



## **Scale-Up and Scale-Out Architectures — IBM Provides Choice with the xSeries**

Analyst: David Reine

### **Management Summary**

Decisions in the consumer marketplace are not based on one single characteristic of the target product. Overall capability is often the most important factor. However, many times a product may be ideal for one environment, but not at all suitable in another. The use, or application, that you have in mind can make a highly-rated product very unsuitable for you.

Take automobiles, for example. The highest rated sports coupe would be of little use to you, if you needed to transport a family with three children and a dog. Speed and style are irrelevant when capacity is required. A minivan, on the other hand, should not even be considered, if the purpose was to find an economical commuting vehicle to get one person to work. The most important aspect of your search might be to find a dealer who not only can provide a wide variety of car styles to choose from, but can also offer a reputable service department as well.

Similarly, a data center's search for a commodity server must take into consideration the application set that the IT staff has in mind for it. Is the purpose to run an enterprise email or Internet server? Is the enterprise looking for a compute engine for its online transaction processing needs? Is the application CPU-intensive or I/O-intensive? The IT staff must first identify what characteristics it needs from its servers and then identify a vendor who can satisfy the various requirements of the many applications that the enterprise may need to execute. Secondly, they must assure themselves that this vendor can provide not only remedial maintenance, but also any professional services that the commodity environment may require. By definition, a commodity server will support any application written for *Windows* or *Linux*. In fact, with virtualization, Linux can run on every server from an x86 compatible, to a RISC-based UNIX machine, to a mainframe. Today, the enterprise needs a partner who can identify what each separate environment requires and be able to support it.

Because of the ubiquitous nature of Windows and Linux, many application providers deliver solutions for the 32/64-bit x86 environment first. This, however, encompasses two separate and distinct architectures: the Intel *Xeon* and the AMD *Opteron*, each with its own advantages. Ideally, the enterprise will find a partner who not only can support both, but can also provide additional value, in terms of a server-surround environment. With its *xSeries*, IBM can do just that, not to mention its capability to support Linux applications on its *pSeries* UNIX servers and its *zSeries* mainframe as well. To see if IBM has an x86 server that is right for you, please read on.

### **IN THIS ISSUE**

- **Evolving Data Center Requirements..... 2**
- **Server Architectures ..... 2**
- **Choosing the Right xSeries Product .... 3**
- **Conclusion ..... 3**

## Evolving Data Center Requirements

Similar goals exist in every enterprise data center: remove complexity, simplify the environment in order to lower the total cost of ownership (TCO) to the enterprise, and improve all service level agreements. There are many ways, however, to accomplish these goals. Recently, IT staffs have chosen the twin paths of server consolidation and operating system virtualization, to deliver a server environment best suited for the enterprise's mix of application workloads.

With consolidation, data centers can remove multiple, tired, 32-bit x86 architectures from their server mix and replace them with the latest 64-bit technology, either by migrating the applications to a 64-bit environment or retaining the ability to execute their 32-bit *Windows* applications through the 32-bit compatibility mode available in both the AMD *Opteron* and Intel *Xeon MP* microprocessors. In addition, there is a steady migration of applications off older RISC processors and onto these new 64-bit x86 servers running *Linux*. Consolidation also fosters the implementation of shared storage, through Storage Area Networks (SANs). Originally implemented with Fibre Channel (FC) technology, some data centers are implementing their SANs with iSCSI, to take advantage of existing TCP/IP infrastructure and to simplify further the data center architecture.<sup>1</sup>

Since these servers are much faster than those purchased in the past, there is an opportunity to consolidate many applications onto a single server. Partitioning a server, whether through virtualization by middleware - such as VMWare's *ESX Server* or Microsoft's *Virtual Server 2005*, or by hardware<sup>2</sup> - enables the IT staff to make more effective use of each processor, while still allowing applications to run in isolation, even on different operating systems. Unfortunately, not all applications conform to a single

processing environment; some are compute-intensive, while others are I/O intensive. Fortunately, today, you can choose from a variety of commodity hardware and software technologies that have been designed to provide high performance for each of these environments.

Intel provides the industry with a variety of x86 solutions, including 32-bit *Xeon* and 32/64-bit *Xeon MP*. AMD provides the x86 server community with a single-core, 32/64-bit *Opteron*, and dual-core *Opteron* as well. A wide range of vendors take these engines and develop a variety of solutions for both scale-up and scale-out environments in 1U and 2U rack configurations, blade servers, and deskside towers. There are other 64-bit architectures, notably, IBM's *POWER*, Sun's *SPARC*, and Intel's *Itanium 2*. This bulletin will examine only the x86 environments.

## Server Architectures

There is a variety of system environments to go along with the multiple processor architectures. The two most prevalent are *scale-up*, for Symmetrical Multiprocessing (SMP) systems<sup>3</sup>, and *scale-out*, for high-performance computing applications<sup>4</sup>. There are technical differences between the two types; however, the most significant variable between the two may be cost: the cost of acquisition and the cost to manage and maintain. It is more expensive to acquire a single 16- or 32-way system than multiple mono- or dual-processor nodes in a rack-mounted cluster; however, it is more costly to manage and maintain many smaller servers. An additional penalty could be felt if a single multi-CPU SMP system fails without a recovery solution. Losing a scale-up system might put the enterprise out of business, while losing a scale-out node might result in shifting the workload to other nodes, or continue to operate with reduced

<sup>1</sup> See **The Clipper Group Explorer** dated April 13, 2005, entitled *iSCSI SANs - Panacea or Placebo?*, which is available at <http://www.clipper.com/research/TCG2005020.pdf>.

<sup>2</sup> Hard partitions are found in some scale-up systems.

<sup>3</sup> With many processors working in a shared environment under a single instance of an operating system (unless the server has been partitioned to run multiple instances).

<sup>4</sup> Based on many similar processing processors, each operating independently; each with a single instance of an operating system (unless a processor has been micro-partitioned to run multiple instances).

performance.

Furthermore, there is a significant amount of ambiguity in the definition of a scale-up/scale-out node. The industry agrees with the fact that a mono- or dual-processor system conforms to a scale-out mode, while a node with more than eight CPUs clearly is scale-up. Between four and eight CPUs, however, there is much debate over how the server should be classified. Generally, the definition will remain within the eye of the beholder, with multiple 4-way or 8-way systems being referred to as a scale-out architecture of scale-up nodes.

The *server-surround* environment is equally important: memory management, I/O performance, and peripheral capability. Product selection is often based upon these components, when commodity servers are being reviewed because this is where the real differentiation lies. (See Exhibit 1, to the right.)

### **Scale-Up x86 Systems**

Scale-up systems with multi-CPU's in rack-mounted nodes, contain anywhere from 8 to 32 processors. They are configured with a single copy of the operating system, unless the systems have been partitioned. Traditionally, the majority of scale-up systems usually run under some form of UNIX on a 64-bit RISC processor, with IBM's *AIX* on *POWER* (e.g., *p595*), Sun's *Solaris* on *SPARC* (e.g., *E25K*), and HP's *HP-UX* on *Itanium 2* (e.g., *Integrity Superdome*) being classic examples. The x86 architecture, however, is now a major player in this environment with a single copy of ESX Server or Virtual Server 2005 overcoming many of the concerns about scalability and performance, by virtualizing system memory into multiple executable partitions, enabling better utilization of system resources.

### **Scale-Out x86 Systems**

Rack-mounted scale-out systems typically consist of multiple 1- or 2U nodes. These may be installed in a rack, or as blades, with mono- or dual-, or even quad-CPU configurations. One major difference with the scale-out mode is that the data center will

### **Exhibit 1 – The Challenge in Selecting the *Right* Environment**

There is a tremendous amount of confusion in the IT world regarding the selection of a scale-up or scale-out architecture. 4-Way nodes are becoming more prevalent in the server mix, bestowing a scale-up image onto a scale-out world. The advent of dual-core CPUs increases this confusion – potentially changing a 4-way server into an 8-way system (depending on the application's ability to use hyperthreading). The answer usually lies within the application set that the enterprise plans for each specific node. In addition, consolidation will often marry compute-intensive programs with I/O-intensive applications in order to maximize processor utilization.

It is imperative that the enterprise uses all of the resources at its disposal to ensure that the IT staff understands the implications of both environments.

need to deploy a separate copy of virtualization software on each node, potentially increasing software costs and the accompanying redundant infrastructure and management. This may be offset by the decreased server acquisition costs. However, because of the multiple nodes installed, a scale-out architecture provides increased availability and resiliency. If you lose a node, you usually only lose a minor portion of the virtual infrastructure, as opposed to losing the entire 16-way or 32-way node.

### **Which to Choose?**

There is a wide variety of x86 solutions from which to choose, based on Intel Xeon and AMD Opteron processors, in both rack-mounted enclosures and blades. It is important to understand the functional requirements of your application set in order to determine the best choice of microprocessor. Will the enterprise applications benefit from the reduced memory latency available from an Opteron solution or perhaps from a larger cache available with Xeon MP? How will virtualization enable your data center to

achieve better utilization of the server environment? Can the application set benefit from the multi-threaded, dual-core capability of Opteron? Can the enterprise benefit from a 32-bit only environment with Xeon? Does your enterprise, in fact, have the systems expertise necessary to make such an evaluation?

### Choosing the Right xSeries Product

Dual/quad processor x86 server nodes are available from many vendors. Acquisition choice should be focused on the value-add that any one of these vendors can provide to your enterprise. From a hardware technology standpoint, IBM has adapted their mainframe experience into the third generation of the *Enterprise X-Architecture* technology for the xSeries offerings. Major value-add features of the xSeries are listed in Exhibit 2, to the right.

#### IBM's Scale-Up Systems

Let's now look at IBM's scale-up offerings, not only for servers, but also for the server-surround and support environment. IBM provides an x86 scale-up architecture within its *xSeries*, with two 4U rack-mounted systems: the 16-way *xSeries 455* running Windows or Linux on 64-bit Itanium 2 and a 32-way *xSeries 445* running those same operating systems in a 32/64-bit Xeon MP environment<sup>5</sup>. Some of the advantages of this scale-up environment are dynamic system resource management, lower software costs, and lower infrastructure costs.

IBM has also introduced the 3U *xSeries 366*, sitting on the edge between scale-up and scale-out with 4-way, clusterable nodes. The x366 is based upon the latest Xeon MP, the *EM64T*, with 32-bit compatibility for legacy applications.

The principal x86 scale-up systems from IBM are a 4-way x365/366 Xeon MP, a 16-way x455 Itanium 2, and a 32-way x445 Xeon MP. The 365 and 366 are 3U nodes, while the 445 and 455 are 4U systems. The

<sup>5</sup> The x445 is also configurable with up to four low-cost Xeon processors in order to improve the price/performance ratio in less demanding compute environments.

#### Exhibit 2 – xSeries Features

- Hot swap redundant power/fans
- Drop down *Light Path Diagnostics* panel to expedite hardware repairs and dramatically reduce service time
- EXA chipset for Intel memory management
- Memory Subsystem with 8 DIMMs
  - Larger memory space for 64-bit addressing and to take advantage of less expensive 1GB DIMMs instead of 2GB DIMMs
  - Online spare memory configuration leaves 6 DIMMs for active memory
  - High-performance DDR2 memory
  - ECC memory
  - Memory Mirroring
  - Chipkill memory on 1GB DIMMs and higher for increased availability and higher reliability
- Calibrated Vector Cooling
- Support for both PCI-X and PCI-Express to position for 4Gbps support
- Flexible Disk choices: 2.5" & 3.5" SCSI, SATA

x366 is the latest entry in the IBM Intel family, utilizing the Xeon EM64T for applications that require 64-bit addressing while retaining compatibility with 32-bit applications. The x366 is based upon the third generation of IBM's *Enterprise X-Architecture* to enable the 366 to operate in mission-critical environments with the superior availability and reliability. Its No.1 rating in the TPC-H 300GB v2 benchmark typifies its outstanding performance.

Designed for consolidation, database, and ERP environments, the x445 provides outstanding scalability with up to 32 Xeon Processors MP configurable in up to four 8-way nodes. The x455 is configurable to 16 ways using Intel's Itanium 2 architecture, and is optimized for 64-bit scale-up solutions.

#### IBM's Scale-Out Nodes

IBM also offers a wide variety of scale-

out nodes ranging from the 1U *Pentium 4*-based *x306* to the *x336* (1U) and the *x346* (2U) 2-way, 32-bit Xeon CPUs. In addition, IBM has just announced the availability of the dual-core *Opteron* in the *x326* (1U), to take advantage of AMD's multi-threaded architecture and integrated memory controller. All of these nodes are configurable within the IBM *eServer Cluster 1350*.

In addition, IBM also provides a complete set of x86 blades for the *eServer BladeCenter*. These include the HS20, a 2-way 32-bit Xeon blade, the HS40, a 4-way blade based upon the Xeon MP processor, and the LS20, a 2-way blade based upon the dual-core Opteron CPU. With a BladeCenter architecture, the enterprise can eliminate some of the excess infrastructure that is common to a scale-out environment.

### **IBM's Server-Surround**

Any evaluation of a proposed consolidation or virtualization project needs to extend beyond the scope of simply the server. The IT staff needs to evaluate the management, networking, and storage capabilities of the proposed solution. They also need to review the services available from the vendors, professional as well as remedial.

*IBM Director* is their workgroup management solution. IBM Director's tools provide customers with flexible capabilities to realize maximum system availability and lower IT costs. With IBM Director, the IT staff can view and track the hardware configuration of local and remote systems in detail and monitor the usage and performance of critical components, such as processors, disks, and memory.

IBM offers a complete line of storage arrays<sup>6</sup>, with the *DS300/400*<sup>7</sup> at the entry level for iSCSI and FC networking solutions, and the *DS4000/6000* available for mid-sized

storage needs.<sup>8</sup> The *DS4800* is designed for 4Gbps I/O, to support the most demanding environments.

In terms of support, IBM also provides a complete range of professional services in order to assist your data center staff in defining the correct solution for your enterprise. They also provide the remedial services that the enterprise may need to maintain the network at the highest standards.

### **Conclusion**

Clearly, there is no single answer to the simplification of an enterprise data center. It depends mainly on the scope of your requirements and complexity of the enterprise infrastructure. The key is the ability to adapt to changing business needs with a flexible solution. Because of its very broad product line, **IBM can approach each enterprise's solution open-mindedly.** They have all of the solutions in their portfolio: 2-way nodes to 32-way nodes; Intel Xeon and Itanium, AMD Opteron, and IBM POWER; scale-up racks, scale-out clusters, and blades. They do not need to tailor the problem to fit a limited product offering, as could be the case with some vendors.

Perhaps more significant, are the services that IBM has to offer to the enterprise. From assistance in defining the proper environment, to selection of the right infrastructure, IBM has professional services that can assist your IT staff in reaching the best solution. Further, IBM can also provide a complete server surround environment so that any remedial care can be provided by a single source, with no finger pointing.

If you require an x86 solution for your enterprise application set, you need to consider IBM's xSeries for your enterprise consolidation, and virtualization projects.



<sup>6</sup> See **The Clipper Group Navigator** dated December 19, 2003, entitled *Putting Your Data in the Right Place - Using Value to Determine the Right IBM Storage*, available at <http://www.clipper.com/research/TCG2003070.pdf>.

<sup>7</sup> See **The Clipper Group Navigator** dated October 12, 2004, entitled *IBM Introduces SAN Storage for the Entry SMB - A Storage Strategy Appears*, available at <http://www.clipper.com/research/TCG2004084.pdf>.

<sup>8</sup> See **The Clipper Group Navigator** dated October 12, 2004, entitled *IBM TotalStorage DS6000 - Upper Class Storage for Middle Class Budgets*, available at <http://www.clipper.com/research/TCG2004085.pdf>.

### ***About The Clipper Group, Inc.***

***The Clipper Group, Inc.***, is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at [www.clipper.com](http://www.clipper.com).***

### ***About the Author***

***David Reine*** is Director, Enterprise Systems for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at [dave.reine@clipper.com](mailto:dave.reine@clipper.com) or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)***

### ***Regarding Trademarks and Service Marks***

**The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager,** and “*clipper.com*” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

### ***Disclosure***

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

### ***Regarding the Information in this Issue***

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.