



89 Fifth Avenue, 7th Floor
New York, NY 10003
www.TheEdison.com
212.367.7400



White Paper

IBM System x and BladeCenter

Leveraging IBM eX5 Systems for Breakthrough Cost and Density Improvements in Virtualized x86 Environments



Printed in the United States of America

Copyright © 2011 Edison Group, Inc. New York. Edison Group offers no warranty either expressed or implied on the information contained herein and shall be held harmless for errors resulting from its use.

All products are trademarks of their respective owners.

First Publication: January 2011

Produced by: Craig Norris, Lead Analyst; Howard M. Cohen, Senior Analyst; Barry Cohen

Table of Contents

Executive Summary	1
Introduction	3
Objective	3
Audience	3
Contents of this Report	4
Modern Data Center Issues	5
Addressing the Issues	5
The Future: Now	5
Where the Value Is: Increased Density is Key	7
High Density Server Consolidation is Not Just for the Enterprise	7
Roadblocks to High Density VM Consolidation	7
Considerations in Server Consolidation: Memory and More Memory	8
Assessing Memory Requirements	9
How IBM Provides Optimized Solutions for x86 Server Virtualization	10
Externalized Memory	10
Comparison Example	11
Workload Optimization	13
Who Can Benefit Immediately	16
Scalability Leaves Room for Growth	17
Conclusions	18

Executive Summary

As applications, workloads, and data multiply in enterprise organizations, inefficiencies in data center operations become more pronounced, resulting in higher costs, whether directly through expenditures such as energy bills, equipment purchases, and licensing, or indirectly through greater claims on resources such as administrative staff.

Addressing the ever-growing amount of data and the number of applications in the data center by adding more and more x86 servers has resulted in uncontrolled server sprawl and its attendant costs in electrical power, real estate, administration and maintenance, and equipment. An often-cited adage regarding data center hardware management costs sums up the situation thus: “Every penny spent on hardware requires a dollar to manage.” Consolidating servers through virtualization on single physical servers has become the de facto approach to reining in server sprawl, thereby minimizing costs, simplifying deployment, and streamlining management with single-point-of-management software.

Many organizations are not fully leveraging the virtual capabilities of their servers, because they configure a median of under ten VMs per server ¹ when they could be configuring many more. In fact, the average VM density falls far short of the true capital and operational savings that virtualization can achieve, which really starts to become very significant at as little as a 30:1 server consolidation ratio, well within the reach of even a modest data center in a smaller business. Organizations also encounter constraints in their ability to scale single VMs in order to accommodate expanding transaction processing or database implementations.

The amount of addressable memory a given server can access is the single largest roadblock to higher VM densities and scalable VMs. Modern servers must be equipped with far larger memory ranges per processor than previous designs. Even now, the memory addressability limitation of many popular servers is one of the bottlenecks driving customers to purchase more systems than is needed to meet their performance requirements, which results in their running more physical servers with underutilized processors. This can be very costly. In addition, it could help explain why IT managers often believe, mistakenly, that high-density virtualization and its associated economy of scale is available only to enterprise organizations.

¹ According to a recent Aberdeen Group report, *Best-in-Class Practices for Virtualizing Microsoft Applications*, August 2010, even the best-in-class organizations in the study only consolidate at an 18:1 ratio.

IBM's approach for today's consolidated data centers is to utilize common x86 server building blocks, yet include innovations designed to further drive down costs for virtualization. With the fifth generation of enterprise X-Architecture™, eX5, IBM has introduced new server technology that can reduce the number of physical servers required for a given environment by 50–66 percent and lower licensing fees by as much as two-thirds over other equivalently priced servers on the market. This is because as many as three DL380 G7 servers (HP's two-socket, 2U, x86 server with a 192 GB maximum memory capacity) are required to accommodate what IBM's two-socket 2U x3690 X5 server, supporting up to 512 GB of memory, can accommodate in a third of the space.

The new 1U IBM MAX5 memory expansion enclosure delivers memory scalability that is unprecedented, decoupling memory from system processors, enhancing performance, and shattering the constraints inherent in the 30-year-old x86 architecture. Combined with the x3690 X5, it offers up to six times the memory of other two-socket x86 servers, and can therefore enable replacement of up to six HP DL380 G7 servers with a single x3690 X5 server. The scalable memory expansion it enables can be harnessed to significantly reduce costs by running lower-cost memory on a greater number of DIMM slots, with headroom to grow in the future.

Not only does this translate into upfront acquisition cost savings, it reduces total cost of ownership via lower software licensing fees, lower power consumption, and maintenance of fewer physical machines. Most importantly, it permits the kind of high-density VM consolidation that yields the most significant savings.

IBM eX5 servers are preconfigured and pretested to ease deployment and enable a virtualized operation to be up and running with the least amount of administrative intervention and business interruption. From the superior performance and enhanced memory subsystem that leverages functions built into the latest Intel® Xeon® 6500/7500 series processors, to mainframe-inspired reliability features that underlie the entire IBM X-Architecture blueprint, IBM designs eX5 systems to be relied upon to support multiple business-critical virtualized workloads.

IBM eX5 technology helps data centers manage growth with pay-as-you-go scalability. It lets even modest data centers in smaller businesses protect their initial investment in eX5 server technology. For instance, with certain products, they can double the number of processors and grow memory as their business needs demand.

Introduction

Objective

With approximately 60 percent of organizations deploying core business-critical workloads on x86-based systems ² and more than half those workloads running on virtual machines (VMs), questions arise as to whether that environment will really be robust enough to deliver the kind of reliability, scalability, and performance required. These are valid concerns. Assembling an Off-the-Shelf (OTS) system, even using top-of-the-line hardware and best-of-breed applications, can often fall short of the mark. Many organizations have taken this approach only to see disappointing results, which can have serious consequences.

Of course, published industry benchmarks are one way to assess the top-end capability of products and/or solutions. However, benchmark configurations are most often designed to test only the outer limits of specific metrics such as processor speed, throughput, capacity, etc. They do not accurately reflect an organization's real-world workload, where an optimal balance must be struck between price, performance, scalability, and reliability. Benchmarks are merely tools and should never be the sole means of assessing hardware and/or software for virtualization solutions.

IBM has achieved outstanding results in terms of scalability and performance in x86 enterprise configurations by applying extensive development expertise and years of engineering breakthroughs and innovation from other technologies to x86-based technology. This white paper examines the considerable rewards to be achieved in data center consolidation through virtualizing IBM x86-based servers, and best practices to consider when adopting this strategy.

Audience

Any IT manager, CIO, systems architect, etc. who is considering an economical two-socket x86-based Windows or Linux infrastructure to run large-scale IT business solutions will benefit from reading this study. Anyone currently running such an operation with less than acceptable results will also find valuable information in this white paper, which may lead them to an effective solution within a reasonable price range.

² *The Value of Memory-Dense Servers: IBM's System x MAX5 for its eX5 Server Family*, Michelle Bailey, IDC, March 2010.

Contents of this Report

This paper contains the following sections:

- **Modern Data Center Issues** — Discusses primary challenges facing data centers today, solutions that have arisen to address those challenges, x86 server consolidation and its current limitations, and IBM's solutions to consolidating x86 servers in the data center.
- **Where the Value Is: Increased Density is Key** — Discusses the widespread adoption of server consolidation via virtualization, the lack of full realization of its true potential, and the reasons underlying that lack. Continues with a discussion of considerations and best practices in consolidating servers, including the vital importance of memory and assessing memory requirements.
- **How IBM Provides Optimized Solutions for x86 Server Virtualization** — Introduces eX5, the fifth generation of IBM's enterprise X-Architecture and its ability to reduce the number of servers required for a given workload by 50-66 percent and lower licensing fees by as much as two-thirds over other comparably-priced servers available on the market.
- **Conclusion** — Presents Edison's observations and conclusions regarding the implementation and results covered in this white paper.

Modern Data Center Issues

As applications, workloads, and data have multiplied, inefficiencies have cropped up in data center operations. These have become quite severe in several areas, including server sprawl and its attendant difficulties in deployment and management, underutilization of server processor capacity, overutilization of electrical power, and memory bottlenecks that restrict performance. All these inefficiencies result in higher costs, whether directly through expenditures such as energy bills, equipment purchases, and licensing, or indirectly through greater claims on resources such as administrative staff.

Addressing the Issues

Scaling out by adding multiple Intel processor-based x86 servers running Linux or Windows is an approach widely adopted in order to facilitate application scalability, as well as reduce costs over larger, more centralized systems. This approach allows for dynamic addition of capacity simply by adding comparatively inexpensive servers to the environment.

More work finds its way to the data center every year as companies push for higher productivity and greater automation. Addressing this by adding more and more servers, data centers have found themselves with explosive or unchecked server sprawl that adds to administrative complexity, consumes ever-increasing amounts of energy for power and cooling, and requires an ever-increasing amount of floor space, miles of cabling, and hours upon hours of maintenance. Many of these servers are typically utilized at about 10 percent of capacity, an inefficiency that makes these servers ideal targets for virtualization.

The Future: Now

A growing trend to rein in server sprawl consists of data center consolidation through server virtualization, where specialized software allows multiple servers running as virtual machines to operate on a single physical server. By consolidating and virtualizing servers, data centers can increase the utilization of hardware and decrease the number of physical assets that require deployment, maintenance, and management. This approach offers the opportunity for organizations to consolidate their applications into a more efficient implementation that simplifies IT, cutting costs related to equipment, energy, physical space, and management. For these reasons, virtualization is becoming one of the central business priorities for IT managers.

However, consolidation of x86-based servers through virtualization has not been without its limitations. Memory capacity is the most critical limitation. Large numbers of VMs can be accommodated by the computational power available in a single processor on a physical host machine, especially as the number of cores increases. Nevertheless, since every workload will have its own instance of operating system and application software, each individual VM still requires the same amount of working memory that it would as a physical machine.

For this reason, modern servers destined for virtualization must be equipped with far larger memory ranges per processor than previous designs. This presents innovation opportunities for memory expansion on x86 systems. As discussed in the following section, memory represents the limiting factor, since the economy of scale in a virtualized data center becomes truly significant at around a 30:1 server consolidation, leveraging the performance capabilities of current processors.

The typical solution, for data centers to purchase additional servers in order to run increasing workloads, is costly and results in underutilization of processor capacity.

Where the Value Is: Increased Density is Key

With more than 50 percent of business server workloads now deployed on virtual machines, server virtualization is quickly reaching maturity and is the foundational platform for the data center. IDC reports that virtualization has become the default build for new server installations.³ In addition, the average size of virtualized workloads increased threefold between 2006 and 2009. Clearly, IT executives and planners are well aware of the value that can be realized by virtualization and data center consolidation.

High Density Server Consolidation is Not Just for the Enterprise

Despite this high level of awareness, the majority of businesses fall far short of seizing upon the full potential of consolidation. They have adopted the strategy and implemented the enabling technology; however, the average consolidation ratio hovers below 10 VMs per server.⁴ Data center consolidation economics significantly increase at much higher VM densities. Many IT professionals mistakenly take this to mean that considerable savings and cost reductions are within reach of only enterprise-scale organizations having massive data center capacity. This is not the case. Raising consolidation ratios to as little as 30 VMs per server can yield serious returns in terms of hardware acquisition costs, maintenance, software license fees, and energy bills. Depending on the workloads, with the proper system in place, this can be more than comfortably achieved on a single two-socket, 2U, physical IBM x86 server. For instance, IBM's latest eX5 servers are capable of supporting hundreds of VMs.

Roadblocks to High Density VM Consolidation

Obviously, if a data center can double production, yet remain well within the capacity of a given system, it could cut in half the equipment to be managed and the data center floorspace and energy required for power and cooling, while reducing the hardware points of failure and attendant service calls, along with the onsite service part inventory. So what is keeping organizations from optimizing their server consolidation ratios?

- **Fear of System Failure** — IT administrators are leery of running many critical business applications on an x86 server.

³ The Value of Memory-Dense Servers: IBM's System x MAX5 for its eX5 Server Family, March 2010. IDC

⁴ According to a recent Aberdeen Group report, *Best-in-Class Practices for Virtualizing Microsoft Applications*, August 2010, even the best-in-class organizations in the study only consolidate at an 18:1 ratio.

- **Traditional Thinking** — While many data centers do purchase multicore x86 servers with a lot of memory, these are typically being deployed to run large, memory-intensive applications, such as databases or enterprise ERP, rather than to leverage their capacity to run hundreds of VMs on a single system. The tradition of utilizing large servers for large applications and small servers for small applications prevails.
- **Memory Limitations** — Not having enough accessible memory is the main impediment to fully leveraging virtualization, one that has nothing to do with either perception or mindset. The following subsections address this vital issue.

Considerations in Server Consolidation: Memory and More Memory

The amount of memory accessible by virtual machines is the greatest impediment to increasing VM densities, according to what most IT organizations report. Industry-standard x86 servers initially architected to support single applications have become inadequate, in terms of addressable memory, to accommodate many virtualization goals. Under-provisioning memory to a VM forces the VM to page, a process where the VM stores and retrieves data from secondary storage, hindering performance.

Aware of this, virtualization software vendors, as well as x86-based hardware vendors, have taken measures to address it. Hypervisor vendors, such as Microsoft and VMware, can reclaim some unused guest memory, increasing efficient memory utilization. Hardware vendors have come out with innovations in server and processor design that have enabled systems to get the most out of the RAM allocated to a VM. For example, Intel's Extended Page Tables (EPT) can offload memory management functions from a hypervisor and can reduce the overhead that hypervisor requires.

Intel's new 6500 series and 7500 series Xeon processors deserve particular mention for their vastly increased memory support and greatly improved Reliability, Availability, and Serviceability (RAS) features. Designed to accommodate high-density virtualization, a pair of 6500 processors supports up to 512 GB of 1,066 MHz DDR3 RAM, while the 7500 in a four-socket server is capable of handling 1 TB. A raft of new RAS features in both Xeon series makes these processors ideal for servers running critical business workloads. Altogether, the 7500 and 6500 series processors deliver exceptional scalable performance and advanced reliability for the virtualized data center.

By default, servers with higher processor counts provide additional memory access, because they typically are physically larger and therefore have greater numbers of DIMM slots and higher overall memory capacities. This has led data centers to purchase expensive systems with high core counts in order to gain greater memory accessibility, which has led to underutilized processors, which is one of the main inefficiencies that server consolidation was intended to address in the first place.

Assessing Memory Requirements

To understand how this becomes a consideration in data center provisioning, take for example a modest data center that is outgrowing its available electrical power as well as its physical space. Leasing more space and purchasing more servers would involve a considerable added expense, and may not even solve the power issue. On top of that, maintenance and management of the expanded number of machines would require allocating additional technical staff members. Projected utilization reports reveal that the average server would see a processing capacity utilization of less than 10 percent.

Consolidating the servers through virtualization would overcome all these challenges, and significant savings could be realized by a 30:1 consolidation ratio, scaling to 60:1 over a two-year period. As an affordable candidate, such a data center would typically consider a two-socket physical server that supports up to 192 GB of memory. A recommended best practice for server virtualization is to allocate about 30 percent of available RAM to accommodate peak workload periods. This results in the available memory of approximately 134 GB. It is determined that the data center's application workloads profile typically run at approximately 4 GB per VM.⁵ Factoring in the overhead for system management and hypervisor software, the equation to calculate the physical host server's memory requirement would be as follows:

$$30 \text{ Virtual Machines} * 4 \text{ GB/Each} = 120 \text{ GB} + 10\text{-}15\% \text{ Overhead} = 132\text{--}140 \text{ GB memory}^6$$

The equation for anticipated growth over the next two years would be as follows:

$$60 \text{ Virtual Machines} * 4 \text{ GB/Each} = 240 \text{ GB} + 10\text{-}15\% \text{ Overhead} = 264\text{--}276 \text{ GB memory}$$

(This projected figure takes into account an alternative growth in VM size rather than number in order to accommodate a rapidly expanding database.) The typical two-socket server supporting 192 GB of memory would suffice to meet the data center's immediate goal of a 30:1 server consolidation ratio, but offers no means of accommodating any additional expansion in either the number or size of VMs. To meet the target two-year VM growth, the data center would have to purchase either a four-socket server or another two-socket server. Either option would not only involve a significant hardware expenditure, but also increase software licensing costs due to the additional sockets.

⁵ Additional memory allocations made possible by using the dynamic memory feature of VMware.

⁶ IBM has a collection of Sizing Guides for calculating the requirements for a variety of virtualization models. Contact your IBM account representative to inquire about them.

How IBM Provides Optimized Solutions for x86 Server Virtualization

IBM has long been aware of the growing trends favoring dynamic infrastructure, such as flexible, expandable servers that can readily adapt to market demands and help reduce costs, manage risk, and improve service. The vendor recognized that x86 servers were no longer mere commodities, instead they demand many of the same levels of availability, reliability, serviceability, and performance as are expected of mainframes and other advanced servers. In 1998, IBM introduced x86 servers designed using the IBM X-Architecture blueprint. This set of design specifications incorporated such x86 innovations as IBM-patented Chipkill™ memory protection, plus light path diagnostics and Predictive Failure Analysis.

With the fifth generation of enterprise X-Architecture, eX5, IBM has introduced a brand new technology that represents the culmination of four previous generations of leadership architecture in the x86 marketplace. eX5 continues to leverage the long IBM technology development heritage, delivering new server technology with superior memory addressability that can reduce the number of servers required for a given workload by 50-66 percent and lower licensing fees by as much as two-thirds.

IBM's approach to providing server solutions for consolidated x86-based data centers is to utilize common server building blocks, including innovations around Intel processor-based solutions unavailable in off-the-shelf servers. These innovations are designed to further drive down costs through maximized memory for virtualization and through simplified deployment. The System x3690 X5 and BladeCenter® HX5 are designed to eliminate the memory bottleneck faced by two-socket 2U servers.

Externalized Memory

An important innovation central to the capabilities of the eX5 architecture results in unheard of efficiency in the use of memory, eliminating the limitations of native chip-side memory for tasks where high-speed performance is vital by allowing the chip to access extended memory. The new IBM MAX5 memory expansion enclosure decouples memory from system processors for unprecedented memory scalability and optimized server performance. The independent memory scaling that it enables delivers up to six times the memory available from other two-socket x86 server vendors, and 50 percent more memory than any other Intel Xeon 6500/7500 (Nehalem EX) processor-based server. A 1U MAX5 provides 32 additional DIMM slots for each eX5 rack server. In this manner, a two-socket x3690 X5 rack server can be expanded to 64 DIMM slots (up to 1

TB total); a four-socket x3850 X5 server can be expanded to 96 DIMM slots (1.5 TB). Similarly, the 30mm-wide MAX5 blade adds 24 DIMM slots to the 30mm HX5 blade server, expanding it to 40 DIMM slots (320 GB).

This scalable memory expansion can also be harnessed to significantly reduce costs by running lower-cost memory in a greater number of DIMM slots, with headroom to grow in the future.

Compared to server designs with smaller memory ranges, eX5 can deliver significant economic benefits in terms of virtualization software costs, which are often licensed per socket. Leveraging the full economy of scale made possible by data center consolidation using MAX5 technology, an organization can reap considerable cost benefits.

Comparison Example

For the data center example described in the previous section, a single IBM two-socket x3690 X5 server, supporting up to 512 GB of memory, would much more than cover the data center's consolidation goal, which is clearly the more cost-effective solution over purchasing additional two-socket servers or going to a four-socket machine. To illustrate, Edison obtained quotes from IBM for the two-socket x3690 X5 and the closest equivalent of the latest two-socket, 2U, server from HP, the HP ProLiant DL380 G7, each including VMware and Host OS licenses. The costs are presented in the following table:

HP DL380 G7		IBM x3690 X5	
Memory Configuration	Total Cost	Memory Configuration	Total Cost
Incl. 192 GB memory (maximum available; 18 DIMMs)	\$ 33,972.08	Incl. 192 GB memory (24 x 8 GB DIMMs)	\$ 29,545.00
		Incl. 192 GB memory (12 x 16 GB DIMMs)	\$ 32,653.00
	\$ 100,000.00 (3 x HP DL380 G7)	Incl. 512 GB memory (maximum natively available; 32 x 16 GB DIMMs)	\$ 54,615.00
	\$203,832.40 (6 x HP DL380 G7)	Incl. 1,024 GB memory (with MAX5 option; 64 x 16 GB DIMMs)	\$100,103.00

As this comparison shows, the purchase cost for a server supporting 192 GB memory is slightly lower for the x3690 X5. The lower price of the first listed IBM server over the second was achieved by taking advantage of the greater number of DIMM slots and using less expensive 8 GB memory, an option not available to the maxed-out DL380 G7 server.

If more memory were required for a data center standardized on the HP server line, an additional two-socket or a higher-priced four-socket server would have to be purchased. However, the x3690 X5 can continue to scale to as much as 2.6 times that capacity. This means that three DL380 G7 servers would be required to support the equivalent number of VMs as the maximum supported by the x3690 server.

In addition, adding the MAX5 externalized memory chassis lets one x3690 X5 server scale to 5.3 times the memory capacity available to a single DL380 G7 server, which means that six of those would be required to achieve equivalent VM capacity.

Clearly, there are notable savings at the outset when comparing the x3690 X5 to the DL380 G7 on a one-to-one basis. Comparing the x3690 X5 configured with full memory capacity to the three DL380 G7 servers required to match that capacity reveals greater savings realized in the upfront purchase costs alone, as shown in the following table:⁷


Software			
VMware (Per Server / 3 years Maintenance and Support)			\$10,000
Client Operating System, e.g., Red Hat Enterprise Linux (Per Server / 3 years Maintenance and Support)			\$3,800
Hardware / Memory Requirements			
	192 GB Memory	Number of Servers Required	512 GB Memory
IBM x3690 X5 (Configured for Full Memory Capacity)	\$55,000	1	\$55,000
HP DL380 G7 (Configured for Full Memory Capacity)	\$34,000	3	\$102,000

⁷ Prices are rounded off from manufacturer's suggested retail price.

Totals (Hardware plus Software and Maintenance, 3 years)			
	192 GB Memory	Number of Servers Required	512 GB Memory
IBM x3690 X5 (Configured for Full Memory Capacity)	\$68,800	1	\$68,800
HP DL380 G7 (Configured for Full Memory Capacity)	\$47,800	3	\$143,400

Looking beyond base server price reveals still further savings in a 512 GB memory configuration: the single x3690 X5 server is incurring a third of the VMware and client operating system costs, a third of the electrical power, and a third of the systems management costs.

These factors together can make the IBM x3690 X5 an estimated 59 percent less expensive than the HP DL380 G7. IBM eX5 is the optimal x86 platform for virtualization, delivering considerable savings even at the prescribed 30 VMs, as illustrated in the following chart:



Workload Optimization

IBM's industry expertise is leveraged to deliver servers optimized for virtualized workloads in order to minimize time to value and maximize reliability. eX5 servers are preconfigured and pretested to ease deployment and ensure that a virtualized operation is up and running with the least amount of administrative intervention and business interruption.

With multiple and often business-critical workloads running on the same server, RAS becomes more essential than ever. IBM System x® servers shield an organization from IT disaster with multiple levels of protection, including:

- **IBM Predictive Failure Analysis** — Using advanced heuristic techniques and self-diagnostics, IBM Predictive Failure Analysis (PFA) helps detect when components are operating outside of normal thresholds. PFA can predict the failure of supported components, so administrators can replace them before they fail. Many vendors offer PFA on hard drives and memory; however, with eX5, IBM includes PFA for more components, such as hard drives and solid-state drives, memory, processors, power supplies, fans, PCIe slots, and voltage regulator modules.
- **Light Path Diagnostics** — This feature potentially saves hours of troubleshooting by constantly monitoring selected components within a System x server. A failure detected by PFA causes a light to illuminate on the front face of the server to alert administrators to a problem. The pop-out/drop-down light path diagnostics panel has a light identifying the failed subsystem. Within the system, an illuminated LED directs engineers or administrators specifically to the failed component, saving the servicer from having to test multiple components to find the failed (or failing) one.
- **FlexNode Automatic Node Failover** — This feature reduces downtime on a BladeCenter HX5 two-node system by ensuring that if one VM in a virtualized server system experiences a fatal error and a shutdown, the system will automatically reboot as a single VM. After reboot, the remaining system continues to operate without the failed VMs.
- **Virtual Fabric for IBM System x** — This is a fast, flexible, reliable I/O solution that helps reduce cost and complexity in both virtualized and non-virtualized environments. It decouples I/O adapters from the server hardware that uses them to enable multiple virtual Ethernet ports from a single physical port, and allows an exact number of ports needed (from two to eight per adapter) to be defined. It also enhances reliability with Intelligent Failure Monitoring to enable automatic failover between physical or virtual ports. This technology allows many VMs to efficiently share small numbers of high-bandwidth adapters. It can reduce hardware expenditures required to adequately provision networking ports, as well as the cabling and power costs to support additional hardware. When used with a 10 Gb Virtual Fabric Adapter, one adapter can act as multiple Ethernet, Fibre Channel, FCoE, and iSCSI adapters simultaneously, with bandwidth apportioned among them in 10 Mb to 100 Mb increments.
- **Redundant Power and I/O** — All eX5 servers feature redundant power to enable high availability for business-critical workloads. Additionally, all BladeCenter servers, including the HX5, feature redundant I/O connections as well.
- **Multiple layers of memory protection** — Because mission-critical workloads require mission-critical hardware, IBM provides all eX5 servers with several redundant

levels of memory protection, from memory scrubbing to IBM Chipkill™ memory correction, to IBM Memory ProteXion,™ to memory rank sparing, to mirroring. These technologies work together to minimize downtime.

- **Memory ProteXion** (redundant bit steering) technology provides multichip error protection and works in conjunction with Chipkill technology — which provides multibit protection per chip — and standard ECC protection, to provide multi-level memory correction.
- **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 so, the Memory ProteXion feature will write the data to new location; if it is not recoverable, scrubbing sends an alert to light path diagnostics, which then notifies IBM Systems Director management software to alert the administrator.
- **Memory mirroring** works much like disk mirroring. The total memory is divided into two channels. Data is written concurrently to both channels. If a DIMM fails in one of the DIMMs in the primary channel, it is instantly disabled and the mirrored (backup) memory in the other channel becomes active (primary) until the failing DIMM is replaced.
- **Memory rank sparing** works somewhat like mirroring, except instead of having an entire memory channel reserved, two ranks per memory card are configured as spares. In the event of a memory failure on the same card, one of the spare ranks is automatically used in place of the failing one. (Rank sparing and mirroring are mutually exclusive.)

As the latest offerings in the System x and BladeCenter technology roadmap, servers with eX5 technology are a major component in a dynamic infrastructure, adding significant new breakthrough innovations to previous System x advantages, such as:

- **Superior Chip Technology** — eX5 servers incorporate the latest advancement in Intel's Xeon processors, the 6500 or 7500 series. The 6500 series is a new class of processors for two-socket servers that delivers up to 58 percent better performance than the comparable Xeon 5600 series, which is what competitors such as HP use in their high-end two-socket, 2U, servers. This latest generation of processors overcomes the memory/performance tradeoff that slows the earlier processors' performance when three DIMMs per channel are configured.

Additionally, the new chip technology vastly reduces the impact on memory performance for nearly all chip-based RAS features, such as mirroring, DIMM sparing, and rank sparing. In previous generations, when a RAS feature is enabled, it typically reduces the performance of the memory subsystem by reducing the

available bandwidth and/or the available capacity. The available memory bandwidth is reduced by up to two-thirds (due to one or two memory channels being unavailable when the particular memory RAS feature is enabled). By contrast, when a memory RAS feature is enabled on the 6500 or 7500 series processors with which the IBM eX5 servers are equipped, memory bandwidth and capacity are not affected. Enabling all those processor features gives the eX5 servers mission-critical RAS.

- **Risk Management** — eX5 servers are designed with mainframe-inspired reliability, offering Predictive Failure Analysis, light path diagnostics, redundant subsystems, and memory protection not ordinarily available on commodity x86 servers, in order to maximize server uptime. HP two-socket servers are comparatively limited in their memory protection capability because of the limitations of the Intel 5600 chip series used.
- **Improved Serviceability** — eX5 servers offer VMControl, a single tool to manage both physical and virtual servers, monitoring hardware and automating smart movement of workloads ahead of potential hardware failure. VMControl reins in the potential for “virtual server sprawl” in an efficient high-density VM environment.

Who Can Benefit Immediately

Because eX5 servers offer so much more memory capacity without necessitating the over-provisioning of processors, customers benefit from acquisition onward with lower-cost server purchase prices and lower per-socket software fees. They can realize greater VM density using only a small portion of their overall system capacity. They can implement a fully redundant infrastructure, deploying two highly reliable servers rather than many less reliable ones, minimizing points of failure and potential downtime.

Besides bringing decades of experience and expertise to bear on x86 server technologies, IBM brings it to bear on implementation as well. Workload optimization, placing applications on systems architecture or using system options that provide the best processing services for a given application at the required service levels, is the lens through which IBM views its installations. This focus on particular workload characteristics and service level requirements plays a big role in driving down operational costs for any organization’s server consolidation initiative, and facilitates a rapid return on investment.

Some broad categories of organizations that could gain immediate substantial benefit from data center consolidation using eX5 solutions are:

- Organizations with scores or hundreds of servers all running with underutilized CPU capacity. This very common situation is encountered either in data centers that have never adopted server consolidation through virtualization, or in data centers

that have virtualized but have found it necessary to over-provision processors in order to obtain sufficient inline DIMMs to meet memory requirements.

- Organizations running scores or hundreds of servers without fault tolerance. This is a hazardous scenario, which risks the loss of critical systems, data, and even business accounts. The eX5 server line's simplicity of deployment offers the opportunity for such a data center to set up multiple VMs on affordable, high-memory servers, with redundancy that leverages the ease with which virtualized servers can be copied and moved.
- Organizations pushing their servers' capacity to the maximum limits of memory. The eX5 line of servers with the MAX5 memory drawer option eliminates memory as a performance bottleneck.

Scalability Leaves Room for Growth

Despite the considerable benefits of processor improvements in respect to virtualization, many data centers now report that the major constraint to increasing VM densities lies in the ability to add memory to a system.

When the time arrives to add capacity, System x and BladeCenter eX5 servers help data centers manage growth by enabling them to scale their servers in tandem with their expanding workloads. This pay-as-you-grow scalability lets them double the sockets or significantly increase available memory, simply by adding additional server chassis. Additionally, the balanced design of eX5 technology enables them to leverage the power of the latest Intel processors without sacrificing response times. The new MAX5 memory chassis can dramatically increase memory scalability without requiring the purchase of additional servers. The X-Architecture advantage allows organizations to tackle tough workloads and gives them room to grow at their own pace.

Even modest data centers in smaller businesses can benefit from the investment protection offered by eX5 server technology. They can keep up as memory intensive applications are deployed, and can build on their existing server investment as their business grows.

Conclusions

Over the past few years, server virtualization has become the de facto standard configuration practice for many organizations. Even when only a single operating system and application instance is installed on a server, that instance is being installed on a hypervisor. For most organizations, server virtualization is being used to consolidate x86 workloads from many underutilized servers to fewer systems. Unfortunately, in the majority of cases, these same organizations are not taking true advantage of the technology's capabilities in order to fully realize the considerable value virtualization offers.

Even in "best-in-class" organizations, customers are restrained by the memory bottleneck and are consolidating servers at only an 18:1 ratio; while the vast majority of organizations are consolidating at below a 10:1 ratio at the most. eX5 servers from IBM can support exponentially greater numbers of virtual machines because of the advanced memory technology in the new MAX5 memory enclosure. For example, IBM System x eX5 series servers running VMware can provide near disaster-proof support for dozens or even hundreds of virtual machines, while providing more than adequate memory and processor headroom for load balancing and performance tuning. The cost reductions this kind of VM density makes possible are far too substantial to be overlooked.

Edison strongly advises organizations to recognize that with current technology they can dramatically improve their data center server consolidation and gain considerable cost savings as a result. By standardizing on IBM System x and BladeCenter eX5 servers running server virtualization at levels of 30 VMs per server or more, organizations will begin to see the true potential and rewards of virtualization.