IDC
Analyze the Future

WHITE PAPER

IBM System X4: Delivering High Value Through Scale Up

Sponsored by: IBM

Kenneth Cayton Jed Scaramella

January 2008

EXECUTIVE SUMMARY

The scalability debate rages on: Is it more efficient to add resources — processors, memory, and I/O — to a single server or simply to add a second, third, or fourth server to accommodate increases in demand for server applications? Of course, the answer very much depends upon the thread-level parallelism of the hosted application workload. With the exception of databases and HPC programs, most traditional x86 applications were not designed for multiprocessor platforms, and so the path of least resistance has been the one application/one server, scale-out model. Servers grew smaller and users adopted sophisticated tools to handle space and management issues. And then, businesses started running out of electrical power in their datacenters.

Power shortages — and related cooling issues — proved to be the forcing function driving users to reexamine their deployment practices. While it's possible to maintain hundreds of smaller servers to deliver enterprise computing needs, it is no longer practical. IDC estimates that for every $1.00 spent on new servers today, an additional $0.50 is spent on power and cooling. Unless drastic changes are made, IDC expects this ratio to increase to $0.70 or more by 2011. IDC's Worldwide Quarterly Server Forecaster shows an expected growth rate for scale-up architecture systems with four sockets or more that is higher than that expected for smaller systems over the period from 2006 to 2011. This change reflects a leveling between the two usage models over the past three years.

IDC has identified two major trends that are helping to drive this direction: virtualization and consolidation. Both of these trends tend to drive fewer, yet larger, systems to take on consolidated workloads that formerly ran on multiple, individual servers; consolidation of workloads can result in 4–12 virtual servers per physical server. The adoption of virtualization technology for x86 servers began largely in test and development environments, helping companies reduce the number of single-core systems required for deploying new applications. In today's multicore processor world where every dual-processor system is an 8-way symmetrical multiprocessing (SMP) machine, almost every server will require this technology simply to use the resources at hand. The natural consolidation that takes place addresses numerous issues, including the following:

- **System utilization.** Systems that are underutilized beg the question of how to better take advantage of this investment in server hardware. The obvious but not simple answer is to put more work on the system.

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

- **Space and energy.** Power and cooling are most understood when cost or physical constraints affect an end user's ability to function normally; for example, when costs exceed budget or power requirements exceed available resources. Floor space can be a huge issue when an expanding IT footprint results in building out datacenter or IT space.

- **Operating expenses.** Since the mass adoption of x86 servers, system management costs have grown significantly. The increasing operating costs have resulted in the diversion of resources and capital from initiatives that are aimed at driving innovation and increasing the value of IT.

Using a scale-up environment has the potential to increase system utilization, lower power and cooling requirements, decrease space requirements, and simplify operations, thus reducing system management requirements.

SITUATION OVERVIEW

In definitional terms, scale-up (vertical) and scale-out (horizontal) performance and capacity increases are relative to application usage models and not specifically tied to the number of processors in a system. Most four-socket or larger systems would fall under the scale-up category, but they could also be used in a scale-out fashion for some large application installations. Scale-up architecture is typically used in reference to systems with four or more sockets running a single instance of an operating system.

The proliferation of the PC in the 1980s gave rise to the client/server computing model and drove the rapid growth of the LAN server market, ushering in a trend toward more scale-out server computing. Then, the emergence of a commercial Internet and Web-based computing in the 1990s skewed even more growth toward horizontal scaling. Initially, the majority of scale-out computing centered around infrastructure and presentation logic roles in multitiered workloads that developed over the same time period.

In the early 2000s, many believed that scale-out computing would totally dominate the server market over time for all but a handful of transaction processing applications. This outlook has proven to be overly optimistic as traditional scale-up systems, including IBM's System z, have leveraged architectural benefits and added support for open operating system platforms (Linux) to remain fairly equal to scale-out systems over the past several years. Additionally, IBM is noting a resurgence in its x86-based scale-up platform as customers gain experience with virtualization technologies. The additional processing and memory resources of the larger servers take advantage of dynamic load balancing tools in virtualization solutions to support a larger number of simultaneously peaking applications. IT customers are able to do this with fewer larger servers rather than with more smaller systems, reducing the complexity and cost of managing IT resources.

The Impact of Virtualization and Consolidation

Virtualization in computing is a broad term that refers to the abstraction of computer resources. It includes making a single physical resource (such as a server, an operating system, an application, or a storage device) appear to function as multiple logical resources. Alternatively, virtualization can include making multiple physical resources (such as storage devices or servers) appear as a single logical resource.

Server virtualization has a long history beginning with the mainframes of the 1960s that supported logical, software-defined partitions. For well over a decade, high-end RISC-based servers with Unix operating systems have utilized virtualization technologies to isolate and manage workloads. The capability for virtualization was lacking in early standard x86-based servers. As noted previously, customer demand for such capability was minimal because the use of one application per server was standard practice. Virtualization has changed rapidly over the past six to seven years and is now broadly accepted in x86-based servers. Virtualization is a key enabler for consolidating multiple applications onto a single system, increasing utilization, and reducing total cost of ownership (TCO).

Figure 1 shows the expected growth of virtualization, comparing physical and logical server shipments from 2005 to 2010. (Note: There is a logical server for each physical server and added logical servers for each virtual server. The result is that when virtualization is present, there will be more logical servers than physical servers.)

FIGURE 1

IDC's 2006 Server Virtualization Model: Physical Server Unit Shipments and Logical Server Unit Shipments, 2005–2010

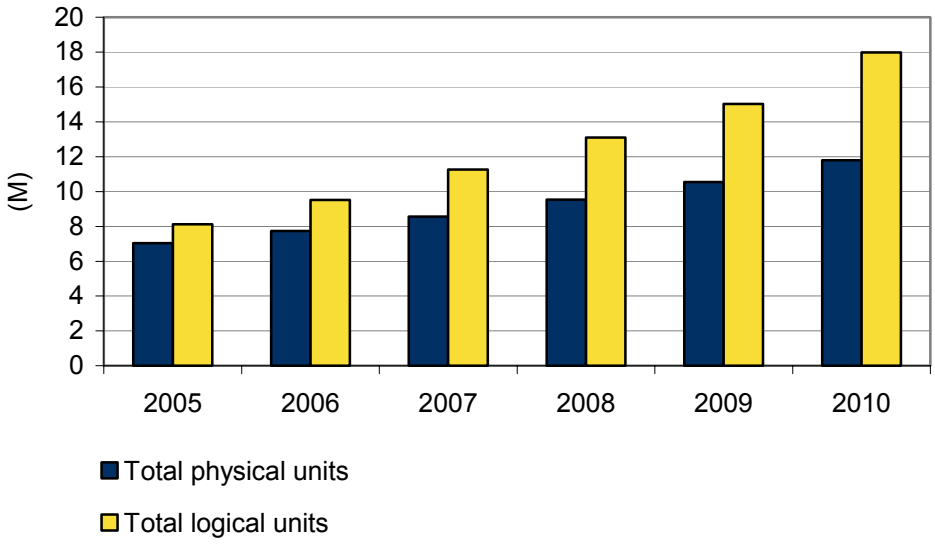

Source: IDC's 2006 Server Virtualization Model

IDC's forecast for the server market estimates that total server market unit shipments will increase at a compound annual growth rate (CAGR) of 5.8% from 2006 to 2011. The total x86 server market is expected to increase slightly faster at a CAGR of 6.3%. Looking at the segment of systems with four or more sockets — servers supporting a scale-up model of computing — IDC sees even faster growth. IDC expects the total market for four or more sockets to increase at a CAGR of 9.3% from 2006 to 2011. x86 servers with four or more sockets are expected to increase at a CAGR of 10.7% from 2006 to 2011.

Figure 2 shows the expected growth rates of x86-based servers with one, two, and four or more sockets over the forecast horizon. One-socket servers are expected to increase at a CAGR of 3.2% from 2006 to 2011. Two-socket servers are expected to increase at a CAGR of 6.8%, and servers with four or more sockets are expected to increase at a CAGR of 10.7%. The takeaway is that scale-up architecture x86-based servers with four or more sockets are expected to grow almost two times as fast as the overall market.

FIGURE 2

Forecast of System Unit Growth by Socket Capability, 2006–2011

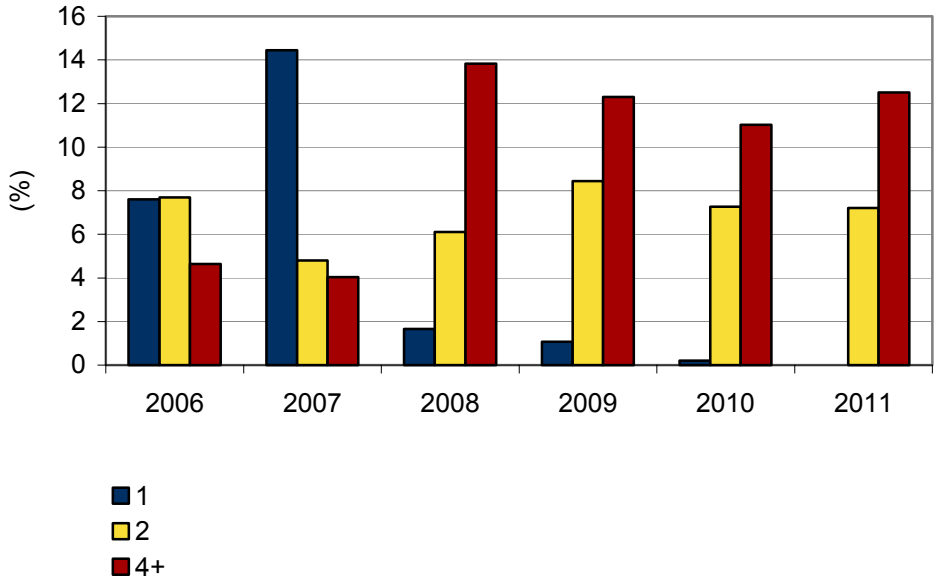

Source: IDC's Worldwide Quarterly Server Forecaster

Trends in Power and Cooling

In addition to trends in virtualization and consolidation, several infrastructure and environmental aspects are driving end users toward scale-up architectures. Some of these datacenter environmental characteristics include power and cooling, floor space, and a simplified environment such as management, cabling, security, updates, and patches.

The number of servers installed has doubled over the past six years. The increased number of systems compounded by servers that draw more power has created power and cooling issues in a high number of datacenters. IDC estimates that for every $1.00 spent on new servers, an additional $0.50 is spent on power and cooling. Unless drastic improvements are made to server systems, IDC expects this ratio to increase to $0.70 or more in the next five years.

The operating expenses associated with power and cooling have become a critical cost factor for companies — and power/cooling concerns are now at the top of the list for IT managers, talking about the top challenges they face in the datacenter. Importantly, end users have used consolidation to reduce the number of systems in operation and to lower their power and cooling requirements and associated costs. Other benefits from consolidation include simplified management, reduced number of cables, simplified security with fewer systems, and easier patch and update management. These benefits lower TCO by controlling both acquisition costs and operating costs.

The Benefits of Multicore, Multisocket Systems

Server processors have moved from single-core to dual-core and now to quad-core. The increased number of cores per processor has effectively created multiprocessing systems on a chip. This offers a boost in performance and an increase in scalability in each system for applications capable of using this increased capability. The result is more scaling capabilities across servers of all socket counts. The impact on software vendors is the need for and opportunity to create applications that can leverage these additional capabilities.

What added value do four-socket and larger systems offer today? Systems with four or more sockets are typically designed to handle larger and more complex critical applications with better reliability, high-availability, and scalability features than systems with fewer sockets. End users find these capabilities important whether they are running large mission-critical applications or multiple applications consolidated on a single platform. Bigger servers can more easily carry the added cost of sophisticated management hardware, redundant components, and hot-swap electrical isolation packaging. IDC's end-user research is finding that customers are consistent in placing comprehensive reliability, availability, and serviceability (RAS) capabilities high on their requirements lists when selecting new systems.

The process of IT transformation continues to accelerate as IT managers inventory the systems they have inherited and decide how to deploy servers for future projects. Usually, large companies and very large enterprises have large datacenters with "n-tier" designs that include multiple computing tiers for Web serving, application serving, and database serving. This is, in itself, a shorthand for a range of deployments that are optimized to support specific sets of workloads, such as:

- Business processing (BP), including online transaction processing (OLTP), enterprise resource planning (ERP), customer relationship management (CRM), and business intelligence
- Decision support (data analysis workloads running against database engines)
- Collaborative (email and groupware)
- IT infrastructure (support for network protocols and file/print)
- Web infrastructure (Web serving, proxy, and cache)
- Technical (scientific/engineering and high-performance computing)

Generally, there is a large inventory of aging SMP servers that may present the most pressing need for replacement or refresh, following many years of upgrades and added capacity. A large number of midrange enterprise and high-end servers were installed seven-plus years ago, during a period when scalable servers saw strong sales, given the trends to prepare for Y2K. Now, IT managers are evaluating how to replace their capacity — through scale-up or scale-out computer deployment or a combination of both types of deployments.

IBM VALUE STORY

Ten years ago, IBM invested in the research and development of a new scalable x86 platform. The company added a core-logic chipset to its X-Architecture initiative conceived to enhance x86 servers with superior power, scalability, control, and service characteristics, calling it Enterprise X-Architecture (EXA). IBM made this investment to offer systems with higher-end functionality than was currently available in that market. That investment paid off with a series of successful products.

Today, IBM has a new generation of scale-up or "pay-as-you-grow" servers using fourth-generation eX4 technology (see Figure 3 for an illustration of the eX4 architecture). By developing this platform, IBM is leveraging its expertise in high-end systems, migrating technology developed in its mainframe and high-end RISC lines to the x86 platform. The X-Architecture provides high performance, scalability, and mainframe-like reliability in a scalable x86-based platform.

FIGURE 3

IBM eX4 Architecture

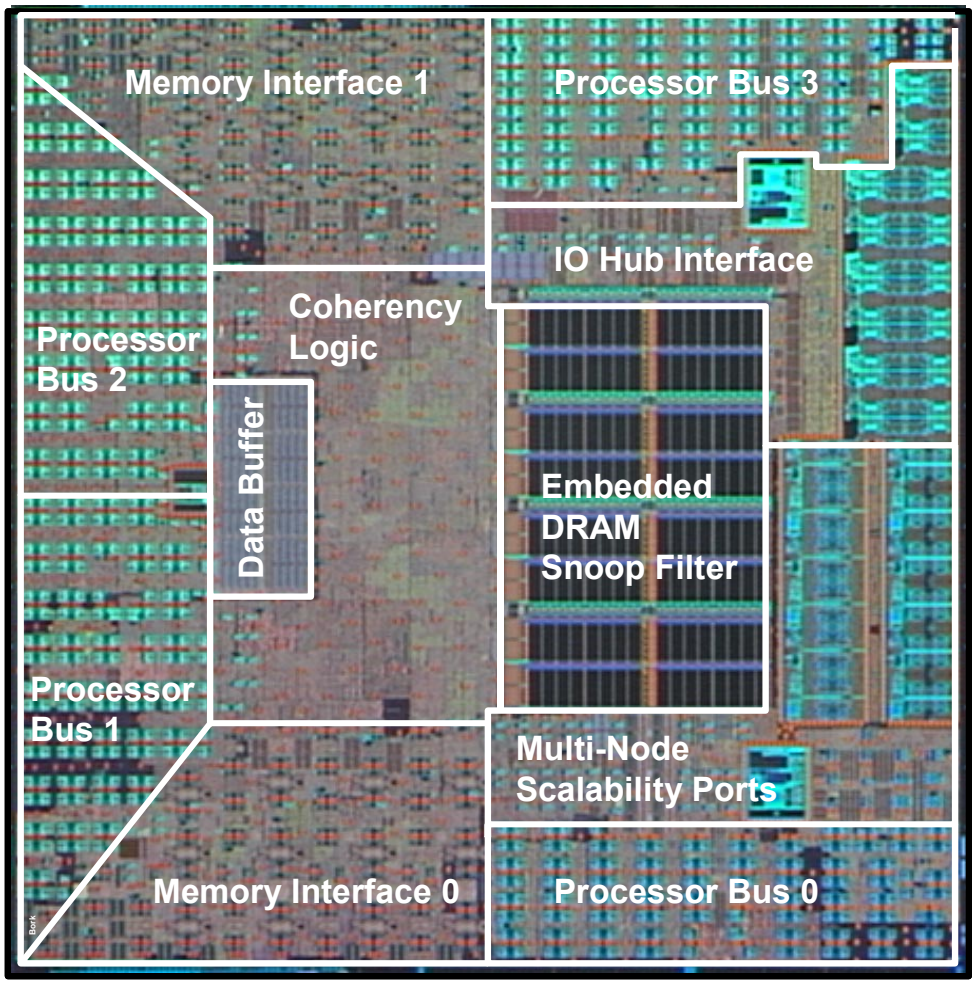

Source: IBM, 2008

Processors: eX4 Chipset

IBM's X-Architecture strategy is based on the idea that the performance capabilities of larger IBM systems can be applied to industry-standard servers, resulting in increased power and availability. IBM's Enterprise X-Architecture is designed to meet IT's business needs for higher levels of scalability, flexibility, availability, performance, and operational efficiencies.

The core of the architecture is the IBM eX4 chipset. The next-generation eX4 chipset supports systems from four sockets up to 32 sockets utilizing the Intel Xeon Series 7300 quad-core and Xeon Series 7200 dual-core processors. A maximum configuration of 64 cores is supported using quad-core processors with a fully populated 16-socket system. There are no x86 OS platforms available today supporting more than 64 cores.

Memory

The memory subsystem is designed with three levels of memory protection for increased reliability. The DDR2 memory is currently lower in power and cost than the FBD memory. IBM uses one Advanced Buffer eXecution (ABX) chip across four DIMM slots to drive the registered DIMM technology, while its competitors rely on fully buffered DIMM to support an equivalent amount of slots. Each four-socket chassis supports 32 DIMM slots and up to 256GB of memory or a maximum of 1TB of memory in a four-chassis, 16-socket system.

I/O Subsystem and Modular Architecture

The system design is modular, using four-socket building blocks to expand capacity, as needed. This creates a pay-as-you-grow environment by providing a system that can expand from one module/chassis with four sockets to four modules with 16 sockets. The I/O subsystem is optimized to create a balanced system allowing almost linear scaling when adding four-socket modules.

Virtualization

Virtualization requirements are key design elements built into the eX4 Architecture. The large memory subsystem is a direct result of this thinking. Due to the complexity, certain higher-end workloads are difficult to operate on virtual servers. The eX4 Architecture is capable of overcoming this obstacle because of the large memory footprint. Additionally, there is an optional embedded hypervisor capability — a USB port — inside the chassis to simplify loading and operating hypervisors.

Systems

The foundation system using the eX4 Architecture is the IBM System x3850 M2 enterprise server. The "building block" of the systems is a four-socket rack-optimized server with a 4U rating. The x3850 M2 utilizes either the dual-core or quad-core Intel Xeon Series 7200 and 7300 with up to 2.93GHz (quad-core) and 2/4 processors per node/module (chassis). The memory capacity, 256GB max PC2-5300 DDR II, is double that of previous generations (32 DIMMs). Seven PCI-Express x8 high-performance I/O expansion slots and two support hot-swap capabilities are also included. The ScaleXpander Option kit offers customers the flexibility to scale from four sockets to 16 sockets.

IBM Services

IBM is a global corporation with a full complement of service and support solutions available across all major regions. This holistic scale is useful for all prospective customers and potentially more valuable for those end users looking for more outside help, which may be in the form of business planning, IT planning, process optimization, operations outsourcing, or system implementation. Some end users will look for and need only a little outside assistance, while others may need substantially more — up to complete outsourcing of their IT operations. Having this global capability is an advantage for IBM when end users are looking for more than just hardware and software.

Another aspect of IBM's global world is that its worldwide distribution network is tied into its worldwide support. This is valuable to companies located in different geographies around the world, but it may be even more important to multinational companies looking for products, support, and services for multiple locations in different geographies.

CHALLENGES AND OPPORTUNITIES

Challenges

With the new generation of the X-Architecture, IBM is delivering tangible benefits to customers. IDC believes that IBM will face certain obstacles in successfully bringing the eX4 product to market. The primary challenge is that in the x86 server market, scale-out architecture has tremendous momentum, which is driving customer perceptions. Many companies have adopted scale-out architectures as their de facto standard over scale-up architectures and have adjusted their IT skill sets accordingly.

Additionally, x86 is perceived as being ideal for front-end and midtier nonbusiness-critical applications, yet not well-suited for business-critical workloads that demand high levels of availability and business resiliency. For this reason, IBM will need to educate the market on the eX4 Architecture's RAS benefits and its cost-effective attributes. The message being delivered to the market by the majority of vendors is that a four-socket system is the largest x86 system required for any modern environment. IBM needs to promote its scale-up solution as a viable alternative, supplemented with an explanation of why a larger server is superior for database hosting, enterprise applications, and large-scale server consolidation engagements.

Opportunities

IDC believes IBM can overcome many of these market challenges by leveraging its experience in the datacenter and high-end IT environments. IBM offers a broad portfolio of IT solutions and a history in scale-up computing, which can be applied here. On a tactical level, IBM can draw from many customer success stories and case studies. On a strategic level, IBM must do more to show the benefits of running enterprise workloads on a scalable, secure, and highly available platform. If that is done, IBM's continuing efforts focused on the scale-up space should be well-received by companies aiming to reduce the complexity of managing their IT infrastructures.

Based on the system improvements in the X-Architecture, IDC also believes the fourth-generation architecture is well-suited for customers' ongoing virtualization and server consolidation initiatives. The increased adoption of virtualization among companies is driving the acquisition of larger systems, as well as the requirement for higher RAS capabilities to protect the applications running on those systems. This particular customer demand is a good fit for IBM's X-Architecture systems.

CONCLUSION

Scale-up technology is alive and thriving. The market for scale-up architecture servers not only is healthy but also is projected to outgrow the general server market over the next five years. The server market is very dynamic and continues to change. This change has occurred more rapidly in recent years and shows no sign of slowing down for the foreseeable future. New hardware capability is constantly entering the market, followed by new, more demanding applications to leverage this new available resource.

For both old and new applications, IT customers will need to evaluate which workloads should continue to run on scalable servers and which are ready to go into scale-out deployments, such as server farms, bladed server enclosures, clusters, and grids. IDC believes that, while many scale-out configurations are being deployed as part of the process of IT transformation, scalable servers will continue to play important roles in the datacenter as foundations, or platforms, for housing the most demanding mission-critical workloads and providing high degrees of RAS and high availability for those workloads. In addition, many will look on their consolidated workloads in the same fashion, requiring a much higher level of RAS capability than typically found in scale-out architecture systems.