A new generation of computing power and flexibility for industry-standard servers
August 2001

Introducing IBM Enterprise X-Architecture Technology

by Mark T. Chapman
IBM Server Group

Executive Summary

In the 1960s, IBM introduced the flexibility of the System/360 and the monolithic mainframe world was changed forever. In the 1980s, IBM introduced the IBM PC—which used off-the-shelf parts—and the concept of a personal computer was changed forever. Now IBM unveils modular, scalable server technology, and the world of industry-standard servers will never be the same.

IBM @server X-Architecture™ technology is a blueprint for extending the benefits of advanced mainframe technologies to our Intel® processor-based servers. These benefits are in the areas of availability, scalability, systems management, service and support. IBM has been delivering on the promise of the X-Architecture model since 1998 with such innovative technologies as Active™ PCI, C2T Interconnect™ cabling, Chipkill™ memory, Predictive Failure Analysis®, Light Path Diagnostics™ and IBM Director Software Rejuvenation, to name a few.

Today, IBM continues to build on the X-Architecture blueprint with Enterprise X-Architecture technologies. They yield revolutionary advances in the I/O, memory and performance of xSeries servers. This peerless new server design creates a flexible "pay as you grow" approach to buying high-end 32-bit and 64-bit xSeries systems. The results: systems that can be scaled quickly, easily and inexpensively.

Enterprise X-Architecture technology enables[1] the following capabilities:

- XpandOnDemand™ scalability
- System partitioning
- PCI-X I/O subsystem
  - Active™ PCI-X
- Remote I/O
- Active Memory™
  - 256GB memory capacity
  - High-speed (DDR) memory
  - Memory ProteXion™
  - Chipkill memory
  - Memory mirroring
  - Hot-add/hot-swap memory
- 400MHz front side bus (FSB)
- XceL4™ Server Accelerator Cache
- Active Diagnostics™

If you are looking for all of these abilities in an industry-standard server today, they are available exclusively from IBM.

These features deliver application flexibility, innovative technology and new tools for managing e-business. They bring to industry-standard servers the kinds of capabilities formerly available only to users of mainframes and other high-end systems. Combined with existing X-Architecture technologies, these innovations result in unprecedented "economies of scalability," unmatched flexibility and new levels of server availability and performance.

This paper explains not only what all of these new technologies are and how they work, but also how the application of this new technology can help you better meet today's challenging business and IT needs.

---

[1] Systems may implement a subset of these capabilities. For example, an entry-level server might forgo SMP Expansion Ports or limit the maximum available memory to less than that supported by the chipset.

Table of Contents

Enhancing Industry-Standard Servers

Over the years, IT professionals have faced many issues regarding the industry-standard server environment, including the need for high availability, scalability, performance, overall reliability and operational costs. In many cases, the low-cost advantage of these servers versus larger systems has outweighed inherent system or platform limitations. But there has always been that need to make tradeoffs. These decisions have become even more difficult as server requirements have increased due to 24-hour e-commerce operations and other needs.

Where once Intel architecture server usage was limited to providing only simple file-and-print services, we have seen its role expand into providing business-critical resources for e-business applications, CRM, ERP, business intelligence, supply chain management, collaboration and other traditional "big iron" arenas. Just as server hardware has had to advance tremendously over the years, the underlying core logic of servers must advance as well to overcome the limitations of previous designs.

It's not enough that the processors get faster and the hard disk drives get larger. All aspects of the system must work together efficiently to generate the kind of performance that today's connected business demands. Yet all the speed in the world doesn't help if your servers are down for maintenance. And how do you solve the problem of server proliferation? Is there a cost-effective way to consolidate the hundreds (or thousands) of single-processor, 2-way and 4-way servers you have scattered throughout your enterprise?

Introducing Enterprise X-Architecture Technology

IBM is addressing these business needs through Enterprise X-Architecture technology. The IBM X-Architecture strategy[2] embodies the belief that the capabilities of larger IBM systems can be applied to industry-standard servers, resulting in increased power and availability. Enterprise X-Architecture capabilities, building on the IBM X-Architecture blueprint, pave the way for unprecedented scalability, flexibility, availability, performance and operational efficiencies in deploying server resources to meet dynamic e-business needs.

Much of the Enterprise X-Architecture offering is delivered through IBM-developed core logic. IBM has more proven product technology and expertise in designing core logic than anyone else in the industry, dating back decades. The IBM XA-32™ and XA-64™ families of chipsets for 32-bit and 64-bit industry-standard servers contain advanced core logic, which is the heart of a computer system. Core logic determines how the various parts of a system (microprocessors, system cache, main memory, I/O, etc.) interact. These new chipsets bring to the next generation of industry-standard servers key advantages, including modular system nodes, system partitioning, high-performance clustering and high-speed *remote* PCI-X I/O support.

Specifically, IBM is producing three classes of chipsets: an entry-level chipset for 32-bit servers, an enhanced 32-bit chipset (both are part of the XA-32 family) and an enhanced 64-bit chipset (the first of the XA-64 family of chipsets). Each delivers different levels of functionality. The basic chipset consists of two chips: a combination *memory and remote I/O* controller and a *PCI-X to remote I/O bridge* controller. The enhanced chipsets add two other chips: a combination *scalability and L4 cache* controller and a *synchronous memory interface buffer*. (See *Figure 1* for illustrations of the entry-level and enhanced chipsets. The entry-level chipset is delineated by a dotted line.)

---

[2] See the white paper entitled "IBM X-Architecture Technology" at http://**ibm.com**/eserver/xseries for more information. From the xSeries home page, select **Library** for links to the different types of documentation available.

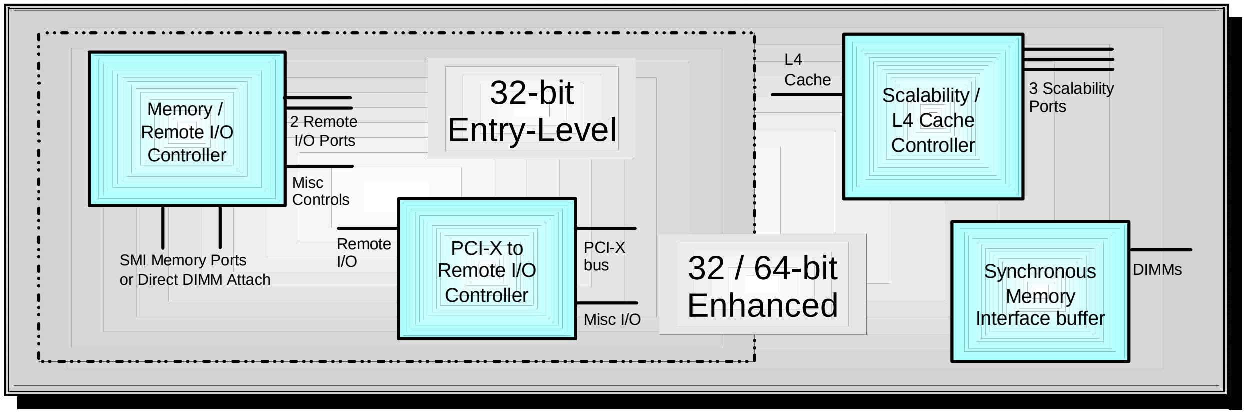

Figure 1. IBM *XA-32/XA-64 chipset components*

The Enterprise X-Architecture paradigm sets IBM @server xSeries servers apart in the industry while maintaining the advantages of compliance with industry standards for processors, memory, I/O, storage and software. It also establishes a revolutionary new economic model for servers through its flexible modular design: *No longer do you have to buy as much server up front as you can afford in order to ensure future growth capacity.*

IBM Enterprise X-Architecture technology truly excels in the multisystem environment, where it allows you to exercise "pay as you grow" scalability. The benefits in areas such as server consolidation and resource allocation are demonstrated by support for system partitioning and easy high-speed clustering.

The Enterprise X-Architecture design is the culmination of decades of experience with mainframe and midrange platform designs, yet it can be packaged in a cost-effective, *industry-standard* design, unlike earlier attempts to scale servers beyond 8-way. It offers leadership capabilities for current 32-bit systems and for the emerging 64-bit server platform.

The new 1.6GHz-and-higher Intel Xeon™ MP (codenamed "Foster") and the second-generation Intel Itanium™ ("McKinley") processors require new chipsets—existing chipsets *cannot* be used with the new processors. This means that server vendors must either develop their own chipsets or wait for a chipset specialist to develop new chipsets. If the server vendor waits for someone else to develop the chipsets it can delay the development of their servers until a chipset is at least ready for testing. On the other hand, if vendors decide to build their own chipsets it requires a massive outlay of cash and manpower, as well as a great depth of experience in designing core logic.

Because IBM had the experience and resources we developed our own 32-bit and 64-bit chipsets, and designed them in parallel. This gave us a head start on other server vendors. (Intel is even using IBM Enterprise X-Architecture servers as the validation platform for its new Xeon MP and Itanium (McKinley) processors to verify that the new processors can be deployed effectively in enterprise-class systems.) It also allowed us to implement the features described in this paper and previously unknown in the industry-standard market, but familiar to users of high-end IBM systems.

Eighty percent of the IBM chipset design is common to both 32-bit (Xeon MP) and 64-bit (Itanium) platforms. Being able to share most of a proven design means not only a shorter time to market but also higher reliability of the chipset. Instead of having to start from scratch developing and testing a 64-bit chipset, we were able to utilize designs that had already been perfected in the 32-bit implementation. This commonality between chipsets also creates a logical migration path from IA-32 based servers to IA-64 (also called Itanium Processor Family) systems when you're ready to begin moving up. (*Figure 2* illustrates the new standard set by IBM.)

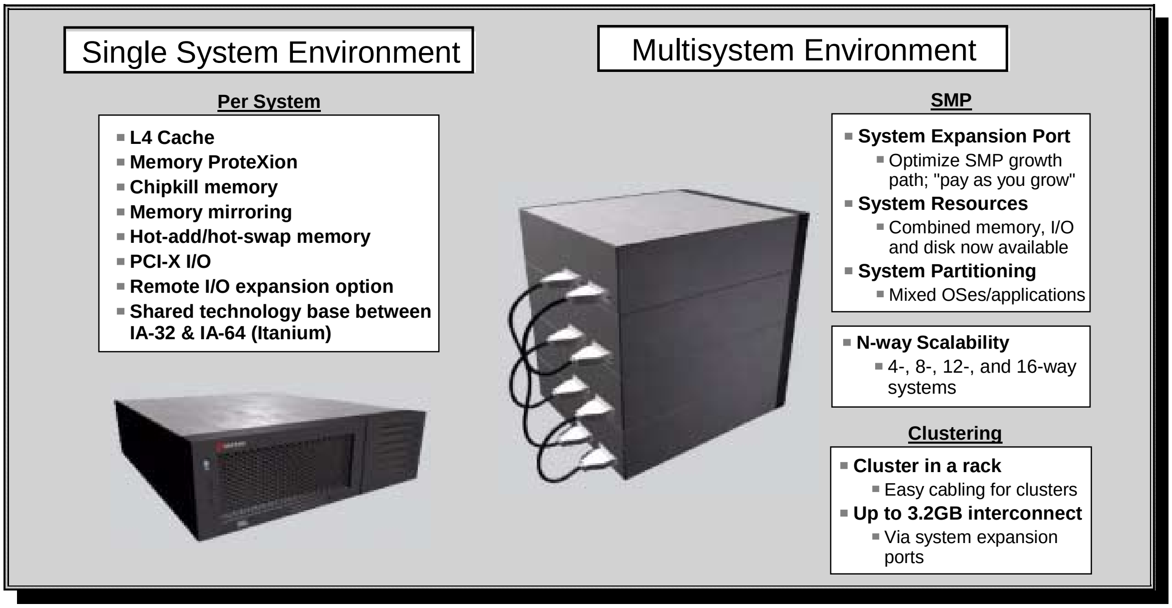

Figure 2. *System capabilities enabled by Enterprise X-Architecture technologies*

In addition to being packed full of new features, the XA-32/XA-64 chipsets themselves are marvels of design. IBM used its patented copper chip and Silicon-on-Insulator manufacturing processes to create chips that are faster, use less power and produce less heat (something your server doesn't need more of) than conventional chipsets. (Less internal heat means greater reliability of system components and less heat vented externally, which results in lower ambient cooling costs. Lower power requirements for chips result in lower server power needs and therefore lower electricity bills.)

The goal of IBM Enterprise X-Architecture concept is to blend the advantages of IBM X-Architecture technology leadership with next-generation Intel microprocessors to create the most potent combination ever brought to industry-standard computing. The resulting Enterprise X-Architecture features extend the X-Architecture blueprint with the following new capabilities, explained in the following sections:

- XpandOnDemand scalability
- System Partitioning
- PCI-X I/O subsystem
- Remote I/O
- Active Memory
- XceL4 System Accelerator Cache
- Active Diagnostics

These and other X-Architecture capabilities contribute to the IBM Project eLiza goal of creating self-managing servers. For more on Project eLiza, see the *How Enterprise X-Architecture Technologies Fit into Project eLiza* topic at the end of this paper.

XpandOnDemand Scalability: Pay as you grow

One of the challenges of optimizing server performance is to provide memory and I/O subsystems that allow new processor architectures to realize their performance potential. Traditional industry-standard server designs begin to encounter performance bottlenecks beyond 4-way scaling, due to processor and I/O bus congestion and inefficient memory utilization. Enterprise X-Architecture design provides support for advanced I/O and memory architectures and includes a high-speed, shared-cache architecture. New levels of scalability for industry-standard servers are achieved with the Enterprise X-Architecture platform using enhanced, high-performance symmetrical multiprocessing (SMP) building blocks that allow effective scalability beyond 4-way

SMP. These technologies provide scalability from 4-way to 8-way to 12-way—and even to 16-way systems—using 4-way "scalable enterprise nodes."

A scalable enterprise node contains processors, memory, I/O support, storage and other devices and operates stand-alone like other computers. Each node may run an operating system different from the other nodes, or if desired multiple nodes can be assigned to one OS image via system partitioning. Nodes are attached to one another through dedicated high-speed interconnections, called *SMP Expansion Ports*, sharing resources for unmatched performance. This gives you the adaptability to run several nodes as either a single large "complex" or as two or more smaller units—and even rearrange the configurations later as needed. We call this innovation *XpandOnDemand scalability* and it gives your business the capacity to grow as never before, with unparalleled flexibility.

You can start out simply and inexpensively with one 4-way SMP node. If it turns out that this is all the computing power you need for that server, you have saved yourself a lot of money by not buying an expensive 8-way or larger system "just in case."

If you eventually need more processing power, simply add a second 4-way node to create an 8-way server using simple cabling. (If creating an 8-way server this way doesn't provide you with enough slots and bays, you can also plug in external remote I/O expansion units, described later, and remote disk storage units, such as the IBM EXP300.)

Later, as your needs increase, three 4-way nodes can be configured as a 12-way system. Finally, four 4-way nodes can be combined into one 16-way server. (*Figure 3a* shows a stylized view of a 16-way server formed this way. In practice, the nodes would be stacked in a rack. *Figure 3b* illustrates two systems connected via the SMP Expansion Ports. Only one cable is required; a second cable provides redundancy and load balancing while doubling the throughput. The third port is used for a four-node configuration.) If at some point your requirements change and you no longer need an 8- or 12- or 16-way server, the nodes can be separated again into stand-alone 4-way systems. This modular design gives you the flexibility to decide how *you* want to grow your IT infrastructure. Unlike previous industry-standard designs, you are not limited to an all-or-nothing choice up front.

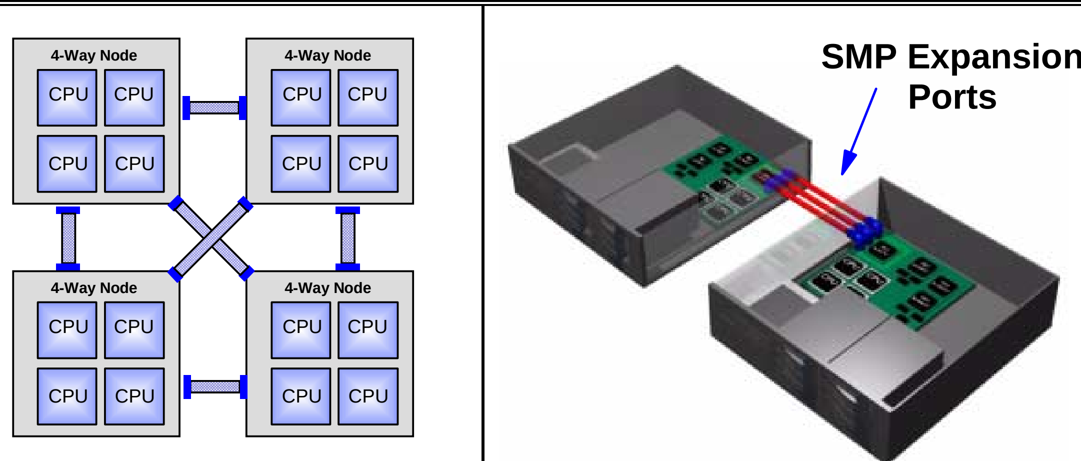

Four 4-way nodes (16-way system)

Figures 3a and 3b. *Scalability via SMP Expansion Ports*

Consider the following: Let's say you are absolutely certain that you will never need more than an 8-way server, so you buy a traditional 8-way system partially populated with four processors and 8GB of RAM installed. The next year you determine that your new server has reached its limits in the original configuration, so you elect to upgrade it to its full potential by doubling the number of processors and the amount of memory. You quickly notice that despite adding four more processors and another 8GB of memory, the performance is not tremendously better than before. This is because the traditional single memory path becomes a bottleneck. This causes contention among the processors, which are fighting for access to the memory, and wastes valuable processor cycles. In addition, your server still has only one system chipset, one front-side bus, and one access path for all PCI adapters, network traffic, etc. In other words, no more system bandwidth than before.

With the Enterprise X-Architecture approach, adding a second 4-way node likewise doubles the number of processors and the amount of system memory, but it does so without the bottlenecks. The Enterprise X-Architecture design provides access between all processors and all memory, regardless of their respective nodes, reducing contention. Not only that, but you also double the number of chipsets, front-side buses, PCI buses, and other resources to help share the data traffic burden. And this is with only two nodes. More nodes translates into more system bandwidth. The contention and resource problems increase in a 12- or 16-way conventional SMP system (assuming you had the cash up front to buy a server that large), but the Enterprise X-Architecture design eases the overload as easily for a 16-way server as for an 8-way system, just by plugging in new 4-way nodes. *And you didn't need to have all the money in your budget at once—the purchases are spread out over multiple budget periods.*

We refer to this design that lets you increase not only processors but all required resources as a balanced system architecture.

In each configuration (4- to 16-way), all of the interconnected nodes can be operated by one instance (copy) of an operating system and applications as a single partition or the nodes can be divided into multiple partitions. (See the next topic for more on system partitioning.)

Likewise, connecting a *cluster* of servers with failover is as easy as connecting two, three or four 4-way nodes—using the *same* SMP Expansion Port cabling among nodes. For scalable clustering, no complicated Ethernet setup is required to create a high-speed interconnection between nodes, because it already exists via the SMP Expansion Ports.

> ***Note:*** It can be confusing to refer to "clustering," because there are different types. One method uses Microsoft® Cluster Service (MSCS) to provide failover and load-balancing capabilities. This type of clustering doesn't benefit significantly from using high-speed communication links, because the internode traffic for MSCS clustering tends to be light and not especially performance-sensitive. A standard 10Mbps Ethernet configuration would be equally suitable for this purpose.
>
> The kind of clustering that would benefit most from high-speed connections involves a distributed application like Microsoft SQL Server or SAP (which can be distributed among nodes without MSCS). Each partition runs its own instance of the operating system and applications, and the multiple instances cooperate to perform the same task. With this clustering method, high-performance communication among the separate application instances is critical, and the SMP Expansion Ports have the potential for much better performance than even Gigabit Ethernet. (A lot depends on how efficiently the program is designed.) On a Microsoft Windows® system the software that allows this to happen is WinSock Direct (WSD). With another OS, using the SMP Expansion Ports for interprocessor communication (IPC) traffic would require some other enabling software. (IBM will supply these drivers as a standard part of the xSeries server offerings, where needed.)

How is the Enterprise X-Architecture solution better than standard distributed clustering?

If you are wondering if all this flexibility comes at the cost of performance, the answer is a resounding *No*! In fact, the high-speed SMP Expansion Ports provide performance unmatched in the Intel architecture arena. So how is four-node scalability better than simply connecting four stand-alone 4-way systems together to achieve a traditional four-server cluster? Although they do share some similarities there are significant differences:

- Four 4-way Enterprise X-Architecture nodes give you greater flexibility. They can be configured as two 8-way servers, an 8-way and two 4-ways, a 12-way and a 4-way, or as one 16-way server. Even after you have connected them a certain way, if your needs change you can still reconfigure the nodes into another arrangement using simple management console commands. You are not limited to one specific configuration.
- System partitioning allows you to set up a partition that utilizes all four nodes with one copy of an operating system and applications, rather than requiring four of everything (and the time and effort needed to individually configure and manage four separate systems) as would be needed with a traditional clustering solution.
- The SMP Expansion Ports allow nodes to talk to one another at up to **3.2 giga*bytes*** per second in bidirectional mode (roughly equivalent to **32 giga*bits*** per second) *per connection*, with each node supporting *up to three* connections to other nodes. 3.2GB is ***32 times*** what is currently available from even Gigabit Ethernet connectivity (or 16 times as much as Gigabit Ethernet configured for bidirectional operation through a LAN switch). Plus, your infrastructure doesn't have to be redesigned to support Gigabit Ethernet for just those few boxes—not to mention the cost savings of not needing Gigabit Ethernet hubs, routers, switches or adapters for just those clustered servers. If you don't use Gigabit Ethernet, consider that the SMP Expansion Ports are ***320 times*** as fast as 100Mbps Ethernet. And the ports are *in addition* to any Ethernet ports installed on the server, leaving any installed Ethernet ports available for normal network connectivity. (In other words, you don't have to tie up slots with Ethernet adapters in each server simply to connect to the other three servers.)
- Because the SMP Expansion Ports allow direct processor-to-processor communication among the nodes—without all the overhead of adapter bus and network traffic being processed by the OS—it results in better overall throughput among systems. (Performance will be discussed more later in this paper.)
- Shared I/O among nodes means that one instance (copy) of an operating system can access all the internal and external adapter slots on all the nodes at once--providing far more I/O than any single server cabinet could hold. (In other words, if you have a server that contains six adapter slots and you attach two additional external I/O expansion units with 12 slots apiece, those 30 slots will all look to the OS like they are in the server cabinet itself.)
- If you discover a workload imbalance (some servers are overworked while others are underutilized), you can easily disconnect a node from one server and move it to another one. This flexibility to redistribute capital resources to quickly accommodate changing needs is especially important in today's fast moving business world.
- Depending on your software, the licensing terms may be more favorable for one 16-way partition than for four separate 4-way servers.
- Should you prefer to cluster the four nodes as four separate servers, the SMP Expansion Ports can still be used for high-speed interprocessor communications (IPC), simply by connecting the same SMP Expansion Port cables. No complex setup is required.

In conjunction with large scale storage systems and a storage area network (SAN) infrastructure, the Enterprise X-Architecture model gives you the ability to provide access to multiple databases or to offer multiple points of access to a single database. Similarly, it enables you to execute one or multiple instances of the database management tools in a clustered configuration (assuming the DBMS supports it). This flexibility contributes to a reduction in the total cost of ownership by reducing the number of software licenses to be managed, as well as simplifying day-to-day operations and support.

XpandOnDemand scalability represents revolutionary advances for Intel-based servers. You'll be able to scale up server performance without any downtime: Adding (or removing) nodes will not require taking down the other nodes (with the proper operating system support). In addition, the modular node design offers the capability of failover among nodes, for maximum availability. If an

Enterprise X-Architecture configuration allows you to *scale up*, a clustering arrangement lets you *scale out*. In fact, both approaches can be used, by clustering (through traditional networking methods) a series of 8-, 12- or 16-way Enterprise X-Architecture systems.

System Partitioning

System partitioning is another of the many mainframe capabilities that Enterprise X-Architecture technology brings to Intel architecture servers. Among the benefits of system partitioning are hardware consolidation; software migration and coexistence; version control; development, testing and maintenance; workload isolation and independent backup and recovery on a partition basis.

Anyone who has worked with mainframes is familiar with the concept of partitioning: System resources, including processor, memory, I/O and storage are virtualized so that all concurrent users appear to have complete access to the system. Yet each user is actually segmented—and protected—from the actions of all other users. If one virtual partition were to lock up, it would not affect the others. In a mission-critical environment, such as a worldwide airline reservation system, for example, it would be disastrous if one such errant partition could crash the entire mainframe. (A partition is also protected from viruses and other security exposures that might affect another partition.) In many client-server environments, this capability is no less essential.

System partitioning, enabled by Enterprise X-Architecture design, describes two types of system partitioning: *physical* partitioning (enabled today) and *logical* partitioning (coming in the future). With physical partitioning, a single multinode server can *simultaneously* run multiple instances of one operating system in separate partitions (as well as multiple *versions* of an operating system or even different types of operating systems). The server can have up to four nodes—interconnected chassis containing independently operating processors, memory and I/O—each capable of running its own operating system and applications. (This is *not* the same as having multiple *hard disk* partitions loaded with different operating systems, where changing OSes requires rebooting to another partition. With *system* partitioning, the operating systems are running *simultaneously* in different nodes on the same server. A partition can span nodes—even to the point of having all four nodes serving one OS.) Each node can be managed independently by software.

For example, a server can continue to run an operating system in one node while you install and test *another* version of that operating system, or a different operating system entirely, in another node on that server—all without having to take the entire server offline. Multiple operating systems can function on the same server without interfering with one another.

Physical partitioning includes three modes: *Fixed*, *static* and *dynamic*.

- *Fixed* partitioning is done while the system is powered off, and involves the cabling together (or uncabling) of two or more physical nodes to modify the partitioning. After recabling, the operating system must be restarted.
- *Static* partitioning requires only the nodes being adjusted to be taken offline. The remaining nodes in the server remain unaffected and continue to operate normally. Static partitioning is performed on node or system boundaries. This means that a partition must have the hardware to function independently (processor, memory, I/O, etc.). It also means that one node can't be subdivided into multiple partitions, but a partition can consist of multiple nodes. Partitioning is done by accessing the offline server from a remote system running systems management software (such as IBM Director) before restarting the operating system. Due to the lack of support for more flexible partitioning in current operating systems this is the type of partitioning that will be available initially on Enterprise X-Architecture based servers.
- *Dynamic* partitioning has the same hardware boundaries as static partitioning. However, it permits hardware reconfiguring (adding or removing hardware) while the partition's operating system is still running. Servers based on the Enterprise X-Architecture design provide hardware support for dynamic partitioning. This capability requires extensive operating system modifications to support online insertion and removal of resources (essentially, plug-and-play

for 4-way processor complexes and individual nodes). Therefore, this form of physical partitioning will await the appropriate OS enhancements.

In the future, Intel processor-based servers running enhanced operating systems will be able to implement *logical* partitioning, with even higher levels of flexibility and granularity in running concurrent operating systems than physical partitioning provides. (Logical partitioning is the type used on IBM @server zSeries—formerly S/390®—mainframes.)

- Servers using logical partitioning will be able to reconfigure a multinode complex partitioned at the individual processor level (with associated memory, I/O and other required resources)—not just at the node level, without shutting down and restarting the hardware and software. More hardware can be added, or removed for maintenance, without powering the system off. When workload demands change, you can also reassign resources from one logical partition to another without having to shut down and restart the system. A simple user interface will guide you through the appropriate steps.

Figure 4 illustrates some sample partitioning implementations. Other configurations are possible.

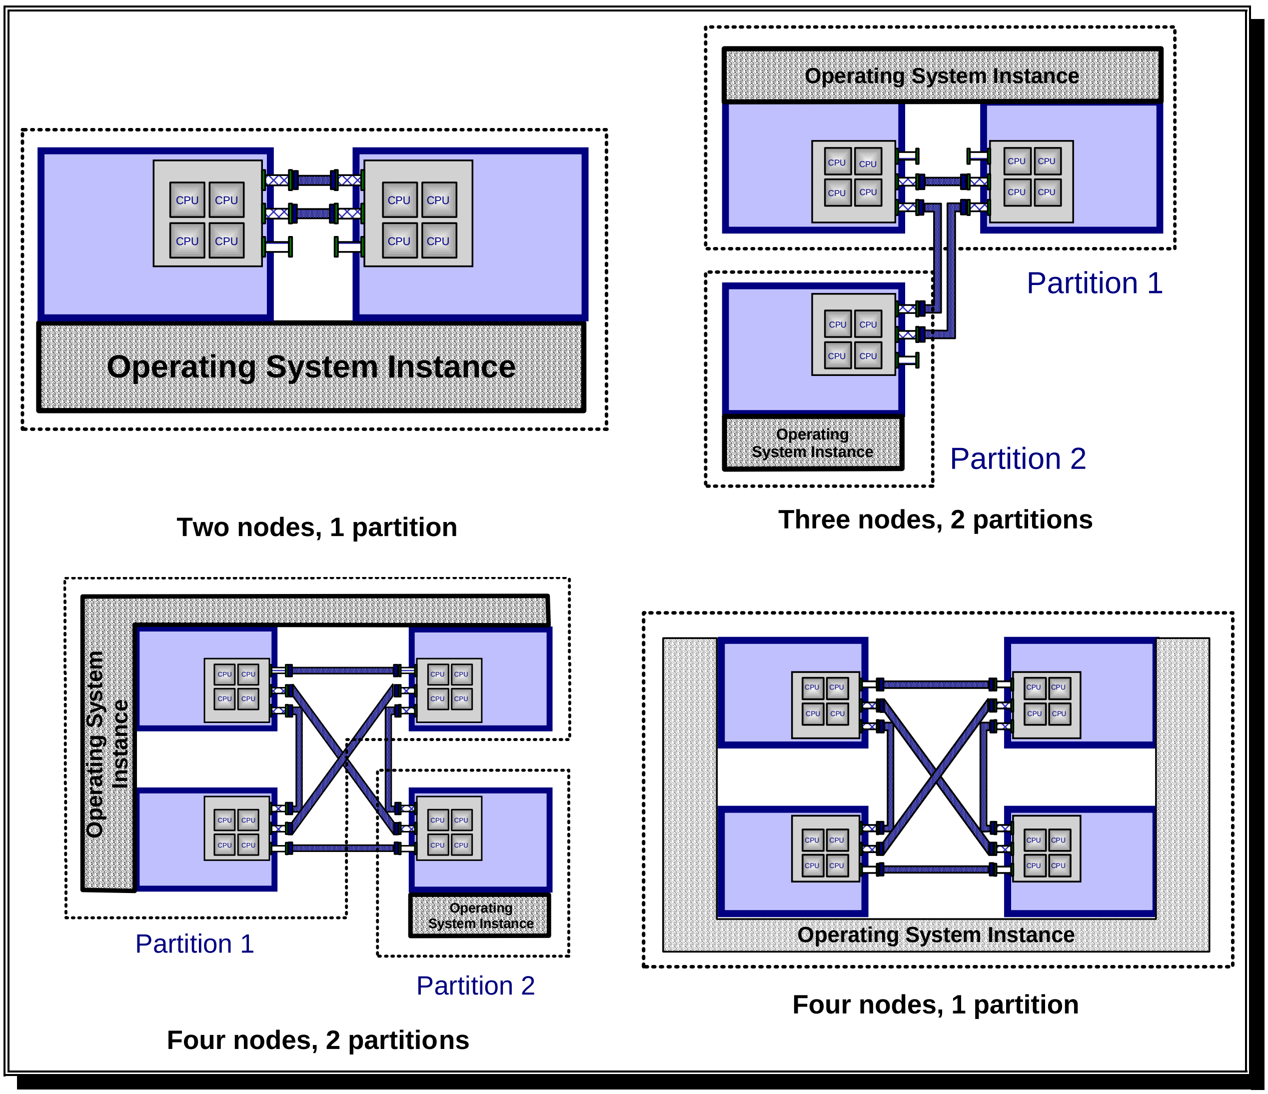

Figure 4. *Partitioning examples*

And that's just the beginning. Enterprise X-Architecture capabilities allow even more granular partitioning (down to the processor level, or lower: multiple partitions per processor) and process control (assigning specific tasks to specific processors). As with dynamic and logical partitioning it's simply a matter of waiting for operating system and "hypervisor" support for these features to catch up to the hardware. One such option available today is ESX Server from VMware. (See the *Other Software Support for Enterprise X-Architecture Technology* topic for more information.)

If you intend to consolidate servers system partitioning offers many benefits. As stated earlier, with system partitioning, the multiple operating systems previously used by multiple servers could all be running simultaneously on one server in one location, rather than being scattered across an organization. (Or many older, slower, single-processor and 2-way servers could be replaced by a few high-speed 4-way scalable enterprise nodes.) System partitioning allows you to eliminate the need to use multiple servers to support different operating systems within your business. One server with system partitioning could act as an application server, both for your marketing department that runs on Windows and for your engineering department that requires a Linux® server. Or you can have most users hosted by Microsoft Windows 2000 Advanced Server while you have a small development group beta testing another operating system in another partition. Or the full power of all nodes could be applied to one operating system instance.

To illustrate, compare the results of having four separate 4-way servers versus one system consisting of four 4-way nodes (assuming that the individual servers and the nodes are configured similarly): Let's say that each node and each stand-alone server has four processors, 8GB of memory, six adapter slots, two hard disk drives, one floppy drive, one CD-ROM drive, one SCSI controller, one Gigabit Ethernet controller, one systems management processor, one copy of an OS and one copy of whatever applications and utility software are installed.

In the case of stand-alone servers clustered together, there would still be four separate servers, with the capability for data sharing and failover—each with its own memory, adapters, software, etc. By comparison, the four-node Enterprise X-Architecture server could look like *one* 16-way server running *one* copy of the OS and applications, with access to 32GB of available memory, 24 adapter slots, eight hard disk drives, four floppy drives, four CD-ROM drives, four SCSI controllers, four Ethernet controllers, and four systems management controllers. In other words, it would be one megasystem with incredible expandability (not even counting the remote I/O capability). Alternatively, those four nodes could be configured as two 8-way systems (with 16GB of memory, 12 adapter slots, four hard disk drives, etc., apiece), with failover. You have the freedom to start out with the nodes configured one way, and then change the configuration later as needed. Some simple cabling changes (and software installation/deinstallation) and you have a whole new server configuration.

If you prefer a clustered solution, the Enterprise X-Architecture design works well in that environment too. You can start out with two 4-way nodes in a cluster, perhaps running Microsoft Cluster Service (MSCS), in a high-availability failover configuration. Each node would have its own operating system and applications. Over time your requirements might increase to the point where you need to add more computing power to the cluster. Just as in the integrated multinode Enterprise X-Architecture configurations, you can cluster a series of 4-way nodes in various combinations and still get the benefits of the 3.2GBps SMP Expansion Port throughput. For example:

- If you are in a high volume, high-availability technical/scientific computing environment, you might prefer a four-node, 4-way, *four-partition* cluster for a small Beowulf-class supercomputer configuration to perform massively parallel computing operations.
- A database environment running IBM DB2® Universal Database™ or Oracle might use a configuration of four 4-way nodes under a *single* partition working as a single 16-way server.
- Another alternative might be a three-node, 12-way single partition running the database as a back-end server, clustered with a one-node (4-way) server running a front-end application.

Enterprise X-Architecture technology gives you the ability to configure your system to best match current demands, while providing unprecedented flexibility to accommodate future needs.

PCI-X I/O Subsystem

The PCI bus has done an admirable job of keeping up with the I/O bandwidth needs of servers. PCI started out with a theoretical limit of 133MBps (mega*bytes* per second)—although the early systems could deliver only 30-40 MBps. The current state of the art in PC I/O buses allows multiple 64-bit 66Mhz PCI bus segments capable of delivering 400-500 MBps each. But even this bandwidth is not sufficient for the emerging world of 10Gbps (giga*bits* per second)—or higher—I/O, including Fibre Channel, Ethernet and InfiniBand. Without another performance boost, PCI would soon be the bottleneck that keeps these high-speed networks from connecting to servers at the networks' maximum speed. I/O bottlenecks already prevent industry-standard servers from the balanced systems architecture that is a characteristic of high-end systems.

Recognizing that the PCI bus is "running out of steam" as an I/O architecture and becoming a limitation to server performance, the industry developed an enhanced bus, called PCI-X, to extend the useful life of PCI until next-generation serial I/O architectures, such as InfiniBand, are ready to take over. PCI-X extends the useful life of PCI for a few more years by increasing the available I/O bandwidth considerably.

IBM has been a leader for many years in establishing industry standards. So too was IBM instrumental in the definition of the PCI-X specification. PCI-X provides a new generation of capabilities for the PCI bus, including more efficient data transfers, more adapters per bus segment and faster bus speeds for server systems. PCI-X enhances the PCI standard by doubling the throughput capability and providing new adapter-performance options while maintaining compatibility with PCI adapters. PCI-X allows all current 66MHz PCI adapters—either 32-bit or 64-bit—to operate normally on the PCI-X bus. PCI-X adapters take advantage of the new 100MHz and 133MHz bus speeds, which allow a single 64-bit adapter to move as much as 1GB of data per second. (The next PCI-X specification (2.0) will support bus speeds of up to 266MHz.) Additionally, PCI-X supports twice as many 66MHz/64-bit adapters in a single bus as PCI.

Active PCI-X

Building upon the success of IBM Active PCI, Enterprise X-Architecture design introduces Active PCI-X with support for both PCI and PCI-X adapters in selected xSeries servers. Active PCI-X provides another IBM solution to increase total server availability. Active PCI-X features can be categorized as follows:

- **Hot-swap** — Allows you to replace without having to shut down and restart the server.
- **Hot-add** — Provides easy upgradeability, by allowing you to add new adapters to the server while it's running. (IBM was the first in the industry to offer this capability.)
- **Failover** — Allows a second—backup—adapter to pick up the workload on the fly if the primary adapter fails.

Remote I/O

A limitation to the PCI and PCI-X industry specifications is that all adapter slots must be in the main system cabinet. This requires a trade-off in server implementation. On one hand, designers want to shoehorn the maximum number of adapter slots into a box. At the same time, they want to shrink the size of the server as much as possible as components become ever smaller, to minimize the rack space required. However, reducing the size of the system beyond a certain point means sacrificing adapter slots. (This is why server vendors offer models with many slots and few drive bays or many bays but few slots. It's either that or increase the overall size of the server—which causes other problems in a crowded data center.)

One obvious solution is to limit the number of slots inside the server chassis while extending additional bus segments to external "card cages" for holding adapters. By moving adapter slots outside of the main system cabinet, the server can be made much smaller. Unfortunately, the PCI-X bus specification doesn't explicitly support external I/O.

Using IBM Enterprise X-Architecture remote I/O (RIO) enclosure support, it is possible to add *dozens* of PCI/PCI-X adapter slots via external I/O expansion boxes to a single server[3], providing incredible I/O scalability. This is yet another example of what we mean by "pay as you grow" scalability. Buy only what you need, when you need it. It allows IBM to continue to shrink server cabinets while increasing I/O scalability via external expansion units.

RIO provides the capability of more available adapter slots per server, but also potentially greater reliability. (Using separate cabinets with their own power supplies and fans provides redundancy and may result in less heat build-up than if many adapters and drives are installed in the main unit.) Another advantage of an external I/O cabinet is that *it can be shared by more than one server.* (See *Figure 5*.)

Remote I/O is based on the time-tested interconnect technology used on IBM @server pSeries and iSeries (formerly IBM RS/6000® and AS/400®, respectively) servers. Remote I/O, in conjunction with the XpandOnDemand system node design, allows you to expand your systems in an orderly, controlled fashion.

In the past, to ensure adequate growth capacity in a server, you had to buy one large monolithic tower with many processor sockets, adapter slots and drive bays—for a huge initial price. Perhaps you had to buy more than you would ever use, to be sure you had enough capacity. Now, with IBM XpandOnDemand scalability solutions, you can start out small with a single rack-mounted node. Later you can add additional nodes, I/O expansion cabinets and disk expansion units as needed (and as your budget allows) to an industry-standard rack. Not only does Enterprise X-Architecture remote I/O give you great expandability, it comes with amazing performance as well. The remote I/O expansion units connect to the system units using a **2GBps** dedicated I/O port.

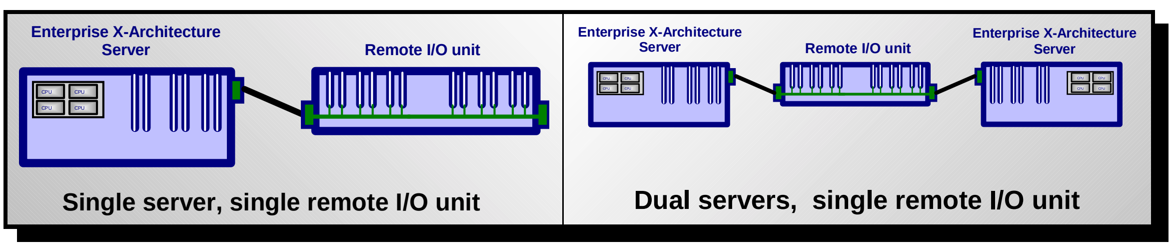

Figure 5. *Remote I/O configuration examples*

Active Memory

IBM has a long history of pioneering advances in memory development, including the first one-transistor memory cell (1966), the first use of cache memory (1968), the first 1 megabit memory chip (1984) and ECC-on-SIMM (1996). IBM was active in the development of the DDR (double data rate) memory specification and was the first to ship servers containing DDR memory (2000). Continuing along this path, IBM is delivering a number of memory technology breakthroughs via Active Memory that increase capacity, performance and reliability:

- **Large memory capacity** — There are any number of reasons to buy additional servers. You may have run out of room to add processors, or you may need more adapter slots (something that's solved via Enterprise X-Architecture remote I/O), or you may need more memory than your server can hold. While some servers are constrained by the number of memory sockets they can hold, others are limited by the maximum amount of memory that can be addressed

[3] Subject to any hardware limitations of a specific server implementation.

by the chipset the server is using. Most servers are limited to 16GB of RAM or less for these reasons. The Enterprise X-Architecture design smashes that barrier by enabling the use of as much as **256GB** of RAM in a 64-bit Itanium-based server (64GB in a 32-bit Xeon MP server). This is enough to hold most databases entirely in memory, with potentially huge gains in performance over databases that are accessed primarily from disk. The concern preventing most users from considering running their databases entirely from memory is the fear of a memory failure causing a system crash, with an attendant loss of data. The possibility of such a failure is the reason for the major enhancements in memory reliability and availability provided by the Enterprise X-Architecture solution.

- **High-speed memory access** — Today, the fastest Intel processor-based servers have a front-side bus (FSB) speed of 133MHz (many servers still use a 100MHz FSB). This determines how fast the processor can access main memory and external cache memory. By contrast, Enterprise X-Architecture technology enables servers to implement the 400MHz FSB of the Intel Xeon (Foster) and second-generation Itanium (McKinley) processors. This means reads and writes of memory by the processor will be at triple today's servers' best speed, due to higher bandwidth and lower latency (waiting for the memory to be ready for the next read/write). In addition, Enterprise X-Architecture design also supports the use of double data rate (DDR) main memory, for even higher performance.
- **Memory ProteXion** — Memory ProteXion helps protect you from unplanned outages due to memory errors far more effectively than standard ECC technology, even while using standard ECC DIMMs. It works somewhat like hot-spare disk sectors in the Windows NTFS file system, where if the operating system detects bad sectors on disk, it will write the data to spare sectors set aside for that purpose. Think of Memory ProteXion as providing hot-spare bits. The error correction is handled by the memory controller, so there is no operating system overhead or support requirement—it's transparent to the OS. Because it operates on standard ECC DIMMs, you don't have to pay extra for this protection either.

  Memory ProteXion (referred to as redundant bit steering on other IBM systems) was originally developed for IBM mainframes and has been in use for many years on zSeries and iSeries servers. IBM reliability testing and analysis shows that a server protected with Memory ProteXion is nearly **200 times** less likely to fail than one using standard ECC memory. For example, given identical servers with 8GB of RAM, users can expect to experience approximately one failure per year for every **132** servers using ECC memory. By comparison, servers using Memory ProteXion should see only one memory failure for each **26,042** servers per year!

  On a 2-way interleaved memory system, two ECC (Error Checking and Correcting) DIMMs contain 144 bits, but only 140 bits are used for data and checksums. The remaining four bits are unused. Standard ECC memory can detect two-bit errors, but it *corrects* only single-bit errors. If multiple bits of a memory chip fail at once, the whole DIMM fails, crashing the server and temporarily leaving the system with reduced memory capacity (until the module is replaced).

  Memory ProteXion, instead of immediately disabling the DIMM, merely rewrites the data to some of those spare bits. This allows correction of as many as four sequential bit failures per pair of DIMMs—eight per *memory controller* (and a server may have more than one controller). The next time the server is restarted it rechecks the memory status. In the case of a soft (temporary) memory error, the hot-spare bits are freed up for reuse. If a hard (permanent) error occurs, the hot-spare bits are again utilized to keep the DIMM operating until it is replaced.

  This advanced technology helps reduce server downtime, resulting in a more robust client-server computing platform. This can be especially important in a large database environment where transaction commit/rollback, reindexing and data synchronization among servers can result in *hours* lost before a crashed database is operational again.

  But what happens if a memory controller has used up the available spare bits and another error occurs before the next server restart? Does the server crash at that point? No, because

Memory ProteXion is only the *first* line of defense against memory errors. If a memory controller runs out of spare bits, it continues on to the second line of defense, Chipkill memory.

- **IBM Chipkill memory** — IBM Chipkill ECC memory (now in its third generation in industry-standard computers) comes into play only if a server encounters so many errors in a short span of time that Memory ProteXion can't handle them all. This should be a rare occurrence, but if it does happen you are still protected. Like Memory ProteXion, Chipkill memory goes well beyond the error-correction afforded by standard ECC memory, providing correction for up to four bits per DIMM (eight bits per memory controller), whether on a single chip or multiple. Also like Memory ProteXion, Chipkill support is provided by the memory controller, so it is implemented using standard ECC DIMMs and is transparent to the OS.

  Chipkill memory was initially developed for IBM mainframes 20 years ago and utilized in NASA's Pathfinder mission to Mars. It's another example of the IBM commitment to provide robust, high-availability systems. The memory controller provides memory protection similar in concept to disk array striping with parity, writing the memory bits across multiple memory chips on the DIMM. In essence, each DIMM acts as a separate memory array. If any one chip fails, it affects only a single bit from a byte of data, because the other bits are stored on other, working chips. The controller is then able to reconstruct the "missing" bit from the failed chip and continue working as usual.

  Between Memory ProteXion and Chipkill memory, memory errors that go uncorrected should be an extremely rare occurrence.

- **Memory mirroring** — The third line of defense against server downtime due to memory failure is memory mirroring. Should a server somehow encounter so many memory errors that Memory ProteXion and Chipkill memory together cannot correct them all, mirroring will keep your system running. Memory mirroring conceptually works much like RAID-1 disk mirroring, in that data is simultaneously *written* to two independent memory cards. (Each card is configured identically.) Memory is *read* only from the memory card designated as active. (See *Figure 6*.)

  If a DIMM begins producing soft errors sufficient to trigger system alerts—giving the administrator plenty of warning that the DIMM is beginning to fail—or if the DIMM fails entirely, the server will automatically shift to using the mirrored (offline) memory card until the failing DIMM is replaced. This allows the system to keep running until a convenient time to perform maintenance on the failing memory module. Memory mirroring allows the ability to hot swap and hot add memory DIMMs. (Because memory mirroring means that only half of the system memory capacity is available to software, mirroring can be disabled in BIOS, if desired.)

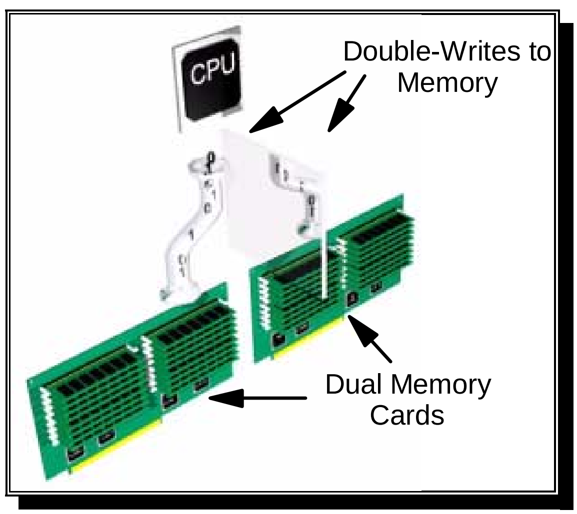

Figure 6. *Memory Mirroring*

- **Hot-add/hot-swap memory** — Hot-*swap* memory allows failing parts to be replaced while the server is still running. Likewise, hot-*add* support allows new DIMMs to be added to a running server as memory demands increase. Selected IBM @server xSeries servers already permit service personnel to hot add new drives, adapters, power supplies and fans as needed. Now we have added memory to the list.

The more memory installed in a server, the higher the probability of a memory-related system error. With servers capable of holding tens and even hundreds of gigabytes of RAM, reliability is becoming more of an issue than ever before. As disk storage increased in capacity over the last two decades users looked for ways to increase performance and protect their data. These needs resulted in disk arrays (such as RAID-0 data striping, RAID-1 disk mirroring and RAID-5 striping with parity), as well as offline storage, such as tape drives. Similarly, the need to protect data has resulted in the memory technologies described here. Because Chipkill memory, Memory ProteXion, memory mirroring and hot swap capability are implemented entirely in hardware, there is no operating system overhead or dependencies. (Hot-add memory does require future operating system support.)

All of these memory protection mechanisms are tried-and-true technologies that have been available on IBM mainframes and other large systems for years. Just as importantly, they are all implemented using industry-standard ECC DIMMs, therefore memory is inexpensive and readily available.

XceL4 Server Accelerator Cache

Another advanced performance-enhancing feature is support for a huge Level 4 (L4) system cache, with up to 64MB[4] of DDR (double data rate) high-speed ECC memory per 4-way central electronics complex (CEC) in Itanium-based servers or 32MB in Xeon MP systems. Years ago when microprocessor vendors were looking for ways to improve system performance, they ruled out using fast but expensive SRAM chips for main memory. Fortunately, they discovered that using a smaller amount of high-speed memory in a separate processor cache (which held a copy of the most frequently accessed data from main memory) provided much of the same performance boost, but at a much lower cost. From that point on microprocessors have included built-in cache memory.

Intel 32-bit and 64-bit processors contain a relatively small amount (128K to 4MB, depending on the processor) of L1, L2 and (in the case of the Itanium) L3 *internal* cache memory. The amount of internal cache is limited by the available room inside the processor module. The more high-speed cache memory there is, the more often the processor finds the data it needs and the less often it has to access the more distant (and thus slower to respond) main memory. (Processor speed is increasing at a much faster rate than is main memory speed, so each year the penalty increases for having to access main memory when a cache "miss" occurs.)

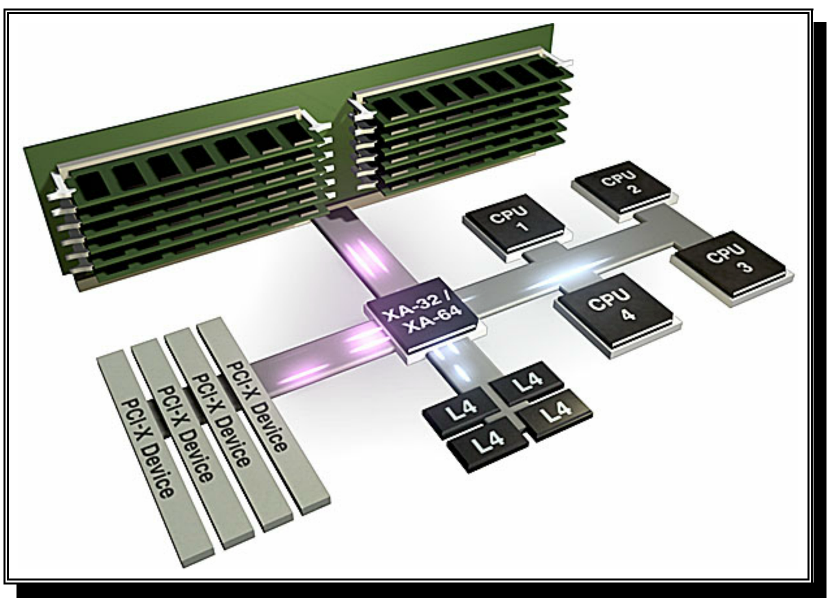

Figure 7. *XceL4 Server Accelerator Cache*

The addition of a large *external* L4 cache (along with the faster and wider front-side bus described previously) reduces memory latency and increases memory bandwidth for unprecedented performance. (See *Figure 7* for a conceptual illustration.) Without an L4 cache, a high-end processor like the Itanium can consume as much as 80-85% of available main memory bandwidth, leaving little for other devices that require memory bandwidth (including the PCI-X, Ethernet, SCSI and Fibre Channel interfaces).

[4] On an Itanium processor-based system; **32MB** on a Xeon MP processor-based server.

By using high-speed DDR memory, and interposing it between the processors and user RAM, the IBM XceL4™ Server Accelerator Cache gives processors and I/O devices alike a considerable boost in performance. How much of a boost? In an industry where vendors tout even a 2% performance advantage over competitors, XceL4 caching increases overall server throughput by as much as **15%-20%**!

Active Diagnostics

Expandability, performance and economy are all important features in a server, but equally important is your ability to prevent or minimize server downtime. A technology new to industry-standard servers that helps in that regard is IBM Active Diagnostics. Although not directly enabled by the XA-32/XA-64 chipsets, Active Diagnostics is another Enterprise X-Architecture feature that will be incorporated in many of the servers that use the chipsets.

Based on the Distributed Management Task Force (DMTF) Common Information Model (CIM), Active Diagnostics allows an administrator to run diagnostics on system resources while users are still working, thereby increasing system uptime and taking IBM customers closer to OnForever levels of high availability. This is a feature that has been sorely lacking in the industry-standard server market, and this void is being filled with Enterprise X-Architecture technology.

IBM, working with Intel and PC-Doctor, introduced industry-standard extensions to CIM to support concurrent diagnostics in 1999. This Common Diagnostic Model (CDM) standardizes how to interface concurrent diagnostics (also called OS-Present Diagnostics or On-Line Diagnostics) with the operating system, making the diagnostic routines available to all CIM-aware diagnostic applications. With the diagnostic extensions to CIM now defined, IBM is working with independent hardware vendors to define methods and create concurrent diagnostic tools and utilities to keep xSeries servers up and running.

IBM Active Diagnostics is our implementation of CDM. The combination of IBM Predictive Failure Analysis® (system health monitoring), Active Diagnostics and hot-swap components means that you may never again have to shut down your xSeries server to run diagnostics or replace hot-swappable components.

Active Diagnostics can be operated through the IBM Director management software to provide a consistent, easy to use management interface for controlling many system functions.

IBM Director

IBM Director[5] is a powerful suite of tools and utilities included with xSeries servers. Built upon industry standards, IBM Director is designed to manage servers in the Intel environment and support a variety of operating systems, including Microsoft Windows, IBM OS/2®, Novell NetWare, Linux and SCO UnixWare. IBM Director supports a variety of industry standards such as DMI, CIM, WBEM, WMI, SNMP, MPM, TCP/IP, IPX, SNA, NetBIOS, SLIP, XML and HTTP, among others.

IBM Director is being updated to support Enterprise X-Architecture capabilities. IBM Director will gain the ability to distinguish between a node and a partition, so that systems management hardware alerts can be pinpointed by both partition and node.

IBM Director consists of three major components: a management server, the management console and agent, plus an optional fourth component, Server Extensions. (*Figure 8* shows how the new IBM Director console may look with the Enterprise X-Architecture enhancements. Note the system partitioning tasks.)

---

[5] For more on IBM Director, read the white paper "IBM Director with UM Services™" at http://www5.pc.ibm.com/ww/me.nsf/Technical+Information%5cxSeries+and+Netfinity?OpenView&WW.

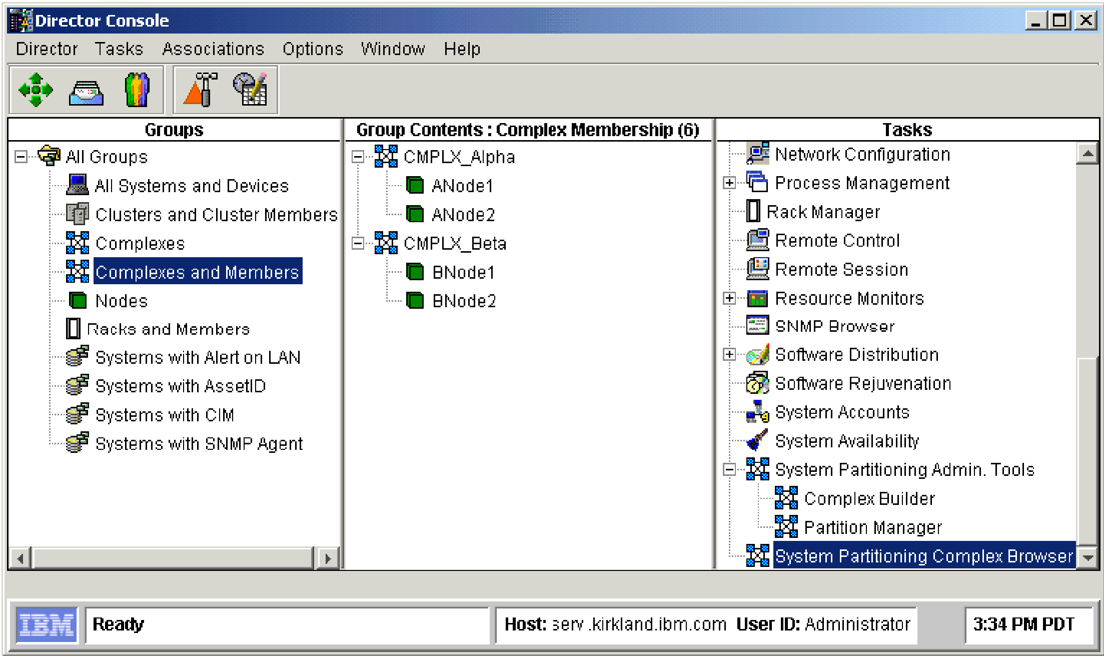

Figure 8. *IBM Director console*

The IBM Director Server Extensions are often referred to as Life Cycle Tools. These no-charge programs, included with xSeries systems, extend the manageability of your server hardware throughout its life cycle to help administrators configure, deploy, manage and maintain IBM @server xSeries servers efficiently. Today, IBM Director Server Extensions include:

- Advanced Systems Management
- Capacity Manager
- Cluster Manager
- RAID Manager
- Software Rejuvenation[6]
- Rack Manager
- System Availability

IBM Director is being updated to add the following new Server Extensions for Enterprise X-Architecture technology:

- **System Partition Manager.** System Partition Manager provides a graphical interface for creating partitions. It allows an administrator to configure a specific server (while it is offline) from a remote system, prior to starting the OS. System Partition Manager uses the network link to the onboard systems management processor to establish the relationships among nodes. These relationships are maintained in a persistent database and can be recalled and activated at any time using the graphical interface. Because System Partition Manager integrates with IBM Director (see *Figure 8*, previously), it is part of a common management infrastructure that is used to manage a running partition. (*Figure 9* shows an early version of the System Partition Manager.)

[6] For more about IBM Director Software Rejuvenation, read the white paper by that name at http://www5.pc.ibm.com/ww/me.nsf/Technical+Information%5cxSeries+and+Netfinity?OpenView&WW .

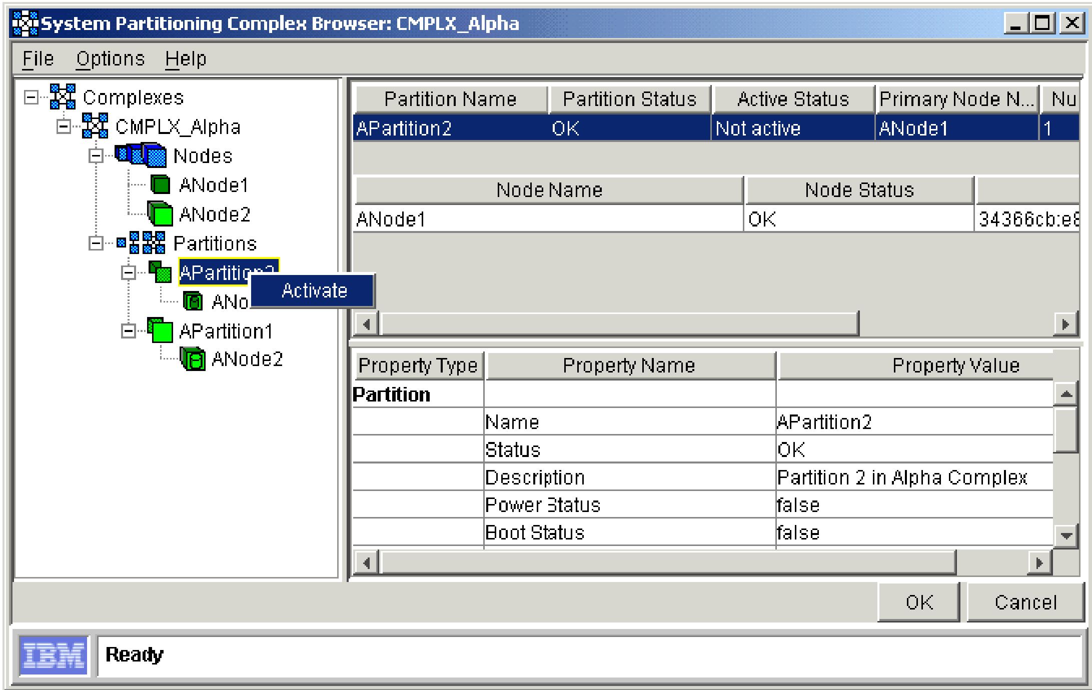

Figure 9. *System Partition Manager*

- **Active PCI Slot Manager.** Active PCI Slot Manager helps optimize I/O performance by offering guidance on the best slots in which to install PCI and PCI-X adapters. It also helps you decide whether adapters should go in the server or a remote I/O expansion unit. (*Figure 10* shows how this tool—still in development—might look in the final version.)

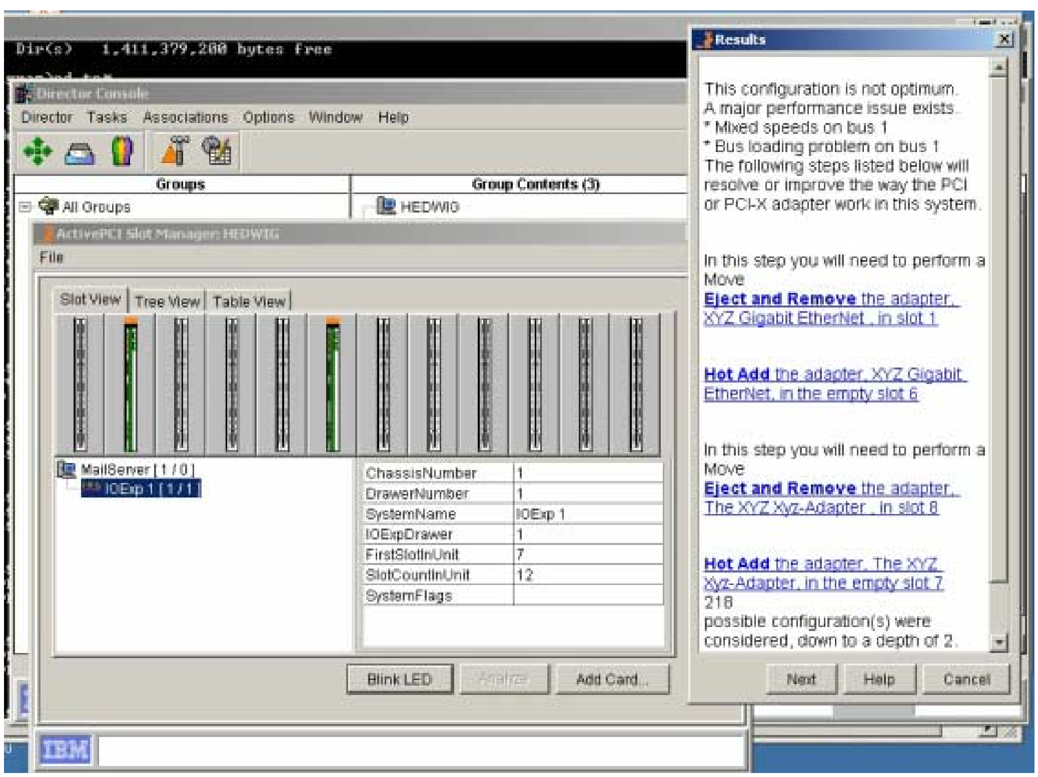

Figure 10. *Active PCI Slot Manager*

Other Software Support for Enterprise X-Architecture Technology

New versions of operating environments are being developed to optimize Enterprise X-Architecture features and performance. Applications need not be modified to take advantage of these technologies. Some of the software enhancements include:

- **Microsoft Windows 2002 Advanced Server/Datacenter Server** — Microsoft Windows 2000 Datacenter Server already supports servers with up to 32 processors and 64GB of memory. Microsoft is adding code that will optimize the performance of Enterprise X-Architecture servers. Microsoft is expected to add support for dynamic partitioning in future server versions of Windows.
- **Linux** — The xSeries development team, the IBM Linux Technology Center and other parts of IBM are working with the Linux community, and with our core Linux alliance partners (Caldera, Red Hat, TurboLinux, and SuSE) to exploit advanced Enterprise X-Architecture features in upcoming releases of Linux. These features include memory optimization, remote I/O support, hot-swap/hot-add capabilities and dynamic partitioning, among others.
- **Novell NetWare** — IBM is working with Novell's Modesto technology team to take full advantage of the Enterprise X-Architecture capabilities. Novell will support high scalability, large memory, PCI-X and remote I/O, and partitioning in all forms.
- **VMware ESX Server** — ESX Server is familiar to many people as an operating environment that allows various versions of Microsoft Windows and Linux (along with their applications) to be run in “virtual machines.” This type of product will allow Enterprise X-Architecture users to set up multiple partitions per *processor*, rather than the current maximum of one partition per *node*, providing much more flexibility in server configurations. IBM plans to test the current version of this software (as well as future implementations).

Future Directions for Enterprise X-Architecture Technology

This paper has described the technologies currently enabled by the Enterprise X-Architecture platform. Yet the Enterprise X-Architecture initiative is an ongoing effort, as is the X-Architecture blueprint in general, with support for new capabilities planned to be added over time. Technologies that may be added in the future[7] include:

- **Logical system partitioning** — Described earlier, logical partitioning would allow hardware nodes to be partitioned at the individual processor level (with associated memory, I/O, etc.), rather than at the node level currently supported by physical partitioning.
- **Support for >16-way servers** — Future IBM chipsets will support greater than 16-way servers. (The architecture itself already supports 32-way servers.)
- **InfiniBand I/O** — InfiniBand architecture is the next generation of advanced I/O following PCI-X. InfiniBand is a robust new industry-standard design for a high-speed, redundant I/O architecture to supersede PCI and PCI-X in servers. IBM co-founded and co-chairs the InfiniBand Trade Association (IBTA), which announced InfiniBand Architecture specification V1.0 on October 24, 2000. The IBTA consists of over 200 member companies.

  InfiniBand is a new high-speed interconnect architecture that redefines how servers communicate with networks, storage subsystems and other servers. It also moves all the data transfer and validation functions from software to hardware, making InfiniBand highly efficient. InfiniBand standardizes how large numbers of server, storage and networking components should be joined and balanced into a single integrated system. It does this without requiring any changes to how existing applications access networks and storage.

---

[7] All statements regarding future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Many of the technologies incorporated into this new architecture originated with IBM S/360 and S/370 mainframe systems, which were the first systems to use channel I/O as a means to achieve a balanced system architecture. InfiniBand is ideally suited for clustering, I/O extension and native attachment in diverse network applications. InfiniBand technology can be used to connect to attached hosts, routers or disk arrays and other devices. InfiniBand also features enhanced fault isolation, redundancy support and built-in failover capabilities to provide high network reliability and availability. Enhancements even extend to cabling, with InfiniBand supporting the concurrent use of both copper and fiber-optic cables—copper for low cost and fiber for long distances.

For xSeries servers, InfiniBand will provide the framework to continue our drive toward high availability, scalability, ease of use and reduced total cost of ownership embraced by the IBM X-Architecture strategy. IBM intends to accomplish this through multiple, high-speed, low-latency and redundant network paths for server access to I/O, as well as through clustering. Servers and I/O devices incorporating InfiniBand technology are expected in 2002. However, widespread adoption is not expected before 2003-2004, when the technology is integrated sufficiently to deliver comprehensive hardware and software solutions. In the meantime, PCI-X will be the advanced I/O bus of choice.

How Enterprise X-Architecture Technologies Fit into Project eLiza

Enterprise X-Architecture technologies mesh well with the IBM Project eLiza initiative. Project eLiza is a multibillion dollar program aimed at creating self-managing servers that require little or no human interaction, with the goal of making computing networks as easy to manage as today's kitchen appliances. Four major categories serve as the foundation for Project eLiza technologies: self-optimizing, self-healing, self-configuring and self-protecting.

Project eLiza incorporates many IBM X-Architecture features, including the high-availability aspects of the IBM OnForever™[8] initiative. Utilizing IBM @server xSeries servers, businesses now have access to a number of critical self-managing capabilities that help predict, identify and repair problems before availability is affected. Some of these features include:

Self-Configuring	Self-Optimizing	Self-Protecting	Self-Healing
• Active PCI / PCI-X • C2T Interconnect cabling • Hot-add components • LANClient Control Manager™ (LCCM) • ServerGuide™ • Software Delivery Assistant • System Migration Assistant • System Partitioning • Update Xpress™ • XpandOnDemand scalability	• Clustering • IBM Director Capacity Manager • IBM Director Active PCI Slot Manager • IBM Director Cluster Manager • IBM Director System Partitioning Manager • Memory eXpansion Technology (MXT)	• Advanced Systems Management processors • Chipkill memory • Electronic Solution Assurance Review (eSAR) • Fibre Array Storage Technology (FAStT) • Hot-swap/redundant components • Memory mirroring • Memory ProteXion • Predictive Failure Analysis (PFA) • ServeRAID™ FlashCopy™	• Active Diagnostics • Automatic Server Restart • Electronic Service Agent • IBM Director Software Rejuvenation • Light Path Diagnostics • Wake on LAN®

[8] For more information on IBM high-availability technologies, read the white paper entitled "IBM OnForever Initiative" at http://**ibm.com**/eserver/xseries. From the xSeries home page, select **Library** for links to the different types of documentation available.

Visit http://ibm.com/servers/eserver/introducing/eliza for more information about Project eLiza.

Conclusion

Today, IBM Enterprise X-Architecture technology offers unprecedented capabilities for industry-standard servers:

- Unparalleled scalability with system nodes, system partitioning, large memory support and remote PCI-X I/O expansion. This in turn creates a new economic model for servers.
- New levels of reliability, availability and serviceability with Memory ProteXion, Chipkill memory, memory mirroring and hot-add/hot-swap memory, Active PCI-X adapter support and Active Diagnostics. These technologies all dovetail nicely with other high-availability features of the IBM OnForever and new Project eLiza initiatives to help reduce downtime and decrease the costs of servicing and managing your servers.
- Unheard of performance improvements from the SMP Expansion Ports—connecting system nodes together into an 8-, 12- or 16-way server (or an "instant" cluster of 4-way nodes)—a large L4 system cache and a 400MHz front-side bus, all working together to speed through the most complex jobs.

In the near future, the Enterprise X-Architecture platform may be enhanced with greater than 16-way scalability, logical system partitioning and InfiniBand I/O support.

Someday, every industry standard server may have these capabilities, but for now the only place to get them all is from IBM.

IBM Enterprise X-Architecture technologies represent the dawn of a new age in industry-standard computing. By bringing the best of the mainframe capabilities to the Intel-based server platform, IBM is able to offer proven, high-performance solutions in a new building-block design that allows for easy expansion. With each new Enterprise X-Architecture feature designed with protection of your server investment in mind, you'll grow your business with the confidence that your technology will not only be able to handle increasing capacity, but also manage the growth for you. When you build your network based on this blueprint, your enterprise will experience revolutionary scalability, availability and flexibility—while decreasing the costs of servicing and managing your servers. Do more with your IT dollars. Do more with your IT staff. Do more with your data center. Do it with Enterprise X-Architecture technology.

Additional Information

Visit our Web site at http://**ibm.com**/eserver/xseries for more information on IBM @server xSeries servers direction, products and services, including IBM Enterprise X-Architecture technologies. From the xSeries home page, select **Library** for a list of the types of documentation available.

IBM Server Group
Department 4VNV
Research Triangle Park NC 27709

Produced in the USA
8-01