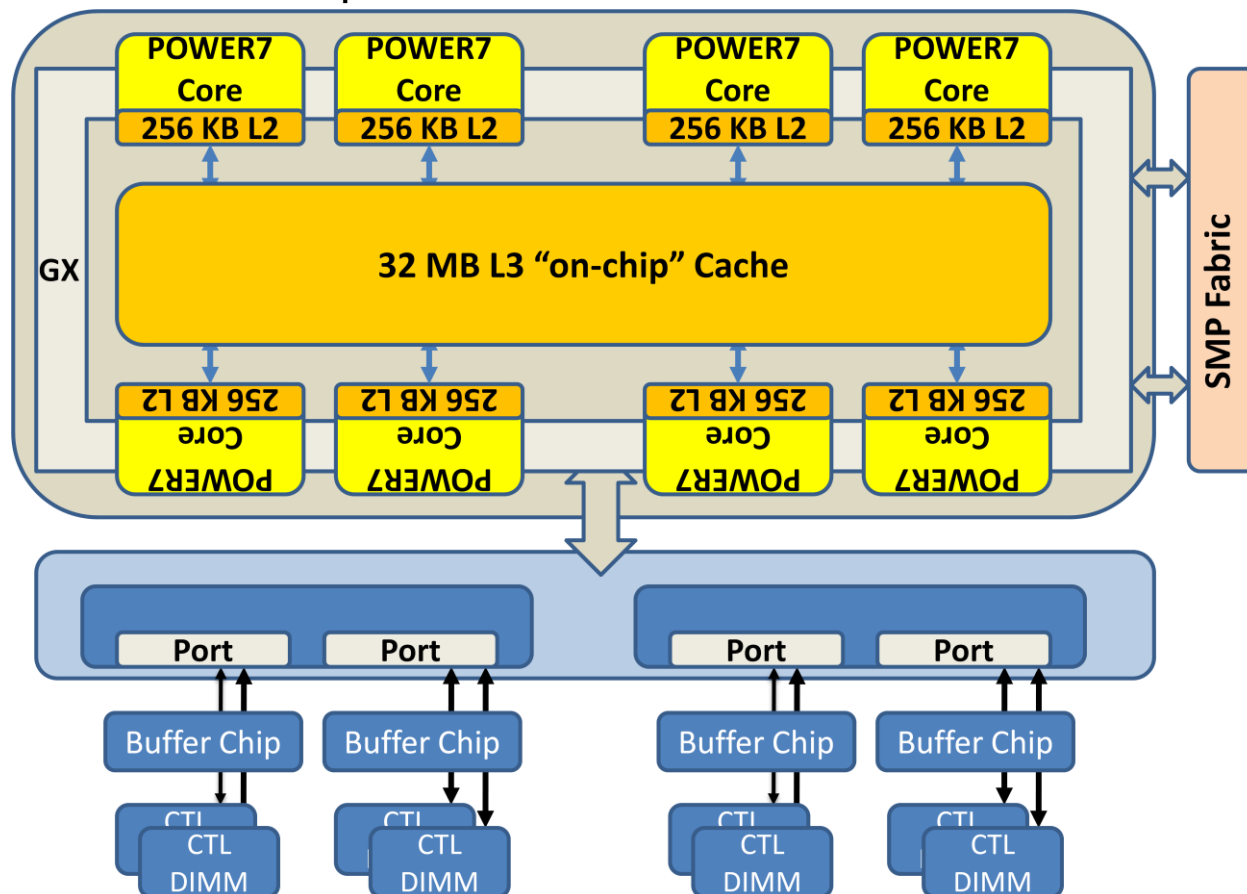
POWER7 Processor Chip

Figure 2: POWER7 Chip structure and Interfaces

There are ten IBM POWER7 server models:

❑ IBM's Power 795 is a new 24-256 core system (3.7, 4 or 4.25 GHz) supporting up to 8 TB of memory, and providing over four times the performance in the same energy envelope as the fastest Power 595 based on POWER 6 technology. It uses IBM's leading-edge EnergyScale technology that varies clock frequencies depending upon workloads. The latest PowerVM™ virtualization software allows customers to run over 1,000 virtual servers on a single physical system. Power Flex is a new high-availability technique to allow running applications to migrate from one system to another in order to perform system maintenance without downtime.

❑ IBM's Power 780 Server supports 16 to 64 processor-core configurations running at 3.8 GHz or up to 32 POWER7 processor cores at 4.1 GHz in TurboCore mode. Running in TurboCore mode, in addition to the faster clock speed, allows the use of double the normal cache size in order to achieve top performance. This model is best suitable for virtualized consolidation of business-critical workloads, including large database-serving and transaction-processing applications.

Consolidating x86 on System p on Power Systems

❑ IBM's Power 770 Server is aimed at users with smaller workloads, but who still require enterprise features such as advanced virtualization and flexible scalability. Two high-performance processor options are available on the Power 770; 16 to 64-core processor configurations running at 3.1 GHz, or 12 to 48 core configurations running at 3.5 GHz.

❑ IBM Power 750 Express Server is a compact 4U (4 EIA units) rack-mount server. It has 1 to 4 sockets using 6-core and 8-core processor options running at 3.0 GHz, 3.3 GHz or 3.55 GHz. This model is best suited for x86 server consolidation running Linux and Windows workloads

❑ On August 2010 IBM announced Express Servers Power 710, 720, 730, and 740. They offer mid-market clients the outstanding performance, energy efficiency, and other benefits of POWER7 technology in compact rack-mount or tower packages. These models are priced within 20 percent of Intel and Windows servers.

❑ The IBM Power 755 server is designed for High Performance Computing, in particular to run highly-parallel, computationally-intensive workloads such as weather and climate modeling, computational chemistry, physical simulations, or petroleum reservoir modeling. It has 32 POWER7 processor cores running at 3.3 GHz and memory capacity of up to 256 GB. Using optional Enhanced Dual Port 12X Infiniband adapters, a clustered configuration with up to 64 Power 755 compute nodes and a total of 2,048 processor cores can be created.

❑ The IBM Smart Analytics System 7700 provides a powerful analytics platform to quickly process vast amounts of data that can be easily deployed and customized. This solution includes several pretested Power Systems 740 Express server configurations with pre-integrated AIX, IBM DB2 and InfoSphere Warehouse software.

Virtualization and Workload Management

PowerVM of System p enables users to use virtualized LPARs with IBM AIX or Linux operating systems. Unlike most of x86 systems, which use a software hypervisor, System p has a firmware-based POWER Hypervisor (PHYP) with micro-partitioning capability, which ensures better resilience and lower overhead. The micro-partitioning offers better scalability and granularity for the virtual machines with the possibility to assign partial CPUs to LPARs, which results in higher utilization. The fine granularity allows for assigning virtual partitions computing power in 1 percent of physical processor increments[1]. The maximum number of virtual machines or partitions is 32 (Power 770, 780) or 64 (Power 795) dedicated partitions (processing core granularity, or up to 160 micro-partitions,

"Unlike x86 systems, which use a software hypervisor, System p has a firmware-based POWER Hypervisor (PHYP) with micro-partitioning capability, which ensures better resilience and lower overhead."

---

[1] However the minimum size of a partition is 1/10 of the physical processor power.

which is significantly more than VMware.

Virtual I/O Server is a special-purpose partition that allows the sharing of physical resources between logical partitions to allow for more efficient utilization, particularly in server consolidation projects). The Virtual I/O Server owns the physical resources (SCSI, Fibre Channel, and network adapters), and allows virtual partitions to share access to them, thus minimizing the number of physical adapters in the system. The Virtual I/O Server eliminates the requirement that every partition owns a dedicated I/O or network resource, and therefore allows for assigning a large number of partitions.

LPAR's CPU resource usage can be managed by the Enterprise Workload Manager, which monitors CPU usage and assigns processors' resources to requests according to Service Level Agreements.

Another interesting feature of PowerVM is Active Memory Sharing[2] which allows multiple partitions to share a common pool of physical memory. The physical memory can be assigned to multiple partitions either in a dedicated or in a shared mode. The system administrator can assign some physical memory to a partition and some physical memory to a pool that is shared by other partitions. In this mode, the system automatically decides over the optimal distribution of the physical memory to partitions, and adjusts the memory assignment based on partition load. Active Memory Sharing increases the memory utilization of a system by decreasing overall memory requirements. VMcontrol system pools support for Linux enhances IBM Systems Director virtualization management.

PowerVM Live Partition Mobility

"PowerVM Live Partition Mobility allows to non-disruptively move a running logical partition (operating system and running applications), from one Power system to another ."

PowerVM Live Partition Mobility allows to non-disruptively move a running logical partition (operating system and running applications), from one Power system to another. This feature avoids planned outages for hardware or firmware maintenance/upgrades, and allows for preventive failure management if a server detects a potential failure, thus, reducing or avoiding unplanned downtime. Live Partition Mobility offers easy migration between servers based on POWER6 and POWER7 processor–technology, and leverages the POWER6 compatibility modes provided by POWER7-based servers.

RAS (Reliability, Availability, and Serviceability) Design and Features.

High-availability systems are designed to have few (or, ideally, no) single points of failure (SPOF) and no single points of repair (SPOR), through the use of redundant components and architectures. Current semiconductor technology is reliable; most failures are caused by bad connections between components. Highly-reliable systems are therefore designed with fewer components and interconnections, as well as component redundancy, and are

---

[2] Active Memory Sharing is only available with the Enterprise version of PowerVM.

built from highly-reliable parts. An 8-core POWER7 server includes 75% fewer processor chips (and chip socket interfaces) as a double-core-per-processor design The L3-cache is integrated into the processor chip, which improves performance (significantly reduced latency) and reduces the number of components and connections.

Power 770, 780, and 795 systems have redundant system clocks and service processors for configurations with more than two Central Electronics Complex drawers (CEC), as well as redundant and hot-swappable cooling blowers and power supplies. A high level of redundancy is employed in the memory and the associated caches. For example, POWER7 cores include redundancy bits in L1-I, L1-D, and L2 caches, as well as L2 and L3 directories, and the Power 770, 780, and 795 main memory DIMMs contain an extra DRAM chip. In addition to the 64 bits error correction code in the main storage (single-bit correction, and double-bit detection) the System p servers use Chipkill, which is an advanced ECC memory technology designed specifically for the NASA *Pathfinder* mission to Mars. Chipkill enables a system to sustain the failure of an entire DRAM chip with no performance degradation. Additional protection is Memory Fencing: Defective memory discovered at boot time is automatically switched off and marked, so affected memory will not to be used on subsequent reboots.

The Power 795 also supports Active Memory Mirroring for Hypervisor to prevent downtime in case of a memory failure in the hypervisor area, which otherwise could bring down all partitions on rare occasions.

Today's information availability demands include more than hardware RAS features: they encompass advanced business continuity, data protection, and end-to-end infrastructure management. System p supports PowerHA SystemMirror for AIX and Linux, formerly HACMP (High Availability Clustered Multi Processing), and HACMP/XD (XD - eXtended Distance). PowerHA scales up to 32 nodes running AIX, or eight nodes running Linux. It constantly monitors the status of servers, networks, and applications to detect failures or performance degradation, and responds by automatically restarting a failing application, selecting an alternate path to storage, as well as taking care of all network connections in the process.

PowerHA XD extends PowerHA's capabilities across geographic sites with Metro-Mirror remote data mirroring and failover using this mirrored data. **PowerHA SystemMirror for AIX** also provides host-based IP-based data mirroring. IBM's PowerHA SystemMirror for AIX Standard Edition monitors operation to detect failures, and provides automated recovery of business application to recovery resources. It exploits IBM's remote mirroring techniques of disk storage subsystems. PowerHA SystemMirror for AIX moves with the automated migration environment, thus, practically eliminating planned outages by transferring users, applications and data to the recovery infrastructure before a scheduled maintenance or upgrade are performed. Enterprise Edition expands the HA/DR support to EMC's SRDF techniques, and provides many usability enhancements including a wizard to configure and deploy host-based mirroring solutions (GLVM configuration wizard). GLVM is the host-based remote-mirroring-over-IP network–solution for AIX that enables users to lower the costs of disaster recovery infrastructure deployments.

Open HyperSwap, is a new DS8700 and DS8800 function within the Tivoli Storage Productivity Center for Replication. Open HyperSwap (comparable to HyperSwap for

Consolidating x86 on System p on Power Systems

System z) is an important feature improving the continuous availability of AIX by managing a set of planned and unplanned disk system outages for Metro Mirror PPRC-capable disk systems. It provides an ability to swap IBM DS8000 volumes (from the primary storage system to the secondary storage system) in seconds, reacting to operator command or in automatic (no operator interaction) fashion upon a storage system failure. It is designed to scale to multi-thousands of volumes. This non-disruptive feature, which allows applications to keep using the same device addresses, is configured and managed by Tivoli Storage Productivity Center for Replication V4.2. System p also supports third-party HA solutions. The Veritas Cluster Server (VCS) software provides automated failover support on both AIX and Linux platforms. VCS also operates on Hewlett Packard and Sun systems, thereby enabling an installation to install a common automation package in a heterogeneous environment. Another third party solution is the EchoCluster for AIX from Vision Solutions, which is simpler and more affordable than PowerHA. SystemMirror and EchoStream from the same company also support Continuous Data Protection (CDP).

Power Flex

Power Flex is a new technique that requires two or more Power 795 systems, PowerVM Live Partition Mobility, and Flex Capacity Upgrade on Demand options. This solution enables clients to shift running applications from one system to another to perform system maintenance without downtime, helping to balance workloads and smoothing peaks in demand and ensuring non-stop operation with stable performance.

Capacity Back-up for business continuity

Customers of IBM's high-end Power Servers may use a specially configured and priced Capacity BackUp system at the recovery site. The CBU server contains a minimum number of active processors, and usually a larger number of inactive processors that can be temporarily activated when needed for disaster recovery. Upon detecting the loss of the main site, a PowerHA SystemMirror cluster can automatically activate the CBU processors to run the business-critical application at the recovery site.

Performance

Despite the slower clock speed, the Power7 delivers four times the performance per socket than its previous generation; part of this performance increase is achieved by processor technology, and part by advanced design such as its integrated Level 3 cache. Power7 technology provides core-to-core performance improvement of up to 50% over the Power6. Additional factors improving performance are DDR3 running 1066 MHz (bandwidth increase by a factor 2.46) memory modules, and workload-optimizing features for performance in both transactions and massive parallel processing, such as TurboCore, maximizing per-core performance for databases, and MaxCore for better parallelization and high capacity.
In High Performance Computing (HPC), POWER7 processors support the AltiVec instruction set, and extended VSX (Vector Scalar Extension) SIMD (single instruction

---

multiple data) acceleration, which can execute up to eight single-precision or double-precision floating point operations per clock cycle per core, to improve parallelism. Active Memory Expansion provides more memory for applications by compressing and decompressing data in-flight. Larger available memory reduces paging to external disk devices.

Dynamic Energy Optimization

The POWER7 processor includes intelligent energy features that help to dynamically optimize energy usage based on performance. The EnergyScale feature in the processor work with the Systems Director Active Energy Manager software to dynamically optimize processor clock frequency according to thermal conditions and system utilization. It can also turn off processor cores or limit the energy draw across one or a group of Power servers, and track environmental data from applications used to monitor air conditioning units, Uninterruptible Power Supplies (UPS), and intelligent Power Distribution Units (PDUs).

According to SPECint_rate2006 the new 256-core IBM Power 795 offers more than five times better energy efficiency compared to servers from Oracle and HP.[3]

Software

AIX 7- Express Standard and Enterprise Editions

AIX 7 comes in three versions, the Express Edition, the Standard Edition, and the Enterprise Edition. The Express Edition is easy to deploy, and hence, best suited for entry servers, blades and targeted at small and mid-sized business clients and server consolidations deployments. The Express Edition supports up to 4 cores per virtual machine with 8GB of memory per core.

The Standard Edition is suitable for the most demanding workloads, with the ability to build large (up to 64) partitions with large amount of memory per core.

The AIX Enterprise Edition combines AIX 7 and enterprise management functions into a single, easy to order offering, thus, simplifying the process of managing virtualized AIX environments. This edition provides real-time information required for effective day-to-day data center operations. Among other functions, it manages workload partitions across multiple systems, manages energy usage, monitors utilization and configuration changes, collects and report virtualized resource usage, as well as providing predictive reports and

---

[3] Efficiency is measured in performance per watt, using SPECint_rate2006 as the measure for performance and the maximum power usage for the IBM Power 795 and from HP QuickSpecs and Sun SPARC Enterprise Site Planning Guides as the measure of energy usage. .
SPECint_rate 2006 results: IBM Power 795 with 256 cores, 32 processor chips, and four threads per core had a peak result of 11,200. HP SuperDome with 128 cores, 64 processor chips and one thread per core had a peak result of 1,648. Sun SPARC Enterprise M9000 with 256 cores, 64 processor chips and two threads per core had a peak result of 2,586.
SPEC® and the benchmark names SPECrate®, SPECint®, and SPECjbb® are registered trademarks of the Standard Performance Evaluation Corporation. For the latest SPEC benchmark results, visit http://www.spec.org.

Consolidating x86 on System p on Power Systems

data to simplify capacity management. It configures and deploys virtual AIX images and manages pools of virtual resources as a single system. Workload Partitions Manager for AIX (WPAR Manager) manages AIX 7 partitions across multiple systems and enables the Live Application Mobility.

Linux

POWER7 supports both Red Hat Enterprise Linux (RHEL 5.5, 6) and SUSE LINUX Enterprise Server (SLES 10 and 11SP1). Both RHEL and SLES run natively on Power systems, offering a scalable fully-supported environment for open source applications. The PowerVM Lx86 feature creates a virtual x86 Linux application environment on Power-processor-based systems, so most 32-bit x86 Linux applications can run without code recompilation. It significantly simplifies the migration of x86 Linux applications to Power servers, enabling customers to exploit the availability and scalability benefits, as well as the savings in energy and management of the series. PowerVM Lx86 is a standard, non-chargeable feature of the PowerVM Editions.
AIX Future developments
IBM made a Statement of Direction for AIX, which includes increasing the scalability to support up to 1024 threads in a single AIX partition, new runtime management options, and integrated regulatory compliance features. Another SOD is for PowerHA SystemMirror for AIX to be tightly integrated with Cluster Aware AIX 7, or, in another words, built-in clustering to simplify configuration and management of scale-out workloads and high availability solutions. Management of AIX systems' pools will be simplified with profile-based configuration management.

Consolidation Case Studies

The following case studies describe users' experience and benefits in migrating from the x86 to the POWER6 platform. The migration cases to POWER7 are still in preparation, but these case studies already strongly demonstrate customers' benefits; the benefits in migrating to POWER7 should be significantly larger.

Local Authorities IT Data Center in Germany

The center wanted to deploy SAP ERP 6.0 and consolidate other applications such as a resident registration system for around 170 smaller local government authorities on distributed x86 systems, as well as an application to process digital passport and identity card applications intended for the German federal printing plant in Berlin.

Customer objectives

- ❑ Set up a future-proof flexible IT platform for hosting SAP applications.
- ❑ Establish Linux as a strategic operating system platform.
- ❑ Consolidate an existing SAP system for facilities management onto the new platform.
- ❑ Ensure failsafe operation of the SAP production systems.

Consolidating x86 on System p on Power Systems

All these translate to better customer service, and lower operating costs. The CoD hardware can be used to seamlessly answer future scalability requirements.

South-American Manufacturer of Agricultural Products

Customer's Initial Situation

The customer ran a JDEdwards Oneworld XE ERP application on an x86 platform running in a Microsoft Windows Server 2000 environment, which was supported by a Microsoft SQL Server 2000 database. The infrastructure scaling was inflexible and unable to match the requirements of international business expansion. For example; to get the JDEdwards solution to meet its needs, the company had to invest in considerable customization and integration work, introducing a number of third-party applications to handle specific functional requirements. As a result, it became almost impossible to upgrade the JDEdwards system, and the company found itself caught in a dilemma: Oracle was no longer willing to support Windows 2000, but the installed JDEdwards solution was not compatible with more recent versions of Windows. The complexity of the systems landscape was also leading to unreliability, and frequent patching of Windows operating system caused planned downtimes.

Software development costs were high, achieving standards-compliance across a mixed environment was difficult, and data-center power-demands were rising. As a company operating in the agricultural sector, it is very concerned about lowering its carbon footprint.

Customer objectives

- Set up a future-proof, flexible, and scalable IT platform.
- Simplify the infrastructure.
- Enhance Business Continuity.
- Better performance.
- Lower energy costs.
- Obtain better support of certified consultants with strong presence in region.

Solution

- Migration to SAP Business Suite 7, running in an IBM AIX environment on IBM POWER6 processor-based blade servers in an IBM BladeCenter. The BladeCenter also contains Intel Xeon-processor-based blades
- Integration with Intel Xeon-processor-based HS22 blade servers running Windows environments for the company's email servers, and other applications
- Deployment of PowerHA with PowerVM for better availability. The PowerVM Live Partition Mobility feature allows the IT team to seamlessly switch virtual server workloads from one blade to another for load-balancing with no impact on the SAP applications in production.

Consolidating x86 on System p on Power Systems

❑ Dual-node IBM PowerHA clusters for the SAP ERP and SAP NetWeaver Business Warehouse production environments–if one SAP application server fails, the other is able to take over the workload.

Reasons for selecting this solution

The SAP on Power Architecture was chosen because it fulfilled these requirements:

❑ Ability to use a single ERP system, which was provided by the SAP Business Suite's functionality.

❑ Flexibility: The SAP Business Suite offers support for a wide range of business processes, and is versatile enough to adapt to changing business requirements without the need for third-party software or excessive development work.

❑ Migration from Windows onto a Unix platform would offer better security and stability with less need for patching.

❑ Hardware vendor capable of providing a reliable, high-performance hardware platform for Unix, SAP and new databases.

❑ Flexibility to integrate x86 blades in the IBM BladeCenter.

❑ Scalability to accommodate future workloads.

❑ Vendor support in migration

The New IT infrastructure:

IBM BladeCenter with IBM POWER6-based blade servers running AIX: JS12 blades with 3.8GHz processors for the SAP applications, and JS23 blades with 4.2GHz processors for the underlying Oracle databases.
Storage was provided by an IBM System Storage DS4700 with twenty high-speed 146 GB disks and ten 300GB disks connected to the BladeCenter by 4Gbps Fibre Channel.
Professional services provided by GTS Maintenance & Technical Support: Hardware Maintenance, IBM-SAP Alliance

User's Benefits

❑ SAP Business Suite facilitates international expansion with multi-company support and easy localization for different markets, helping to reduce software development costs by 80 percent.

❑ Compliance with international standards is estimated to be 40 percent easier.

❑ Versatile IBM BladeCenter architecture enables to run both AIX and Windows environments–supporting all major business systems in a single energy-efficient chassis, cutting data center power costs by 20 percent.

❑ Vastly improved availability, performance, flexibility, and future scalability.

Market View

In February 1990, IBM introduced the RISC System/6000–a family of workstations that were among those using Reduced Instruction Set Computer (RISC) processor technology

Consolidating x86 on System p on Power Systems

and Advanced Interactive Executive (AIX), IBM's implementation of the UNIX operating system. It was not easy to enter the UNIX market dominated by Data General, Digital Equipment Corporation, Hewlett-Packard and Sun Microsystem with a new, unknown product. Twenty years later, three of the companies mentioned above had disappeared or were acquired. However, IBM's server family gradually but continuously evolved to become the market leader. In the last 10 years, IBM Power servers have increase their market share (based on IDC figures) from ca.17.5% to more than 40%, taking market share from other UNIX vendors.

The last year was particularly successful for IBM, fueled by the uncertainty created by Oracle acquiring Sun Microsystem. The forecast for the Power p series looks very positive; Oracle managed to annoy its customers and partners through a dramatic increase of maintenance prices[4], and HP's decision to base all non-x86 platforms on Itanium technology is somehow risky because of the relative low popularity of Itanium. A slow-down of Power's momentum is not foreseeable.

IBM Global Financing

In October 2010, IGF announced extending its North American financing offer designed to entice HP and Oracle-Sun customers to switch to IBM's new POWER line of servers program; it is now available in Europe, AP and selected growth market countries.

In Europe, the program includes interest-free payment deferral for credit-qualified Sun-Oracle and HP clients. Clients can therefore benefit from an improved business case and accelerated project break-even while avoiding payments until 2011 on their new Power Systems. For leases entered after October 5, 2010, the term of the fair-market-value lease is 39 months, with a 90 day interest-free deferral. All IBM Power Systems servers are eligible for these terms. This program has no set end-date, and can be withdrawn from the market at any time. Credit offerings are available to qualified customers only. Certain offerings may not be available in all countries.

In AP and Growth Markets, the program also includes IBM Storage, and IBM will provide clients with cash-back opportunities for specified old Oracle-Sun and HP equipment, which can be applied to new Power Systems or IBM Storage lease payments. Additionally, credit-qualified clients can benefit from a payment deferral and avoid payments until 2012 on their new Power Systems or IBM Storage equipment, in order to accommodate a seamless move to IBM systems.

Summary and Conclusions

Over the last few years, and particularly since the economy crisis in 2008, IT organizations have faced rising energy costs, difficulties in getting sufficient electrical power, and shortages of skilled IT professionals to manage their infrastructure. All these

---

[4] Oracle now charges ca 12% of the hardware cost per year for Sun hardware maintenance–double or more than in the past. Source: *Oracle hardware support plan stings Sun VARs* - 07 May 2010 SearchITChannel.com

issues and more can be answered by effective server consolidations, which, in addition, can contribute to a better environmental footprint. Organizations must begin to think differently about the data center, addressing the needs of a changed economic climate.

> *The new Express Servers, IBM Power 710, 720, 730 and 740, announced in August 2010 and available since September, offer mid-market clients the outstanding performance, energy efficiency, and other benefits of POWER7 technology at an attractive price*

The times in which a new application required a new hardware-server (and additional servers for development, testing, etc.) are over; server virtualization plays a crucial role in this development. Effective consolidation requires advanced virtualization techniques, and high availability to protect "more eggs in the basket". POWER7 ensures that the data center is equipped to address these requirements and is well positioned for future rapidly-growing businesses.

IBM's Power series minimizes deployment efforts and on-going administration with deeply integrated, pre-configured and optimized systems. It scales up, scales out and scales within to allow for flexible just-in-time procurements or upgrades. The figures speak for themselves: Power 7 offers four times the energy efficiency and twice the performance at similar prices, compared to Power6.

The new Express Servers, IBM Power 710, 720, 730 and 740, announced in August 2010 and available since September, offer mid-market clients the outstanding performance, energy efficiency, and other benefits of the POWER7 technology at an attractive price. These models are the ideal solution for SMEs to migrate from x86 to POWER servers. Whether running on Red Hat Enterprise Linux or Novell SUSE Linux Enterprise Server, more and more companies are gaining competitive advantages and optimizing their IT investment by running Linux applications on IBM Power Systems. The Power platform offers benefits of IGFs flexible financing options, significant TCO reductions, minimizes availability risks, and allows for higher server utilization through its advanced virtualization capabilities. Moving Linux and Windows workloads onto Power systems not only saves administrative and energy costs, but increases availability and improves SLAs as well.