



IBM System x Virtualization Solutions

*By Mark T. Chapman
IBM Systems and Technology Group*

Executive Overview

The IT challenge is clear: Simplify and optimize your infrastructure so your business can respond efficiently and in a cost-effective manner. IBM's vision of the new enterprise data center is an evolutionary model that redefines the economics of IT and can dramatically improve operational efficiency. It also can help reduce and control rising costs and improve provisioning speed and data center security and resiliency—at any scale. Virtualization can help clients realize this vision. Virtualizing on IBM® System x™ enterprise servers is one of the best ways to simplify your IT and lower power, cooling, floor space and management costs, while increasing efficiency.

Server virtualization makes the most of your hardware, reducing the number of servers required, enabling high x86 server utilization rates, and improving the manageability of workloads. Given the right software, virtualization can be accomplished on almost any system. However, there are distinct advantages to virtualizing on larger, scale-up servers optimized for virtualization. Having more processors, more memory, and more I/O to divvy up among virtual machines reduces the virtualization overhead percentage for more efficient utilization of resources.

In today's data center, utilization is key to the economics of server-based computing. Minimizing the number of servers translates to lower hardware acquisition costs, lower power and cooling costs, a reduction in data center space requirements, and reduced time and effort for server maintenance. Scale-up servers enable you to maximize the consolidation of workloads onto a small number of highest-performance systems, while providing the reliability you need for business-critical applications.

IBM System x enterprise servers are ideal platforms for virtualization. Our large memory footprints and the capability to scale efficiently from 4 to 16 sockets (8 to 96 cores), give you an industry-leading virtualization foundation. System x servers are built with the IBM X-Architecture™ blueprint, which melds both industry standards and innovative IBM technology to produce the fastest, most flexible, most available x86 servers in the industry. Some models even include VMware embedded in hardware to minimize deployment time and improve performance.

With virtualization, you run multiple applications and operating systems on a server, so availability becomes paramount. System x enterprise servers have a unique memory subsystem that provides multiple layers of redundancy, maximizing uptime. Our servers also have integrated reliability, redundancy, and alerting features that deliver mainframe-inspired availability.

To simplify the management of multiple virtual servers, our System x enterprise servers include proactive management capabilities, such as IBM Systems Director and IBM Virtualization Manager. They let you discover, visualize and manage both physical and virtual systems from a single console and can integrate with popular virtualization environments, including VMware, Xen¹, and Microsoft® Hyper-V™.

Ultimately, using virtualization you can reduce IT complexity and save money by better utilizing data center resources, reducing the number of physical servers needed, reducing power and cooling requirements, and avoiding unnecessary infrastructure expansion.

System x enterprise servers leverage decades of IBM development and virtualization expertise to give you a proven foundation for your virtualized environment. They're highly scalable, and designed for availability and efficiency—to help you get the most out of every IT dollar spent.

This paper will demonstrate the compelling advantages of using IBM System x enterprise servers for your virtualization needs.

¹ "Xen" as used in this paper refers to various implementation of the Xen open source virtualization engine. IBM virtualization tools support multiple instantiations of Xen; however, not all of them. Refer to the IBM virtualization solution documentation or consult an IBM representative or Business Partner for details.

X-Architecture Virtualization Technologies

A number of specific technologies incorporated into the X-Architecture blueprint, chosen from among the best the industry has to offer, as well as custom-designed by IBM, contribute to System x enterprise server superiority as a virtualization platform. The following topics briefly describe these features and their advantages to you.

Processor Virtualization Technologies

The latest Intel® Xeon® processors and AMD Opteron processors include new features specifically designed for virtualization.

Intel's **Virtualization Technology (VT)**, incorporated into their latest quad-core and six-core processors, integrates support from the processor, chipset, BIOS, and enabling software, to improve traditional software-based virtualization. Taking advantage of offloading workloads to system hardware, these integrated features enable virtualization software to provide more streamlined software stacks and "near native" performance characteristics.

Similarly, **AMD-V** virtualization technology reduces overhead and provides a hardware boost to enable faster switching between virtual machines (VMs), so you can host more VMs per server and maximize the benefits of virtualization.

Of course, every server vendor uses Xeon or Opteron processors in their servers. That's just the starting point. We went much further.

IBM eX4 Chipset

Unlike other server vendors, we didn't go with a plain vanilla chipset from Intel or ServerWorks. We chose to design our own. The fourth-generation X-Architecture eX4 chipset, used in the x3850 M2 and x3950 M2 servers, provides a number of features the competition can't offer that aid in both performance and virtualization.

One of the core features of the eX4 chipset that gives the platform its performance advantage is the snoop filter. The Xeon Coherency Protocol or "snoop" is an operation that occurs whenever a processor in an SMP system needs to update a memory address during normal operation. The snoop function comes into play when one processor gets ready to manipulate a piece of data. It asks the remaining processors in the SMP complex to verify they have not modified the same piece of data without writing back from their cache. This verification process increases traffic on the front-side bus (FSB).

The eX4 chipset contains **324MB** of EDRAM within the Northbridge chip that operates as a filter for the snoop requests. The filter copies all data as it is written to the processor cache, allowing the chipset to respond directly to the processor snoop requests. This reduces the overall traffic across the FSB and helps to improve system performance versus other architectures and can translate into more virtual machines supported with better VM performance. (Intel now offers a first-generation snoop filter in the chipset used by many server vendors; however, it contains only 64MB of RAM—**80% less** than the eX4 fourth-generation solution. The larger the snoop filter the greater the efficiency.)

In addition, the standard Intel Xeon processor includes L1, L2, and sometimes L3 caches. This wasn't enough for us, so we added a huge L4 cache, for even better performance. (See below.)

Finally, unlike our competitors—who require you to buy a monolithic all-or-nothing 8-socket server up-front—we allow you to start with a four-socket chassis. Then, if you need more sockets, memory, or I/O, you can add a second, third, and even fourth chassis, for a total of up to **16**

sockets in 16U. HP tops out at 8 Xeon processors. Dell and Sun don't even offer 8-socket Xeon processor-based servers. To go beyond 8 sockets, HP requires you to switch to an incompatible non-x86 *dual-core* processor (Itanium 2).

System x3850 M2 and x3950 M2 Expand-on-Demand

Using the mainframe-inspired, IBM-designed eX4 chipset, the **x3850 M2** and **x3950 M2**, with XpandOnDemand™ modular expansion capability, are the only switchless x86 servers from a Tier-1 vendor capable of scaling to **16** sockets and **96** cores.



The **x3850 M2** is a 4U Intel Xeon processor-based 4-socket rack server, with support for up to **256GB** of hot-add/hot-swap DDR2 memory, **256MB** of L4 cache, **7** high-speed x8 PCIe adapter slots (including **two hot-plug** slots) and **4** 2.5-inch HDDs.

If you ever find the need for more resources, simply install the **ScaleXpander Option Kit** to turn it into an **x3950 M2**, then attach another x3950 M2 chassis to the first using an XpandOnDemand cable. This simple act immediately **doubles** your processors, memory, adapter slots, HDD bays and L4 cache capacity. Still not enough? Add up to two more x3950 M2 chassis, for a whopping total of up to **16** processors/**96** cores, **1TB** of RAM, **1GB** of L4 cache, **28** I/O slots (**8** hot-plug), and **16** HDDs in only **16U** of rack space. No other x86 server on the market offers this level of flexibility, scalability, or performance.

To put it into perspective, this configuration offers at least *twice* the processor sockets, *twice* the processor cores, *twice* the memory, and *twice* the adapters (including hot-plug slots the competition doesn't offer) of other x86 servers. Some applications—virtualization and enterprise databases, for example—require the most memory-rich configuration possible. Anything less significantly limits performance. And now with six-core processors you have 50% more cores sharing the same memory and I/O as you had with four cores. That makes system expandability even more crucial. With IBM's more scalable, higher-capacity servers, you can get more of the benefits you virtualize for: more workloads, more savings, and more agility.

If you're concerned that all this "external" processor and memory connectivity will hurt performance, don't be. The innovative IBM Xcel4v™ Dynamic Server Cache, combined with the custom-designed eX4 chipset, boosts performance to the point where the x3850 M2 and x3950 M2 offer groundbreaking results in 4-, 8-, and 16-processor benchmarks—such as TPC-C™, TPC-E™ or TPC-H™; SAP or SAP SD; vConsolidate or VMmark; to mention a few—whether running Microsoft® Windows® Server or Linux®, with IBM DB2® or Microsoft SQL Server®, and so on.

Embedded Hypervisor

Selected models of the x3850 M2 ship with an **IBM 4GB USB 2.0 Flash Key** installed preloaded with **VMware ESXi 3.5**. This is an embedded version of VMware ESX 3.5, fully contained on the flash drive and requiring no disk space—not an "ESX Lite."

VMware ESXi 3.5 includes all the performance, scalability and compatibility features of ESX 3.5, including full **VMFS** support across FC SAN, iSCSI SAN, and NAS, and **4-way VSMP**. Because it runs from flash memory, it's extremely fast and ideal for diskless configurations. It also offers enhanced security, because it runs without an operating system-based console and is updated much like firmware. Licensing works the same as for "standard" ESX 3.5.

If you prefer **Microsoft Hyper-V** to VMware, we have you covered there as well. The x3850 M2 and x3950 M2 fully support Microsoft's hypervisor software installed on the system HDD.

System x3755 and HyperTransport Technology

The faster the processors, the more memory there is, and the faster the data moves among the processors and between the processors and memory, the faster the system. Using AMD HyperTransport Technology, the x3755 maximizes throughput in these critical areas.

HyperTransport technology is a high-speed, low-latency, point-to-point link, designed to increase the communication speed between circuits. It eliminates performance bottlenecks by accessing memory and L2 cache at the same speed as the processor, and by providing an integrated memory controller. This combination permits the fastest possible access to the CPUs and memory. The x3755 incorporates the latest version of HyperTransport technology (3.0), which provides an aggregate bandwidth of up to **41.6GBps**.



Due to the design of Opteron 8000 Series processors, 1- and 3-processor configurations produce less-than-optimal performance out of the box. Because of this, the competition forces you into 2- and 4-processor configurations—whether that’s what you want or not.

IBM took a different approach for the x3755. Instead of limiting your options, we chose to solve the problem. The result is the IBM **CPU Pass Thru** card. Installed in the unused second or fourth CPU/Memory card slot, respectively, it enables full performance in 1- and 3-processor configurations. This offers you the flexibility of lower-cost configurations when you don’t need more.

The x3755 also offers **Xcelerated Memory Technology™**. Because of this innovation, *all 128GB* (8 DIMMs per processor) in the x3755 run at the memory’s full **667MHz**. Competitive Opteron processor-based servers must run memory at **533MHz** when more than **4** DIMM slots are populated per processor. This provides a memory performance advantage of up to **15%** for the x3755.

High Availability

Being able to consolidate many servers into one is only an advantage if the server is able to keep running for long stretches of time, and be serviced in minimal time when necessary. The x3755, x3850 M2 and x3950 M2 offer advanced availability features to maximize uptime.

Light Path Diagnostics

Light path diagnostics enables a technician to quickly identify and locate a failed or failing system component, such as a specific fan or memory DIMM. This enables quick replacement of the component, which helps increase server uptime and lower operating costs. The front of the server has an LED indicator light to show possible component failures. If the LED indicates an error condition, pressing a button will cause a panel to pop out and drop down for easy viewing without the need to open the server cover or remove the server from the rack. The light path diagnostics panel tells the servicer which component requires attention. In addition, many components have their own identifying LEDs. For example, each of the memory modules has an LED next to the DIMM slot, as do all processors, all adapter slots, all fans, the power supply, the voltage regulator modules and the service processor. This allows the servicer to easily identify exactly which



component needs servicing. By following the “light path,” the component can be replaced quickly, and without guesswork. (**Note:** In the event of a failed DIMM or processor, the system will restart and mark the component as bad while offline, thus allowing the system to continue running, with reduced capacity, until serviced.)

Predictive Failure Analysis

Predictive Failure Analysis® (**PFA**) is designed to allow the System x server to detect impending failure of supported components *before* actual failure, and alert the administrator through the IBM

Systems Director management software. This gives you the ability to replace the failing component before it fails, resulting in increased uptime.

Most competitive systems offer PFA only on processors, memory, and HDDs (if that many). In contrast the **x3755** offers PFA support for processors, memory, HDDs, voltage regulator modules (VRMs), fans, and power supplies. The **x3850 M2/x3950 M2** servers provide PFA support for all those components, plus PCIe slots.

Memory ProteXion Technology

These days, many competitive enterprise servers offer ECC memory protection, along with “advanced ECC” protection (a technology patented by IBM that provides *multi-bit* protection per chip). In the past, these servers also offered memory mirroring as a failsafe, in case one or more DIMMs failed. But, for whatever reason, competitive systems have dropped memory mirroring, reducing your protection against server downtime. IBM continues to offer this essential feature.

On x3850 M2 and x3950 M2 servers, IBM Memory ProteXion[®] technology (also called redundant bit steering) provides *multi-chip* error protection, and works in conjunction with IBM Chipkill[™] advanced ECC technology and memory mirroring to provide three levels of IBM Active Memory[™] protection. For even higher availability, the eX4 design offers a fourth level of protection: “memory scrubbing.”

Memory scrubbing is an automatic daily test of all system memory. It detects and reports memory errors that might be developing before they cause a server outage. Memory scrubbing and Memory ProteXion work together. When a bit error is detected, memory scrubbing determines whether the error is recoverable. If the error is recoverable, Memory ProteXion technology will write the data to a new memory location; if not, scrubbing sends an alert to light path diagnostics, which then notifies IBM Systems Director to alert the system administrator. Using memory mirroring and the hot-plug memory capability, the administrator can replace failing memory with no downtime.

The combination of on-chip memory scrubbing, Chipkill protection, Memory ProteXion, and memory mirroring provides your memory data with the ultimate in protection. The competition doesn’t offer this many layers of protection.

Performance, Price-Performance, and Performance-per-Watt

As important as the preceding features are, if they don’t translate into better overall performance or better energy-efficiency there’s little advantage to virtualizing and consolidating your data center. x3755, x3850 M2 and x3950 M2 systems offer those capabilities in abundance.

Performance

Servers need high performance to power all those virtual machines and the wide range of application types within them. Due to the lower latency of the eX4 chipset, the Xcel4v Dynamic Server Cache, high-speed I/O and other features, the x3850 M2 and x3950 M2 servers continue the tradition of performance excellence produced by Intel-processor-based IBM enterprise servers. All told, these servers have amassed a total of well over 100 #1 benchmarks.

Table 1 shows a few recent ones for **quad-core** x3850 M2 and x3950 M2 servers:

<ul style="list-style-type: none">• 4p SAP SD (running VMware). The x3850 M2 is:<ul style="list-style-type: none">– 5% faster than the 4p quad-core HP DL580 G5 (Xeon)• 4p TPC-C²:<ul style="list-style-type: none">– 27% faster than the 4p quad-core HP DL580 G5	<ul style="list-style-type: none">• 8-processor SAP SD⁵. The x3950 M2 is:<ul style="list-style-type: none">– 24% faster than the 8p quad-core HP DL785 (Opteron)• 8p TPC-C⁶:<ul style="list-style-type: none">– 12% faster than the 16p dual-core Unisys
---	--

² http://ftp.software.ibm.com/eserver/benchmarks/news/newsblurb_x3850M2_DB2_tpc 101607.pdf. As of October 15, 2007.

<ul style="list-style-type: none"> • 4p vConsolidate (running VMware)³: <ul style="list-style-type: none"> – Almost 12% faster than the quad-core HP DL580 G5 • 4p VMmark (running VMware)⁴: <ul style="list-style-type: none"> – 14% faster than the 4p quad-core HP DL580 G5 (Xeon) – Nearly 8% faster than the 4p quad-core Dell R900 (Xeon) – Able to support 12% more VMmark virtual machines (running 9 tiles) than either the HP or Dell server (running only 8 tiles) 	<p>ES/000/one (Xeon)</p> <ul style="list-style-type: none"> • 8p TPC-E⁷ (8U, 2-chassis): <ul style="list-style-type: none"> – 21% faster than the 16p (37U) dual-core NEC Express5800/1320Xf (Itanium 2) • 8p vConsolidate (running VMware)⁸: <ul style="list-style-type: none"> – 118% faster than the 4p quad-core HP DL580 G5 (Xeon) • 16p SAP SD⁹ — The x3950 M2 set a world record of: <ul style="list-style-type: none"> – 10,600 SAP SD Benchmark users, also demonstrating 75% performance scalability from 8 to 16 processors • 16p TPC-E¹⁰ (16U, 4-chassis): <ul style="list-style-type: none"> – Almost 11% faster than the 32p (37U) dual-core NEC Express5800/1320Xf (Itanium 2) – Less than half the cost per transaction of the NEC
---	--

Table 1. Quad-core x3850 M2 / x3950 M2 leadership benchmark results

And the new **6-core** x3850 M2 and x3950 M2 models are just as dominating, as *Table 2* demonstrates:

<ul style="list-style-type: none"> • 4-processor TPC-C¹¹ — The x3850 M2 is: <ul style="list-style-type: none"> – 7% faster than the 4p six-core HP DL580 G5 • 4p TPC-E¹²: <ul style="list-style-type: none"> – 8% faster than the 4p six-core Dell R900 • 4p VMmark (running VMware)¹³: <ul style="list-style-type: none"> – World record 24-core result of 19.10 @ 14 Tiles 	<ul style="list-style-type: none"> • 8-processor TPC-C¹⁴. The x3950 M2 achieved a world record, demonstrating: <ul style="list-style-type: none"> – 1,200,632 tpmC — the first x86 server to break 1 million tpmC • 8p SAP SD¹⁵: <ul style="list-style-type: none"> – 9,200 SD users – 75% more than the quad-core HP ProLiant DL785 (Opteron)
---	--

Table 2. Six-core x3850 M2 / x3950 M2 leadership benchmark results

Price-Performance and Performance-per-Watt

Consolidating many servers into one via virtualization enables savings of various kinds: data center floor space, power and cooling, centralized management (simpler management, as well as travel reduction for administrators and servicers), and so on. In addition, System x servers themselves offer additional features to enhance price-performance:

The story of System x energy efficiency starts with the IBM-designed power supply. Typical industry-standard power supplies are only 60-70% efficient. This means *300-400 watts of every kilowatt consumed is expended as waste heat, rather than used productively.* By contrast, the

³ http://ftp.software.ibm.com/eserver/benchmarks/news/newsblurb_x3850M2-x3950M2_vConsolidate_082708.pdf. As of August 27, 2008.
⁴ http://ftp.software.ibm.com/eserver/benchmarks/news/newsblurb_x3850M2_VMmark_040308.pdf. As of April 3, 2008.
⁵ http://ftp.software.ibm.com/eserver/benchmarks/news/newsblurb_x3950M2-8P_sap_062508.pdf. As of June 25, 2008.
⁶ IBM System x3950 M2, eight quad-core Intel Xeon Processor X7350 2.93GHz (8 processors/32 cores/32 threads), 841,809 tpmC, \$3.46 USD / tpmC, avail. April 1, 2008. Unisys ES7000/one Enterprise Server, eight dual-core Intel Xeon Processor 7140 3.4GHz (8 procs/16 cores/32 threads), 520,467 tpmC, \$2.73 USD / tpmC, availability May 1, 2007. Unisys ES7000/one Enterprise Server, 16 dual-core Intel Xeon Processor 7041 at 3.0GHz (16 processors/32 cores/64 threads), 749,839 tpmC, \$3.33 USD / tpmC, availability June 12, 2006.
⁷ http://ftp.software.ibm.com/eserver/benchmarks/news/newsblurb_x3950M2_tpce_041108.pdf. As of April 11, 2008.
⁸ <http://www.principledtechnologies.com/clients/reports/IBM/IBMvCon0808.pdf>. As of August 27, 2008.
⁹ http://ftp.software.ibm.com/eserver/benchmarks/news/newsblurb_x3950M2-16P_sap_052008.pdf. As of May 20, 2008.
¹⁰ http://ftp.software.ibm.com/eserver/benchmarks/news/newsblurb_x3950M2_tpce_041108.pdf. As of April 11, 2008.
¹¹ http://ftp.software.ibm.com/eserver/benchmarks/news/newsblurb_x3850M2_SQL_tpcc_091508.pdf. As of September 15, 2008.
¹² http://ftp.software.ibm.com/eserver/benchmarks/news/newsblurb_x3850M2_4P_tpce_091508.pdf. As of September 15, 2008.
¹³ The x3850 M2 was configured with the Intel Xeon Processor X7460 at 2.66GHz (4 sockets/6 cores per socket/24 cores total) and 80GB of PC2-5300 DDR II memory DIMMs (twenty-four 2GB and eight 4GB memory DIMMs). As of September 18, 2008.
¹⁴ http://ftp.software.ibm.com/eserver/benchmarks/news/newsblurb_x3950M2-8P_DB2_tpcc_081908.pdf. As of August 19 2008.
¹⁵ http://ftp.software.ibm.com/eserver/benchmarks/news/newsblurb_x3950M2-8P_sap_090808.pdf. As of September 8, 2008.

x3850 M2 and x3950 M2 enterprise rack servers are as much as **91%** energy-efficient. Therefore, only **90** watts are wasted out of each kilowatt.

To fully exploit this efficiency and help you maximize computing performance per watt, IBM offers servers equipped with leading performance-per-watt quad-core processors from Intel and AMD and six-core processors from Intel. For example, in the x3755 **AMD PowerNow!** technology with optimized power management dynamically reduces processor power draw based on workload. This can result in a power savings of up to **75%** per processor at CPU idle. Older servers lack this capability, wasting power.

x3755, x3850 M2 and x3950 M2 systems come equipped with **low-power 1Gb technology DDR2** DIMMs, rather than the power-hungry Fully Buffered DIMMs used by competitive systems. eX4 systems also use **Buffer on Board** technology, which together can save up to **37%** over competitive systems that use FBDIMMs. The cost-savings can be substantial, considering that memory consumes 25-40% of data center power.

To reduce energy consumption further, **IBM Calibrated Vectored Cooling™** minimizes the amount of energy that a server’s fan system consumes. It triggers only those fans in a critical “cooling zone” of the server to switch to full speed when temperatures rise in that zone—thus cooling hotspots without requiring *all* fans to spin at full speed. This means less energy consumed, less wear on the fans, and reduced noise. In addition, **2.5” hard disk drives**, used in the x3755, x3850 M2 and x3960 M2, conserve **40%** or more of the energy used by 3.5” drives¹⁶. The optional **IBM Rear Door Heat eXchanger** can remove up to **55%** (50,000 BTUs/15 KVa) of the waste heat from a 42U rack (via sealed chilled-water lines) before the heat exits the rack, thus reducing the load on data center CRAC units and enabling increased rack density. **IBM Systems Director Active Energy Manager™** provisioning software (standard on the x3755, x3850 M2, x3950 M2 and other System x servers), enables you to monitor power usage hourly, daily, weekly or monthly, and to use the data to cap power and thermal limits at the system or rack level.

For more information about these and other technologies in the IBM Cool Blue™ portfolio, go to <http://ibm.com/systems/x/advantages/energy/overview.html>.

All of these features produce extremely power-efficient servers¹⁷ (Table 3):

<ul style="list-style-type: none"> • 4-processor vConsolidate (running VMware). The quad-core x3850 M2 consumed: <ul style="list-style-type: none"> – 13.6% less energy than the quad-core HP DL580 G5 	<ul style="list-style-type: none"> • 8-processor vConsolidate (running VMware). The quad-core x3950 M2 consumed: <ul style="list-style-type: none"> – 19.6% less energy than two 4-processor quad-core HP DL580 G5 servers
--	--

Table 3. Quad-core x3850 M2 / x3950 M2 leadership power-efficiency results

That degree of power-efficiency also greatly benefits their performance-per-watt numbers¹⁸ (Table 4):

<ul style="list-style-type: none"> • 4-processor vConsolidate (running VMware). The quad-core x3850 M2 had: <ul style="list-style-type: none"> – 29% better performance-per-watt than the quad-core HP DL580 G5 	<ul style="list-style-type: none"> • 8-processor vConsolidate (running VMware). The quad-core x3950 M2 had: <ul style="list-style-type: none"> – Almost 36% better performance-per-watt than the 4-processor quad-core HP DL580 G5
---	--

Table 4. Quad-core x3850 M2 / x3950 M2 leadership performance-per-watt results

To put it in other terms, here is how the competition stacks up watt for watt:

- IBM x3850 M2 server has a power rating of 1,440W based on 220V, in a robust configuration
- HP DL580 G5 has a power rating of 2,400W based on 220V, in a comparable configuration

¹⁶ Seagate Technologies: <http://www.seagate.com/content/pdf/whitepaper/TP-534.pdf>.

¹⁷ <http://www.principledtechnologies.com/clients/reports/IBM/IBMvCon0808.pdf>. As of August 27, 2008.

¹⁸ <http://www.principledtechnologies.com/clients/reports/IBM/IBMvCon0808.pdf>. As of August 27, 2008.

- Dell PowerEdge R900 has a power rating of 1,570W based on 220V, in a comparable configuration

This means that:

- HP DL580 G5 consumes up to 67% more power than IBM.
- Dell PowerEdge R900 consumes up to 9% more power than IBM.

With IBM's x3850 M2, you could save up to **\$790 per server, per year**, over the HP DL580 G5—an annual savings of **40%**¹⁹ in energy costs.

Virtualization Software

Virtualization, naturally, doesn't involve only hardware. One IBM tool, **IBM Virtualization Manager** allows you to manage physical and virtual machines from a single console. Combining IBM Systems Director and Virtualization Manager, you can manage VMware ESX Server, Xen, and Microsoft Virtual Server environments today. (Hyper-V support is planned²⁰.) Virtualization Manager also integrates VMware VirtualCenter and IBM Systems Director for advanced virtual machine management.

Security concerns will eclipse TCO as the driver of *recentralization* and server consolidation will spread to desktops via virtual clients. To this end, IBM's **Virtual Client Solution** allows users to enjoy all of the benefits and personal control of a stand-alone desktop—including print capabilities, USB drive support, and audio—while reducing many of the challenges related to current stand-alone desktop environments. These include limiting susceptibility to theft and viruses, ease of deployment of new users, extended downtime during a hard drive failure, or having to rebuild their preferences and settings after each client “refresh.”

The increased spread of SANs, both Fibre Channel and iSCSI, is driving the need for **storage virtualization** as well. Through consolidation and virtualization, the enterprise can achieve a simpler, more scalable, more cost-efficient information infrastructure that aligns more flexibly with emerging business goals. IBM storage virtualization offerings can help transform the economics of enterprise storage by enabling users to simplify their infrastructure, protect their data and efficiently manage information. IBM **System Storage**[™] offerings allow you to virtualize disk and tape storage. Storage utilization is a key issue for many customers. The volume of data customers have to store is growing quickly.

In order to keep control of storage costs, improved storage utilization has taken on increased importance. The IBM **System Storage SAN Volume Controller** helps improve storage utilization, makes better use of existing storage, controls growth, and is designed to improve application availability by implementing changes to storage and moving data without taking applications down. The SVC also helps simplify management, enabling greater efficiency and productivity for the storage management staff. For information protection during disasters, the SVC offers network-based replication. The SVC can easily be added to an existing SAN, and supports over 130 different disk systems, plus all major server environments and all major switch vendors. The SVC does *not* require the purchase of new storage or switches. It can be upgraded without disruption and includes brand new functions, such as Space-Efficient Virtual Disks, the SVC implementation of “thin provisioning,” which improves storage utilization and dynamic provisioning; Space-Efficient Flash Copy, the SVC implementation of “snapshots,” provides high value functionality for disk backup and test environments; and Virtual Disk Mirroring, which is designed to protect critical application data from failure of a disk system, providing high availability for critical data,

¹⁹ Using the United States Department of Energy nationwide average electricity cost of \$.094 per kilowatt hour.

²⁰ Systems Director can manage the *physical* resources of a system with Hyper-V today; recognition of Hyper-V *virtual* resources, etc., will follow later.

Additionally, the **IBM Virtualization Engine™** TS7700 can help reduce costs, including power, maintenance, operations and support staff. Through the implementation of a fully integrated tiered storage hierarchy of disk and tape, the benefits of both technologies can be leveraged to help enhance performance and provide the capacity needed for today's tape processing requirements.

Bringing together the best in infrastructure management, virtualization, and productivity software, IBM storage software utilizes the best in storage technology to answer your need for information on-demand. Offerings include the **IBM System Storage Productivity Center**, which allows storage administrators to configure multiple storage devices—such as SVC—and **Tivoli Storage Manager**, which can help reduce the risks associated with data loss while contributing to reduced complexity, cost management and compliance management for corporate and regulatory data retention and availability requirements. Additional software offerings include IBM's **Data Mobility Solutions** and other offerings that back up, restore, protect and optimize your storage infrastructure and data.

Along with IBM's existing comprehensive portfolio of storage hardware, the IBM **DS5000** storage array is a new offering designed for high performance and scalability and supports VMware and Microsoft Hyper-V. The DS5000 storage series delivers the reliability and availability attributes to further push midrange storage into the enterprise. With redundant components, automated path failover and extensive online configuration, reconfiguration and maintenance capabilities, the DS5000 series enables your data to be available 24x7x365. The DS5000 series also offers multiple replication options and a persistent cache backup to capture cache data in case of a power outage.

The DS5000 series delivers industry-leading performance and is equally adept at supporting transactional applications, such as databases and OLTP; throughput-intensive applications, such as HPC and rich media; and concurrent workloads for consolidation and virtualization. When requirements change, the DS5000 series can upgrade host interfaces, performance, capacity, cache and reconfigure it on the fly—enabling it to keep pace with your growing company.

IBM's new enterprise data center vision is a holistic approach designed to help you achieve dramatic improvements in efficiency and flexibility. This can best be achieved by implementing virtualization not just at the server and storage level, but also at the middleware layer. With **IBM WebSphere Virtual Enterprise (VE)** software, you can virtualize your Java and J2EE application environments, whether they are based on IBM WebSphere®, Apache Tomcat, BEA Weblogic, JBoss, or PHP application servers. With WebSphere VE software running on System x servers, you can achieve greater cost savings, respond even faster to changing business requirements and manage the health and performance of your environments more effectively.

In a VMware ESX Server 3.5 environment, WebSphere Virtual Enterprise delivers:

- Dynamic clusters.
- Service policy goals with autonomic request flow management (using the Autonomic Request Flow Manager component).
- The use of the Application Placement Controller (APC) to place long-running, or batch, jobs when used in conjunction with WebSphere Extended Deployment Compute Grid.
- The ability for multiple application servers, or nodes, to be supported within an individual virtual or physical machine.
- Enhanced isolation characteristics and resiliency.
- Controls so that an application within a virtual machine does not consume all resources in the virtual machine and impact other applications running in the virtual machine.
- Prevention against potential over-commitment by virtual machines of physical resources (for example, memory).
- Health management of applications: fixing the problem, or routing work around the problem.

- Reduction in virtual machine sprawl and increased management efficiencies: by increasing the utilization of virtual machines and consolidating multiple applications in virtual machines using policy-based workload management. It enables optimization goals to be met across the entire set of virtual machines in the infrastructure.

WebSphere VE provides the ability to separate applications from the physical infrastructure they run on. Workloads can then be dynamically placed and migrated across a pool of application server resources, allowing the infrastructure to dynamically adapt and respond to business needs, and allowing requests to be prioritized and intelligently routed to respond to the most critical applications and users.

IBM Virtualization Services

IBM's services and research organizations are working with more clients in the deployment of advanced enterprise technologies than any other vendor in the world. Here is a sampling of the experienced we have gained working with clients around the world:

- We had more than **700 customer engagements** in the 4th quarter of 2007 alone—just as the buzz about being “green” was taking off—not because it was in vogue, but because we have an end-to-end solution that encompasses all of what a data center needs to think about.
- IT Optimization—with over **10,000 engagements**—has been a cornerstone of our work with clients for years.
- This collaboration isn't limited to just technology or systems. We are also the clear leader in SOA engagements, working with more than **6,550 enterprises** to implement SOA in some fashion.

What does all this mean? Fundamentally, it is all about the breadth and depth of our experience, brought to bear on the emerging needs of our clients, both now and in the future. And confidence that we have been there and worked through the early stages of adopting new technologies as well as addressing the hard operational issues that face our customers today.

Here is a subset of the ways IBM can partner with you to help evolve your data center to facilitate consolidation, reduce your space requirement, and lower your energy costs:

- **IBM Implementation Services for Open Source Virtualization** helps clients streamline their server environments by using Xen open source virtualization software to create multiple virtual server instances on individual physical systems. IBM consultants analyze the client's existing server and application infrastructure elements, install and configure Xen, and provide administrative training for the client's IT staff
- **IBM Data Center and Facilities Strategy Services** can help you determine your data center facility capacity, availability and security requirements. It also can help you identify a suitable location for your facility, and help you evaluate the best way to acquire, operate and manage the data center—with a smart design that can carry your organization into the future.
- The **Data Center Energy Efficiency Assessment** measures the energy usage of your cooling, electrical and building systems, compares your energy efficiency to an industry standard and identifies opportunities to improve. It helps you create business-case financial justification to help prioritize improvements for energy savings, giving you a framework to make infrastructure decisions.
- **IBM Server Product Services** for IBM System x servers can help you shorten the time to plan, implement and deploy new servers and software, and reduce the risks associated with integration. You gain the ability to adopt new technology more quickly and improve end-user service levels.

With help from IBM, you can install, configure and exploit the capabilities of VMware ESX Server, Xen, and (soon) Hyper-V in a timely, cost-effective manner. You gain access to hands-on training

and IBM specialists who have the proven skills, experience and expertise required to help you implement VMware ESX Server in the IBM System x environment more smoothly and rapidly.

Summary

The perfect virtualization server must offer a number of critical capabilities:

- It must be fast enough to run many virtual machines concurrently
- It must offer enough memory, I/O and other resources that your VMs aren't constrained to fewer sessions than your processor can handle
- It must run those VMs as efficiently as possible, not wasting cycles on unnecessary overhead
- It must be reliable and available enough that your virtual machines stay up and running as much as possible
- It must be supported by a portfolio of virtualization software—including storage virtualization—that takes full advantage of the system capabilities
- It must be economical enough—not only in terms of server hardware costs, but also in terms of software licensing, power and cooling—that it makes economic sense to consolidate

IBM's enterprise servers offer all of these capabilities in spades:

- Consistent #1 benchmark results across a wide spectrum of benchmarks
- Double the memory, I/O, processors and processor cores of the x86 competition
- A hypervisor embedded in hardware—quick loading and ideal for diskless configurations
- A wide array of high-availability features that the competition can't match
- A broad portfolio of virtualization software
- Extremely energy-efficient hardware components plus Calibrated Vecteded Cooling and other Cool Blue technologies to help cut your power and heating bills considerably

There isn't another server in the industry that can match these capabilities. For virtualization, the only sensible answer is IBM System x enterprise servers.



For More Information

IBM System x Servers

IBM System x and BladeCenter Power Configurator

IBM Standalone Solutions Configuration Tool (SSCT)

IBM Electronic Service Agent

IBM ServerProven Program

IBM Technical Support

IBM Configuration and Options Guide

ibm.com/systems/x

ibm.com/systems/bladecenter/powerconfig

ibm.com/servers/eserver/xseries/library/configtools.html

ibm.com/support/electronic

ibm.com/servers/eserver/serverproven/compat/us

ibm.com/server/support

ibm.com/servers/eserver/xseries/cog

Legal Information

© IBM Corporation 2008

IBM Systems and Technology Group

Dept. U2SA

3039 Cornwallis Road

Research Triangle Park, NC 27709

Produced in the USA

September 2008

All rights reserved.

For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven. Telephone support may be subject to additional charges. For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

IBM, the IBM logo, the e-business logo, Active Memory, Calibrated Vectors Cooling, Chipkill, Cool Blue, DB2, Predictive Failure Analysis, System Storage, System x, Systems Director Active Energy Manager, Virtualization Engine, WebSphere, X-Architecture, Xcel4v Dynamic Server Cache, Xcelerated Memory Technology, and XpandOnDemand are trademarks of IBM Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://ibm.com/legal/copytrade.shtml>.

InfiniBand is a trademark of InfiniBand Trade Association.

Intel, Intel logo, Core, Pentium, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

Linux is a registered trademark of Linus Torvalds.

Microsoft, Hyper-V, SQL Server, and Windows are trademarks or registered trademarks of Microsoft Corporation.

TPC Benchmark, TPC-C, and TPC-E are trademarks of the Transaction Processing Performance Council.

UNIX is a registered trademark in the U.S. and/or other countries licensed exclusively through The Open Group.

Other company, product and service names may be trademarks or service marks of others.

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. IBM has not tested these products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Some machines are designed with a power management capability to provide customers with the maximum uptime possible for their systems. In extended thermal conditions, rather than shutdown completely, or fail, these machines automatically reduce the processor frequency to maintain acceptable thermal levels.

MB, GB and TB = 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will depend on considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available. When referring to variable speed CD-ROMs, CD-Rs, CD-RWs and DVDs, actual playback speed will vary and is often less than the maximum possible.