


Windows Server 2003, Datacenter Edition on the IBM **@**server xSeries 445

Describes the features of the
xSeries 445

Introduces the capabilities of
Datacenter Edition

Explains why Datacenter
Edition and the x445 are an
ideal match



David Watts
Antony Collins
Russ Weiner



International Technical Support Organization

**Windows Server 2003, Datacenter Edition on the
IBM @server xSeries 445**

August 2003

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (August 2003)

This edition applies to Windows Server 2003, Datacenter Edition, running on the IBM @server xSeries 445, machine type 8870.

© Copyright International Business Machines Corporation 2003. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	v
Trademarks	vi
Preface	vii
The team that wrote this Redpaper	vii
Become a published author	ix
Comments welcome	ix
Chapter 1. Features of the x445	1
1.1 Overview	2
1.2 X-Architecture	2
1.3 Enterprise X-Architecture	3
1.4 NUMA architecture	4
1.5 XpandOnDemand	7
1.6 SMP Expansion Module	8
1.7 Memory implementation	9
1.8 IBM ServerProven®	10
1.9 Microsoft HCL	11
1.10 Comparing the x445 to the x440	11
Chapter 2. Features of Windows Server 2003, Datacenter Edition	13
2.1 Description of Windows Server 2003, Datacenter Edition	14
2.1.1 Windows Server 2003 new and updated features	14
2.1.2 Requirements	19
2.1.3 Comparing Datacenter to Enterprise and Standard Editions	20
2.1.4 Features not available in Windows Server 2003, 64-bit editions	22
2.2 Why Datacenter?	23
2.3 High availability	23
2.4 Clustering	25
2.5 Server consolidation	27
2.5.1 IBM's role in server consolidation	28
2.5.2 Datacenter Edition clustering in a consolidated server environment	29
2.6 The Datacenter High Availability Program	30
2.6.1 Datacenter test programs	30
2.6.2 Datacenter Certified Applications program	32
2.6.3 Datacenter service provider and vendor programs	32
2.6.4 Datacenter support and maintenance	33
Chapter 3. Datacenter and the x445	35
3.1 Why Microsoft Datacenter from IBM?	36
3.2 Why the x445 and Datacenter Edition?	36
3.3 NUMA support in Windows Server 2003	38
3.3.1 Using SRAT to support hot-add memory	40
3.4 How to engage IBM	40
3.5 Solution assurance	41
3.6 Services	42
3.7 HARQ and the JSQ	44
3.8 High Availability Guarantee Program	44

Abbreviations and acronyms	47
Related publications	49
IBM Redbooks	49
Online resources	49
IBM Web pages	49
Microsoft Web pages	50
Help from IBM	50

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law. INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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



This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Chipkill™	ibm.com®	Redbooks (logo)  ™
Domino™	Lotus®	ServerProven®
DB2®	Netfinity®	Tivoli®
DFS™	Notes®	TotalStorage®
@server™	OnForever™	X-Architecture™
 server™	Perform™	xSeries®
eServer™	Predictive Failure Analysis®	
IBM®	Redbooks™	

The following terms are trademarks of International Business Machines Corporation and Rational Software Corporation, in the United States, other countries or both:

Rational Software Corporation®

Rational®

The following terms are trademarks of other companies:

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

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

Preface

This IBM® Redpaper describes two flagship products from IBM and Microsoft® and explains why these products are a perfect fit for one another. The IBM @server™ xSeries® 445 is the next-generation scalable enterprise server, with up to sixteen processors and up to 128 GB of memory. Microsoft Windows® Server 2003, Datacenter Edition is the high-end operating system for Intel®-processor servers. These two products are ideal for enterprise customers who require an extremely large single-image business computing environment.

This Redpaper is divided into three chapters:

Chapter 1, “Features of the x445” on page 1 covers Enterprise X-Architecture™ and NUMA, the architectures upon which the x445 is based, and describes the hardware technology that makes up the various subsystems in the server.

Chapter 2, “Features of Windows Server 2003, Datacenter Edition” on page 13 covers the features of the operating system, why customers should use Datacenter Edition over other Windows operating systems, and where clustering, high-availability, and server consolidation are appropriate in enterprise environments.

Chapter 3, “Datacenter and the x445” on page 35 examines the relationship between IBM, Microsoft and the customer when a Datacenter Edition-based solution is being implemented.

This Redpaper is aimed at customers, Business Partners and IBM employees who need to understand why Datacenter Edition and the xSeries 445 are an ideal combination for high capacity business computing environments.

The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

David Watts is a Consulting IT Specialist at the International Technical Support Organization in Raleigh. He manages residencies and produces IBM Redbooks™ on hardware and software topics related to IBM xSeries systems and associated client platforms. He has authored over thirty Redbooks and Redpapers; his most recent include *Implementing Systems Management Solutions using IBM Director* and the *IBM @server xSeries 445 Solution Assurance Product Review Guide*. He has a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM for over fifteen years. He is an IBM @server Certified Specialist for xSeries and an IBM Certified IT Specialist.

Antony Collins is an xSeries Technical Specialist in Sydney, Australia. He has seven years of experience in Windows systems and operational management. His areas of expertise include solution design, systems management and xSeries server architecture. As an xSeries Specialist for IBM Australia, Antony provides pre-sales xSeries support and design assistance to IBM professionals, IBM Business Partners and IBM customers. Antony is a Certified Specialist for xSeries, Windows NT® and 2000 MCP, and also a Certified Novell Administrator (CNA).

Russ Weiner is an xSeries Technical Support Specialist in Seattle, Washington. He has twenty years of experience in information technology as a programmer, network administrator, and in mainframe operations, the help desk environment, and data security. He has worked at IBM in the Advanced Technical Support organization for five years. His current area of focus is field technical support for xSeries servers.



The redbook team (l-r): Antony, David, and Russ

Thanks to the following IBM employees for their contributions to this project:

Donn Bullock, Global Brand Manager, xSeries 445, Raleigh
Mark Chapman, xSeries Marketing, Raleigh
Henry Chung, Technical Project Manager, Datacenter Development, Seattle
Ray Isler, WW Service & Support Planning, Technical Support Marketing, Raleigh
John Kranz, Datacenter Certification, IBM Center for Microsoft Technologies, Seattle
Michael S. Lee, IBM Center for Microsoft Technologies, Seattle
Grace Lennil, Technology Development, IBM Center for Microsoft Technologies, Seattle
John McAbel, WW Database/Cluster Product Manager, Beaverton
Paul McGrath, xSeries Market Management Services & Software, Raleigh
Gregg McKnight, Distinguished Engineer, xSeries Development Performance, Raleigh
Franklin Nanney, WW Service & Support Planning, Technical Support Marketing, Raleigh
Michael L. Nelson, eServer™ Solutions Engineering, Raleigh
Michael Parris, WW Service & Support Strategy, Technical Support Marketing, Raleigh
Bob Zuber, WW xSeries Product Marketing, Raleigh

Thanks to the following Microsoft employees for their contributions to this project:

Sandy Arthur
Al Bergstein
Eileen Brown
Tim Cerling
Jim Drennan
Ji-Young Kim
Thomas Kuklovszky
William Lyon
Bernd Pfann
Jim Steinmark

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this Redpaper or other Redbooks in one of the following ways:

- ▶ Use the online **Contact us** review redbook form found at:

ibm.com/redbooks

- ▶ Send your comments in an Internet note to:

redbook@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 662
P.O. Box 12195
Research Triangle Park, NC 27709-2195



Features of the x445

This chapter describes the IBM @server x445 server. We discuss the server's components and compare it to the x440. This chapter will assist in understanding the hardware and design of the server. This chapter includes the following topics:

- ▶ 1.1, "Overview" on page 2
- ▶ 1.2, "X-Architecture" on page 2
- ▶ 1.3, "Enterprise X-Architecture" on page 3
- ▶ 1.4, "NUMA architecture" on page 4
- ▶ 1.5, "XpandOnDemand" on page 7
- ▶ 1.6, "SMP Expansion Module" on page 8
- ▶ 1.7, "Memory implementation" on page 9
- ▶ 1.8, "IBM ServerProven®" on page 10
- ▶ 1.9, "Microsoft HCL" on page 11
- ▶ 1.10, "Comparing the x445 to the x440" on page 11

1.1 Overview

The next-generation scalable enterprise server, the xSeries 445, powered by Enterprise X-Architecture, leads the market for 8-way and 16-way processing with XpandOnDemand processor and I/O scalability, high performance, and OnForever™ availability for mission-critical database, ERP and CRM solutions.



Figure 1-1 The xSeries 445 server

The key features of the x445 are as follows:

- ▶ XA-32 second generation chipset.
- ▶ Models with the Intel Xeon MP processors, up to 2.8 GHz and 2 MB L3 cache, upgradable to 4-way, 8-way and 16-way. Future plans also have the x445 expandable 32-way.
- ▶ Entry-level models with Intel Xeon DP, up to 3.0 GHz, upgradable to 4-way.
- ▶ Active Memory with, Memory ProteXion, memory mirroring, hot-swap and hot-add memory.
- ▶ XceL4 Server Accelerator Cache, now 64 MB per 4-way (128 MB per 8-way system).
- ▶ 2 GB DDR memory standard, expandable to 64 GB per 8-way system.
- ▶ Six 64-bit Active PCI-X slots per chassis.
- ▶ Integrated LSI Logic 1030 dual-channel Ultra320 SCSI with integrated RAID-1 support.
- ▶ Integrated dual Gigabit Ethernet.
- ▶ Connectivity to an RXE-100 external PCI-X enclosure for an additional 12 PCI-X slots which can optionally be shared between two x445s.
- ▶ Remote Supervisor Adapter standard. Plans are to make the x445 upgradable to the Remote Supervisor Adapter II.

1.2 X-Architecture

By building robust capabilities into xSeries servers and working on new technologies for the future, IBM X-Architecture technology addresses these requirements and helps prepare you for the next wave of e-business. By modifying and using proven server technologies from the other successful IBM product lines, the xSeries servers are able to leverage these through IBM X-Architecture technology.

As introduced in 1998, IBM X-Architecture technology outlined capabilities for performance leadership, technology innovation, increased reliability and business-critical service and support capabilities. We continue to build IBM X-Architecture technology into our systems and set the bar for Intel processor-based servers across the industry.

The IBM xSeries technology agenda is straightforward: leverage proven IBM technologies to bring enterprise capabilities to the Intel processor-based server platform. The X-Architecture technologies can be grouped into the following categories.

- ▶ Core logic : advanced technologies that optimize system performance and scalability
- ▶ Enterprise storage solutions: designing subsystems and I/O bandwidth to match increasing storage and network requirements
- ▶ Availability: system, storage and clustering technologies to increase availability
- ▶ Systems management: technologies for comprehensive local and remote server and storage management within heterogeneous environments
- ▶ Technology-enabled service and support: tying advanced system and management technologies directly to service and support infrastructures

In short, the IBM X-Architecture technology is the IBM design blueprint for the Intel processor-based xSeries servers. We accelerated the advancement of xSeries servers to meet customer demands, as well as those of e-businesses. This has introduced the next step forward: Enterprise X-Architecture.

More information on X-Architecture technology can be found at:

<http://www.pc.ibm.com/us/eserver/xseries/xarchitecture/principals.html>

1.3 Enterprise X-Architecture

IBM continues to build on the X-Architecture blueprint with Enterprise X-Architecture.

Enterprise X-Architecture technologies yield revolutionary advances in the I/O, memory and performance of xSeries servers. This server design creates a flexible “pay as you grow” approach to buying high-end 32-bit and 64-bit xSeries systems that can be scaled quickly, easily and inexpensively. Enterprise X-Architecture technology is designed to achieve unparalleled levels of availability, scalability and performance in industry-standard enterprise computing.

Enterprise X-Architecture technology enables the following capabilities:

- ▶ XpandOnDemand scalability
- ▶ System partitioning
- ▶ PCI-X I/O subsystem
- ▶ Active PCI-X
- ▶ Remote I/O
- ▶ Active Memory
- ▶ High-speed (DDR) memory
- ▶ Memory ProteXion
- ▶ Chipkill™ memory
- ▶ Memory mirroring
- ▶ Hot-add/hot-swap memory
- ▶ XceL4 server accelerator cache

These features deliver application flexibility 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 unmatched “economies of scale” and new levels of server availability and performance.

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. 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 (processors, 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.

The Enterprise X-Architecture paradigm sets IBM xSeries servers apart in the industry while maintaining the advantages of compliance with industry standards for processors, memory, I/O, storage and software. Enterprise X-Architecture also establishes a revolutionary new economic model for servers through its flexible modular design. The x445 Enterprise X-Architecture-based server allows customers to XpandOnDemand, the ability to grow when you need to, using the existing infrastructure.

More Information on Enterprise X-Architecture technology can be found at:

<http://www.pc.ibm.com/us/eserver/xseries/xarchitecture/enterprise>

1.4 NUMA architecture

SMP designs are currently being challenged by the demand for additional Intel processors within a single system. In SMP, CPUs have equal access to a single shared memory controller and memory space. Because processor speed has dramatically increased to be much faster than memory, for most CPU intensive workloads the single front-side bus and single memory controller are often a bottleneck, resulting in excessive queuing delays and long memory latencies when under heavy load.

Processor instruction and data caches have been introduced to help reduce the likelihood of a memory controller bottleneck. A performance boost is often obtained from using large caches because of two reasons:

- ▶ Caches have lower latency and faster access time than a memory access. Because a cache can be accessed quicker than memory, the processor waits less time before receiving a response from a request.
- ▶ Caches also improve performance because they reduce queuing time of accesses which miss the caches and require a physical memory access.

For most commercial applications, process cache hit rates are usually greater than approximately 70%. In this case, the cache greatly reduces memory latency because most processor memory requests are serviced by the faster cache. The caches act as filters and reduce the load on the memory controller, which results in lower queuing delays (waiting in line) at the memory controller, thereby speeding up the average memory access time.

Another bottleneck in many SMP systems is the front-side bus. The front-side bus connects the processors to the shared memory controller. Process to memory requests travel across the front-side bus which can become overloaded when three or more high speed CPUs are added to the same bus. This, in turn, leads to a performance bottleneck and lower system scalability. Large processor caches also help improve performance because they assist in filtering many of the requests that must travel over the front-side bus (a processor cache-hit does not require a front-side bus for memory transaction).

However, even with a large L3 cache, the number of memory transactions which miss the cache is still so great that it often causes the memory controller to bottleneck; this happens when more than three or four processors are installed in the same system.

Non-uniform Memory Access (NUMA) is an architecture designed to improve performance and solve latency problems inherent in large (greater than four processors) SMP systems. The x445 implements a NUMA-based architecture and can scale up to 32 processors using multiple *SMP Expansion Modules*.

The SMP Expansion Modules each contain up to four CPUs and sixteen memory DIMMs. Each module also contains a dedicated local memory controller and 64 MB XceL4 Level 4 cache. The additional fourth level of cache greatly improves performance for the four processors in the SMP Expansion Module because it is able to respond to a majority of processor-to-memory requests, thereby reducing the load on the memory controller and speeding up average memory access times.

As shown in Figure 1-2, each SMP Expansion Module can be connected to other SMP Expansion Modules using three independent 3.2 GBps scalability links (shown in red). These scalability links mirror front-side bus operations to all other SMP Expansion Modules in the system and are key to building large multiprocessing systems.

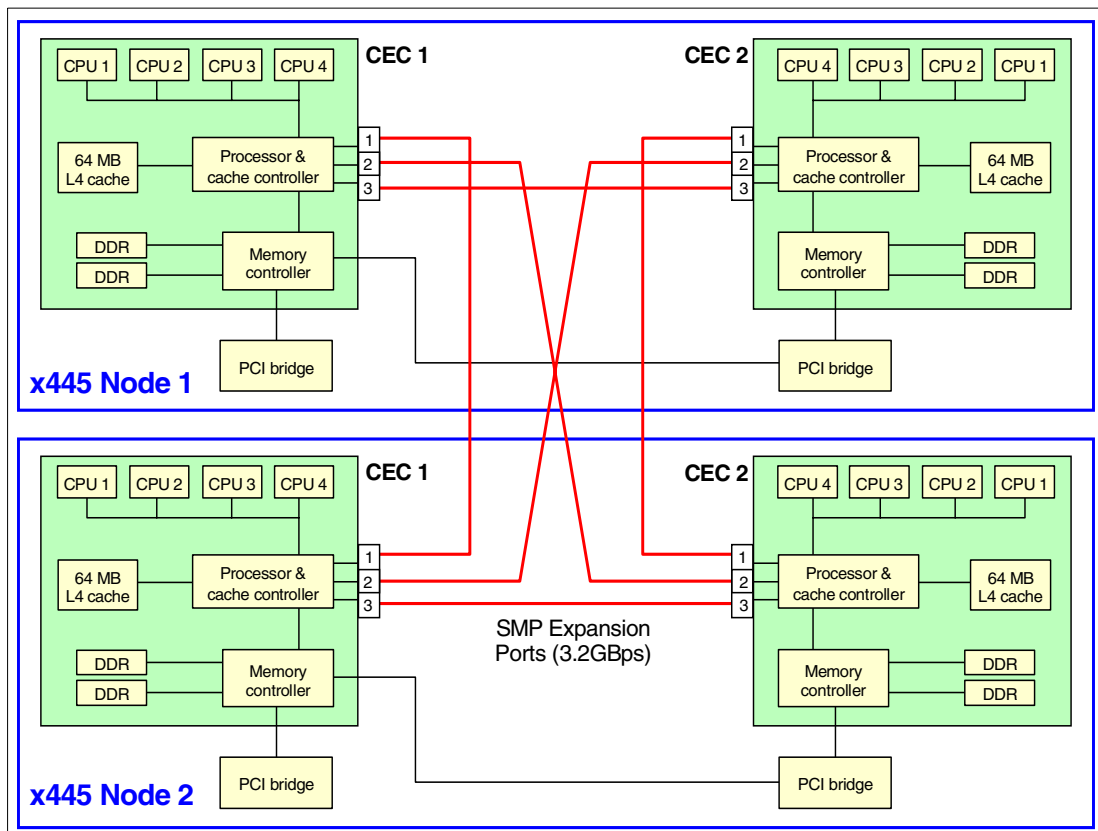


Figure 1-2 Two 8-way x445 servers connected as a 16-way configuration

By mirroring transactions on the front-side bus across the scalability links to other processors, the x445 is able to run standard SMP software. All SMP systems must perform processor-to-processor communication (also known as “snooping”) to ensure all processors receive the most recent copy of requested data. Since any processor can store data in a local cache and modify that data at any time, all processor data requests must first be sent to every

other process in the system so that each processor can determine if a more recent copy of the requested data is in that processor cache.

Snooping traffic is an important factor affecting performance and scaling for all SMP systems. The overhead of this communication becomes greater with an increase in the number of processors in a system. Also, faster processors result in a greater percentage of time spent performing snooping because the speed of the communications does not improve as the processor clock speed increases, since latency is largely determined by the speed of the front-side bus.

Its easy to see that increasing the number of processors and using faster processors results in greater communication overhead and memory controller bottlenecks. But unlike traditional SMP designs, which send every request from every processor to all other processors, greatly increasing snooping traffic, the x445 has a more optimal design. The XceL4 cache in the x445 improves performance because it filters most snooping operations.

The IBM XceL4 cache improves scalability with more than four processors because it also caches remote data addresses. So, before any processor request is sent across the scalability link to a remote processor, the memory controller and cache controller determine whether the request should be sent at all. To do this, each cache controller keeps a directory of all the addresses of all data stored in all remote processor caches. By checking this directory first, the cache controller can determine if a data request must be sent to a remote processor and only send the request to that specific SMP Expansion Module where the processor caching the requested data is located.

The majority of data requests are not found in the XceL4 cache and can be sent directly to the memory controller to perform memory look-up without any remote processor-to-processor communication overhead. The directory-based coherency protocol used by the XceL4 cache greatly improves scalability of the x445 because processor-to-processor traffic is greatly reduced.

Performance is also improved in an 8-way configuration. With two SMP Expansion Modules, the x445 has two EXA memory controllers and two 64-MB XceL4 caches. These two memory controllers share the memory request load from all eight processors and greatly reduce memory contention typically found in single memory controller SMP designs. Think of walking into a very busy airport to get your ticket. Wouldn't you rather see multiple ticket agents available to process the large number of travelers waiting ahead of you? Also, if there were an increase in the number of travelers, wouldn't you want to see more ticket agents to absorb the load? Similarly, the x445 employs a new ticket agent (memory controller) for each four processors added to the system. Therefore, a 16-way configuration would have four EXA memory controllers and a total of four 64 MB XceL4 caches to absorb the additional processor to memory request load.

By using this cache and providing fast local memory access to all processors, typical bottlenecks that are associated with large SMP systems using a single memory controller design are eliminated. Operating systems such as Windows 2000 perform well on this type of architecture, but dramatic performance gains can be achieved when the operating system is NUMA-aware, as with Windows Server 2003 Datacenter Edition.

The NUMA architecture of the x445 server and its support by Windows Server 2003 is further discussed in 3.3, "NUMA support in Windows Server 2003" on page 38.

1.5 XpandOnDemand

One of the challenges of optimizing server performance is to provide memory and I/O subsystems that allow the new processor architectures to operate fully optimized. Traditional industry-standard server designs begin to encounter performance bottlenecks beyond 4-way scaling, due to memory controller and front-side bus congestion. The EXA 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 EXA platform using enhanced, high-performance SMP building blocks to allow the effective scalability beyond 4-way SMP. These technologies provide scalability from two-way to 4-way, 8-way and 16-way systems, using the x445 server.

You can start out with one 4-way SMP node with six PCI-X slots. If it turns out that this is all the computing power you need for that server, you have saved yourself expense by not buying a more expensive 8-way (or larger) system. If you eventually need more processing power, you only need to add a second SMP Expansion Module to the existing x445 server to create an 8-way system.

To grow to 16-way without having to deploy a brand new server with operating system and applications, you can connect a second x445 8-way server to the original one to form a single-image 16-way system with 12 PCI-X slots. Need more PCI-X slots? Simply add the RXE100 external I/O enclosure. This total configuration is shown in Figure 1-3.

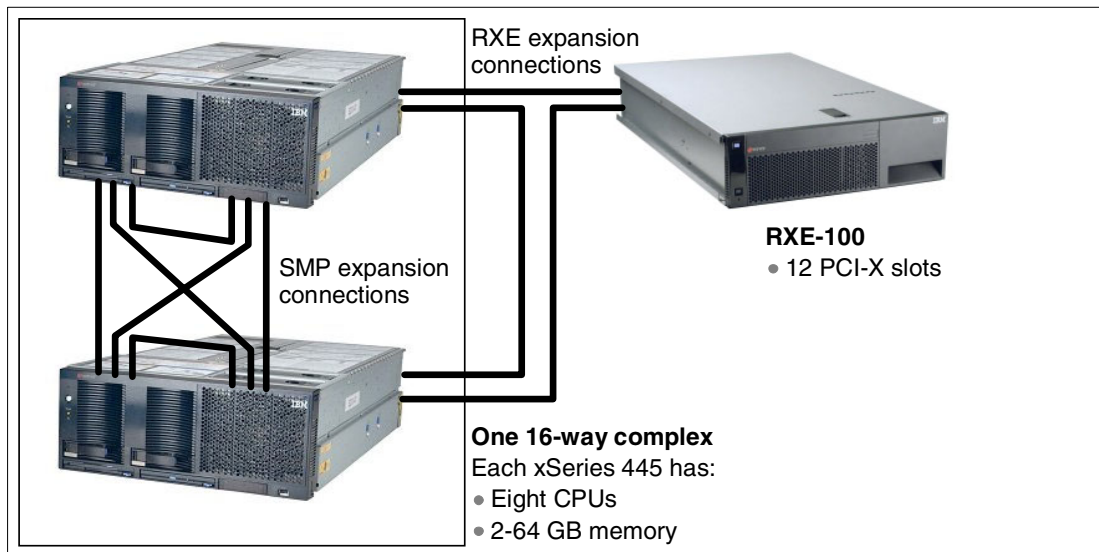


Figure 1-3 x445 XpandOnDemand to 16-way, 128 GB of RAM, and 18 PCI-X slots

With the EXA approach, adding a second SMP Expansion Module doubles the number of processors and the amount of system memory, but it does so without the bottlenecks. With the EXA design, you double the number of processors, chipsets, memory, memory controllers, and front-side buses, which helps spread the burden of the data traffic. This occurs with two SMP Expansion Modules. Connecting two x445 servers together translates into more system bandwidth and more system capacity. The contention and resource problems increase in a 16-way traditional SMP system but these are lessened with the x445 EXA design.

XpandOnDemand scalability represents revolutionary advances for Intel-based servers. You will be able to scale up server performance without any downtime while your initial purchase is secured.

1.6 SMP Expansion Module

The SMP Expansion Module, as shown in Figure 1-4, contains many of the major components for the x445 server:

- ▶ CPUs (two or four Xeon MP or two Xeon DP processors)
- ▶ XceL4 level four cache
- ▶ Memory controller
- ▶ System RAM (16 sockets)

The SMP Expansion Module is also referred to as the Central Electronics Complex (CEC). These terms are interchangeable.

All x445 servers ship with at least two processors and 2 GB RAM. The SMP Expansion Modules provide the NUMA-based architecture as discussed in 1.4, “NUMA architecture” on page 4. When configured as a two-way or 4-way Xeon MP server, the x445 will contain one SMP Expansion Module. When configured as an 8-way server, the x445 is configured with two SMP Expansion Modules. The two SMP Expansion Modules are installed one above the other in the server.

Xeon DP models: When configuring the x445 using the Xeon DP processors, a two-way Xeon DP x445 server will contain one SMP Expansion Module. When configured using four Xeon DP processors, the x445 contains two SMP Expansion Modules.




Figure 1-4 xSeries 445 SMP Expansion Module

Each SMP Expansion Module is also equipped with LEDs for light path diagnostics. LEDs are available for each memory DIMM and for each processor to aid in the quick identification of failed or failing components. LEDs are also located on the top of the SMP Expansion Module to identify the operational components of Active Memory.

1.7 Memory implementation

The x445 implements NUMA by placing the DIMM sockets in the SMP Expansion Modules, meaning that in a 16-way configuration, there are four memory “local regions” of memory.

The sixteen sockets in each SMP Expansion Module are divided into four banks and two ports. A bank is the pair of DIMMs required for two-way interleaving. The ports and banks are shown in Figure 1-5.

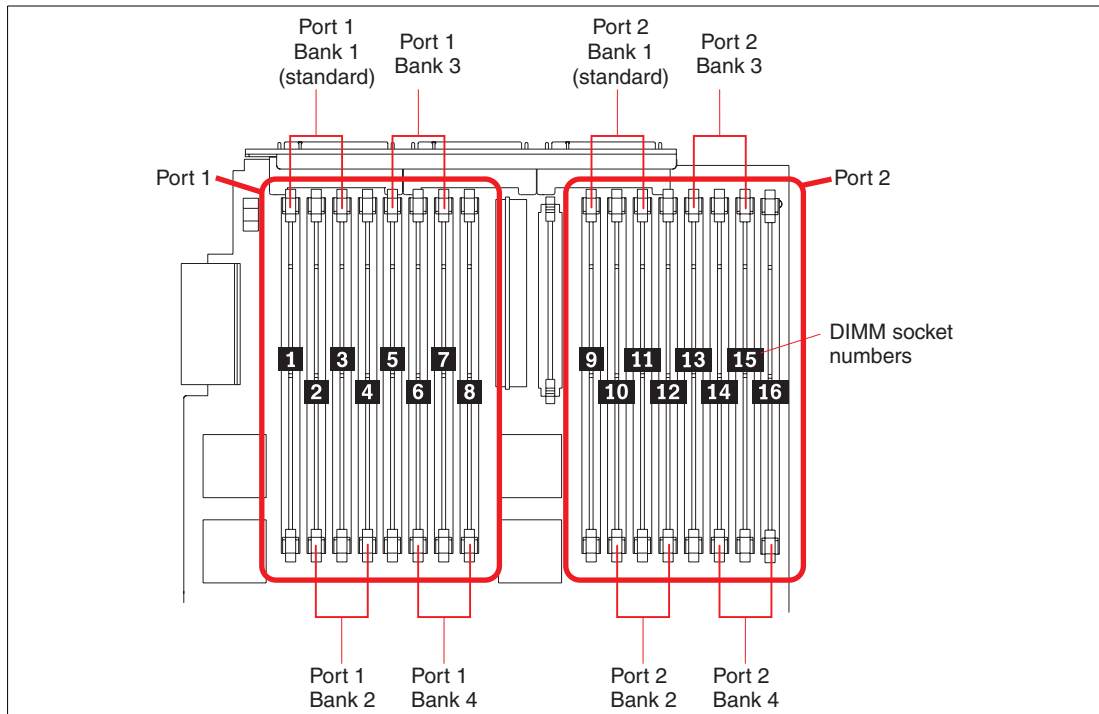


Figure 1-5 Memory DIMM sockets on the x445 SMP Expansion Module

There are a number of advanced features implemented in the x445 memory subsystem, collectively known as *Active Memory*:

- Memory ProteXion

Memory ProteXion, also known as “redundant bit steering”, is the technology behind using redundant bits in a data packet to provide backup in the event of a DIMM failure.

In the event that a chip failure on the DIMM is detected, the memory controller can re-route data around that failed chip through the spare bits (this is similar to the hot-spare drive of a RAID array). It can do this automatically without issuing a Predictive Failure Analysis® (PFA) or light path diagnostics alert to the administrator. After the second DIMM failure, PFA and light path diagnostics alerts would occur on that DIMM as normal.

- Memory scrubbing

Memory scrubbing is an automatic daily test of all the system memory that detects and reports memory errors that might be developing before they cause a server outage.

When a bit error is detected, memory scrubbing determines whether or not the error is recoverable. If it is recoverable, Memory ProteXion is enabled and the data that was stored in the damaged location is rewritten to a new location. The error is then reported so that preventative maintenance can be performed. As long as there are enough good locations to allow the proper operation of the server, no further action is taken other than recording the error in the error logs.

If the error is not recoverable, then memory scrubbing sends an error message to the light path diagnostics, which then turns on the proper lights and LEDs to guide you to the damaged DIMM. If memory mirroring is enabled, then the mirrored copy of the data in the damaged DIMM is used until the system is powered down and the DIMM is replaced. If hot-add is enabled in the BIOS then no rebooting is required and the new DIMM will be enabled immediately.

- ▶ **Memory mirroring**

Memory mirroring is roughly equivalent to RAID-1 in disk arrays, in that memory is divided in two ports and one port is mirrored to the other half. If 8 GB is installed, then the operating system sees 4 GB once memory mirroring is enabled (it is disabled in the BIOS by default). Since all mirroring activities are handled by the hardware, memory mirroring is operating system independent.

- ▶ **Chipkill memory**

Chipkill is integrated into the XA-32 second-generation chipset and does not require special Chipkill DIMMs. Chipkill corrects multiple single-bit errors to keep a DIMM from failing. When combining Chipkill with Memory ProteXion and Active Memory, the x445 provides very high reliability in the memory subsystem. Chipkill memory is approximately 100 times more effective than ECC technology, providing correction for up to four bits per DIMM (eight bits per memory controller), whether on a single chip or multiple chips.

If a memory chip error does occur, Chipkill is designed to automatically take the inoperative memory chip offline while the server keeps running. 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. The controller is able to reconstruct the “missing” bit from the failed chip and continue working as usual.

Chipkill support is provided in the memory controller and implemented using standard ECC DIMMs, so it is transparent to the OS.

- ▶ **Hot-swap and hot-add memory**

Two new features have been added to the x445: hot-swap and hot-add memory. These technologies are mutually exclusive; enable one or the other through the system BIOS (followed by a reboot).

The hot-add memory feature enables you to add DIMMs without turning off the server. This feature is supported only in those servers using Windows Server 2003 Enterprise or Datacenter Editions. The current implementation is restricted such that one of the two ports must be completely empty before you can add DIMMs. DIMMs must be added in pairs, and you can also only add memory to the upper SMP Expansion Module.

The hot-swap memory feature allows you to replace DIMMs of the same type, size, and clock speed as the failed DIMM without turning off the server. Hot-swap memory is operating system independent. Before you can use the hot-swap feature, you must enable memory mirroring. You can also only replace DIMMs that are installed in the upper SMP Expansion Module.

1.8 IBM ServerProven®

The IBM ServerProven program was first announced in 1996 to ensure that leading vendor products were compatible with the then Netfinity® servers and today, the xSeries range. When a product enters the ServerProven program, it must pass a rigorous testing and certification process. Once this process has completed successfully, the vendor will achieve the IBM ServerProven status for their product. The IBM ServerProven program includes IBM and non-IBM hardware, software and middleware.

When you see the ServerProven logo, you can be assured that the product displaying this emblem has been tested and certified to run reliably on the IBM xSeries Servers. This allows IBM customers to purchase a complete, verified and tested solution. While ServerProven reduces installation problems and the need for post-installation support, ultimately warranty and software obligations are the responsibility of the vendor, not IBM.

For additional information on IBM ServerProven program participants, refer to the following Web site:

<http://www.pc.ibm.com/us/compat/serverproven>

For additional information on the IBM ServerProven program, refer to following Web site:

<http://www.pc.ibm.com/ww/eserver/xseries/serverproven>

IBM technicians have worked hand in hand with their Microsoft counterparts to ensure that the IBM x445 server and its most popular hardware and software configurations have been ServerProven with the Windows Datacenter 2003 operating system. This is essential to ensure that the IBM Datacenter Solution Program has industry-leading hardware and software solutions to provide the highest level of availability for today's mission-critical environments

1.9 Microsoft HCL

The Windows Datacenter Server Hardware Compatibility List (HCL) is a listing of servers from approved Datacenter OEMs such as IBM, that have passed through a rigorous testing process to ensure they meet Datacenter Server compatibility requirements.

The hardware testing that IBM must complete ensures to the customer that the following components will work reliably together in the Datacenter solution.

- ▶ All hardware components and all hardware drivers.
- ▶ Applications or utilities that have kernel-mode software, primarily anti-virus software, NDIS Intermediate drivers and similar types of software.
- ▶ Backup, encryption, or other types of products with kernel mode drivers.

Further information on the Microsoft HCL criteria for Microsoft Datacenter can be found at:

<http://www.microsoft.com/whdc/hwdev/platform/server/datacenter/DCHCT.msp>

1.10 Comparing the x445 to the x440

The x440 was the highly successful predecessor to the x445 server. It is therefore important to understand the differences between these two models. Refer to Table 1-1 on page 12 for a comparison of the x445 server with the x440.

Table 1-1 Comparison of the IBM x440 versus IBM x445

Feature	x440 server	x445 server
IBM chipset	First generation XA-32	Second generation XA-32
SMP Scalability	<ul style="list-style-type: none"> ▶ Two-way to 16-way SMP 	<ul style="list-style-type: none"> ▶ Two-way to 16-way SMP ▶ Future 32-way support
Processors	<ul style="list-style-type: none"> ▶ Intel Xeon MP 1.5 GHz/ 1 MB, 1.9 GHz/1 MB and 2.0 GHz/2 MB ▶ Intel Xeon DP 2.4 GHz/ 512 MB (up to 4-way SMP only) 	<ul style="list-style-type: none"> ▶ Intel Xeon MP 2.0 GHz/ 1 MB, 2.5 GHz/1 MB and 2.8 GHz/2 MB ▶ Intel Xeon DP 3.0 GHz/ 512 MB (up to 4-way SMP only)
Xcel4 cache	<ul style="list-style-type: none"> ▶ 32 MB Xcel4 Server L4 Accelerator Cache per SMP Expansion Module ▶ 64 MB Xcel4 cache in 8-way configuration 	<ul style="list-style-type: none"> ▶ 64 MB Xcel4 Server L4 Accelerator Cache per SMP Expansion Module ▶ 128 MB Xcel4 cache in 8-way configuration
Memory capacity	<ul style="list-style-type: none"> ▶ 16 GB RAM per SMP Expansion Module ▶ Up to 32 GB RAM (64 GB in a 16-way) 	<ul style="list-style-type: none"> ▶ 32 GB per SMP Expansion Module ▶ Up to 64 GB (128 GB in a 16-way)
Type of memory used	<ul style="list-style-type: none"> ▶ PC133 SDRAM Memory ▶ 4-way Interleaving 	<ul style="list-style-type: none"> ▶ DDR Memory ▶ Two-way Interleaving
Memory features	<ul style="list-style-type: none"> ▶ Active Memory includes: <ul style="list-style-type: none"> – Chipkill memory – Memory mirroring – Memory ProteXion (redundant bit steering) 	<ul style="list-style-type: none"> ▶ Active Memory includes: <ul style="list-style-type: none"> – Hot-add memory – Hot-swap memory – Chipkill memory – Memory mirroring – Memory ProteXion (redundant bit steering)
Internal hard drive capacity	Two Ultra160 hard drives	Two Ultra320 hard drives
SCSI /RAID	<ul style="list-style-type: none"> ▶ Adaptec Ultra160 dual channel SCSI ▶ Optional RAID 	<ul style="list-style-type: none"> ▶ LSI 1030 dual channel Ultra320 SCSI ▶ Integrated RAID-1 (drive mirroring)
Ethernet	<ul style="list-style-type: none"> ▶ Single port ▶ Broadcom BCM5700 10/100/1000 BASE-T 	<ul style="list-style-type: none"> ▶ Dual port ▶ Broadcom BCM5704 10/100/1000 BASE-T
PCI-X capacity	<ul style="list-style-type: none"> ▶ Six active (hot-swap) PCI-X adapter slots ▶ 66 MHz-133 MHz, 64bit 	<ul style="list-style-type: none"> ▶ Six Active (hot-swap) PCI-X adapter slots ▶ 66MHz-133MHz, 64bit
RXE-100 support	<ul style="list-style-type: none"> ▶ Supported attachment to One x440 only. Sharing an RXE among two x440 servers not available 	<ul style="list-style-type: none"> ▶ Supported sharing between two x445s ▶ Can be shared between 2x 4-ways, 2x 8-ways and 2x 16-way, x445 servers ▶ Maximum one RXE-100 for each x445 server
CD-ROM	CD-ROM	DVD-ROM
Rack size	4U rack optimized	4U rack optimized



Features of Windows Server 2003, Datacenter Edition

This chapter describes Windows Server 2003, Datacenter Edition, its requirements and its features. While our focus is the 32-bit version of Datacenter Edition that runs on the IBM xSeries 445 Server, there is some limited mention of the 64-bit version of the product.

Most of the features we discuss are either new to Windows Server 2003, or improvements over the Windows 2000 version. They fall into categories such as Directory Services, file and disk management, network and communications, clustering, and performance. In many cases, we provide high level examples of usage to help you understand the positive impact of the Windows Server 2003 release. We will also compare Datacenter Edition features to those of Enterprise, Standard, and Web Editions.

The topics covered in this chapter are:

- ▶ 2.1, “Description of Windows Server 2003, Datacenter Edition” on page 14
- ▶ 2.2, “Why Datacenter?” on page 23, discussing compelling reasons for moving to Datacenter Edition
- ▶ 2.3, “High availability” on page 23, discusses the role of IBM Enterprise Services for Microsoft Technologies (ESMT)
- ▶ 2.4, “Clustering” on page 25 describes the types of clusters and equipment typically used to implement Datacenter Edition clustering
- ▶ 2.5, “Server consolidation” on page 27 discusses the roles of IBM and how Datacenter can be used in a server consolidation environment
- ▶ 2.6, “The Datacenter High Availability Program” on page 30 explains the difference between a Windows Datacenter installation compared to other Windows operating systems.

Also of note is Windows System Resource Manager (WSRM), which is described in the Redpaper *Implementing Windows System Resource Manager*, REDP3701.

2.1 Description of Windows Server 2003, Datacenter Edition

There are four members of the 32-bit Windows Server 2003 family:

- ▶ Windows Server 2003, Datacenter Edition
- ▶ Windows Server 2003, Enterprise Edition
- ▶ Windows Server 2003, Standard Edition
- ▶ Windows Server 2003, Web Edition

Although our focus is Datacenter Edition, we will discuss many of the new features of Windows Server 2003 and enhanced features of Windows 2000 Server since Datacenter Edition uses most of them. We will also review Datacenter Edition-specific features.

Microsoft also removed a number of features from Windows 2000 Datacenter Server as the company made refinements and improvements while developing the Windows Server 2003 version. These features were deemed to waste system capacity or were considered inappropriate in a high-end server operating system.

2.1.1 Windows Server 2003 new and updated features

If you consider the usefulness or application of the features below, you will note that Microsoft has invested great effort to make Windows Server 2003 a first class product. Microsoft has developed a series of white papers that detail the planning, deployment, and use of these features:

<http://www.microsoft.com/windowsserver/techinfo/overview>

Directory services features

Microsoft released the Active Directory with Windows 2000. This next version adds functionality in several areas:

- ▶ Cross-forest trust

The Active Directory is the directory service for Windows Server 2003. It is included in all editions except for Web Edition. It stores information about objects on the Windows network.

The Active Directory now allows forest-to-forest trusts. Once set up, all domains in Forest X will trust all domains in Forest Y, for example, and Forest Y will trust all domains in Forest X. However, even if Forest Y trusts Forest Z, Forest X does not trust Forest Z since trusts are not transitive.

- ▶ Installation of a replica from media

The source of an initial replication can now be created from files created during the backup of a Domain Controller (DC) or Global Catalog server and used on a target DC using a CD, DVD, tape, or file over the network.

You are ready, for example, to run DCPROMO to create your new remote site Domain Controller. With Windows Server 2003, you no longer need a live connection back to headquarters. Just insert the CD that you created prior to leaving and you are set. With the ability to install a replica from various media types, you do not have to worry about that slow, busy pipe back to HQ.

- ▶ Use of Active Directory Migration Tool (ADMT) with Windows Server 2003 now included

The Active Directory Migration Tool allows password migration from Windows NT 4.0 or Windows 2000 to Windows 2000/2003 Server family domains. Through batch file scripting or command line, administrators can migrate users, groups, and computers. The Web Edition does not utilize ADMT.

You could, for example, create a trust between two domains, then use ADMT to copy machine accounts and users from domain A to domain B. There are plenty of rules involved and while the source domain can be NT 4.0, the target domain must be Windows 2000/2003. When you run ADMT, you can select wizards to migrate User, Group, and Service accounts. There is also a Trust Migration wizard and more.

- ▶ **Metadirectory Services Support (MMS)**

MMS augments the Active Directory. It is used for identity management and allows large organizations to have a single logical view of an organization's identity information. It is available on Windows Server 2003, Enterprise and Datacenter Edition.

MMS is a stand-alone product and service providers are typically engaged to deliver this product. It is used to pull data from various operating systems and applications such as Active Directory, Novell eDirectory, Domino™, Exchange 5.5, SQL Server, DB2®, and various flat files to create a single logical view that is centrally administered.

- ▶ **Lightweight Directory Access Protocol 3 (LDAP 3) support**

LDAP is the primary access protocol for the Active Directory and LDAP V3 is now supported. The Windows Server 2003 Active Directory is LDAP-based.

File or storage management services features

There are a number of new file and storage features in Windows Server 2003. Shadow copies and the File Replication Service will no doubt become popular for many administrators.

- ▶ **Volume Shadow Copy Service (VSS)**

This is a general infrastructure for creating point in time copies of data. SAN vendors can now write standardized code through hooks into Windows Server 2003, which will allow terabytes of data to be copied across the SAN in minutes for mounting by another server that could process the data for data mining, backup, or development testing.

- ▶ **Storage Area Network Support**

Windows Server 2003, Enterprise and Datacenter Editions do not automatically mount visible logical unit numbers (LUNs). This reduces the potential for data corruption. This feature is turned on by default in the Enterprise and Datacenter Editions. It is turned off by default in the Web and Standard Editions.

- ▶ **Open file backup**

Microsoft's backup utility now supports open file backup. Backup uses shadow files to implement this feature. In Windows 2000, you had to close files before you backed up.

- ▶ **Shadow copies for shared folders**

This is a new feature that provides point in time copies for network folders to allow users to access previous versions of files. To use it, simply right-click a file while in Explorer. You can set up volume shadow copying to occur at intervals throughout the day.

- ▶ **File Replication Service (FRS)**

This is a new, full-featured MMC snap-in to control file replication. If you need headquarters and six remote sites to see the same data and yet you do not want all sites to pull the data across the WAN, you can use FRS to create identical replicas so that each remote site pulls data from the locally created replica.

- ▶ **Folder redirection of My Documents**

You can redirect a user's My Documents folder to their home directory. Additionally, you can have your redirected folders in just one network location since the Group Policy can now set it.

- ▶ Command line interface

Many new command line utilities have been added, including the ability to open a Telnet session and issue commands on a remote machine. Windows Server 2003 has greatly increased command line usage and functionality.

- ▶ Encrypting File System (EFS)

You can encrypt your files on NTFS and authorize others (agents) to access those files. You can also encrypt off-line files and store encrypted files in Web folders. Use the **ESFINFO** command if you forget who your recovery agent is.

- ▶ Virtual Disk Service

This is a new set of command line utilities and management applications that will allow Windows to manage any vendor's SAN disk. Various vendors of SAN hardware can allow disk management through a uniform interface in Windows Server 2003 when they write a VDS "provider".

Clustering features

While clustering can now scale significantly in Windows Server 2003, there are other features that will save administrators a great deal of time, such as driver propagation in print server clustering. No more manual propagation of print drivers. With the Datacenter Program for Datacenter Edition, clustering customers should be confident that their solution will deliver high availability.

- ▶ Cluster Services

With Windows Server 2003 32-bit clustering, both the Enterprise and Datacenter Editions support eight nodes (servers). Datacenter Edition scales in memory to 64 GB per node while Enterprise Edition is limited to 32 GB. Datacenter Edition scales to 32 CPUs per node while Enterprise is limited to eight CPUs.

Note: This is different from Network Load Balancing, which supports a 32 node maximum and is generally restricted to functions such as Web serving. NLB and server clustering are not supported on the same server.

- ▶ Print server clustering

Windows Server 2003 driver propagation is now automatic for all nodes of the cluster.

With Windows 2000, you had to manually carry out this task. With Windows Server 2003, it is ensured that the clients will be able to get the needed driver from the particular node to which they are attached.

- ▶ Network Load Balancing (NLB)

All versions of Windows Server 2003 come with NLB. It allows for a way to distribute incoming TCP/IP traffic among members (servers) of the NLB cluster.

NLB, formerly Windows Load Balancing Service (WLBS), was previously only available on the Advanced Server and higher versions of Windows 2000. Now, even the standard version of Windows Server 2003 has this functionality. Besides unique IP addresses for each Web server, for example, all NLB servers share a common virtual IP that allows NLB to distribute incoming traffic to provide a connection-oriented form of load balancing.

NLB is intended to serve front-end applications such as Web-based or Internet applications and servers. Server Clusters, formerly known as MSCS in previous versions of Windows Server, is intended for back-end applications and services such as databases, ERP, and messaging.

Cluster Services or "Server Clusters", not NLB, is our primary focus when we discuss clustering in 2.4, "Clustering" on page 25. There is a 32 server maximum limit for NLB.

Networking and communications features

There has been a wide range of feature improvements in the area of networking and communications. One that will prove important over time will be IPV6.

- ▶ 802.1x wireless network support

All versions of Windows Server 2003 support wireless packet correction, performance and other enhancements over Windows 2000. There is now Wireless 802.1x and Bluetooth printing support. Bluetooth is a short range wireless connectivity technology. Both IBM and Microsoft are Promoter members of the Bluetooth Special Interest Group.

- ▶ Connection Manager

This provides support of local and remote access to your service provider using a network of access points.

In Windows Server 2003, there is an enhanced Connection Manager Administration Kit (CMAK). In Windows 2000, it allowed you to predefine connection profiles for remote access. Under Windows Server 2003, administrators can now provide more than one VPN server for connections, and there are many other features including split-tunneling to allow client-side VPN connections to route intranet-based traffic over the VPN connection while isolating Internet traffic to a client's local connection.

- ▶ Internet Protocol Version 6 (IPV6)

IPV6 Internet standard protocols are the next suite of network layer protocols for the Internet. This is 128-bit addressing and is available as a networking option. IPV6 resolves many of the problems of IPV4 such as address depletion and extensibility. Windows Server 2003 also supports coexistence with IPV4.

- ▶ e-mail services (POP3, SMTP)

A new feature, POP3 services provide e-mail transfer and retrieval services. Use POP3 to store and manage basic e-mail accounts.

This is a free mail server which does not compete with Exchange. It offers the most basic functionality but would be useful to small establishments. Windows NT and 2000 did not offer this.

- ▶ Internet Authentication Service (IAS)

The proxy component of IAS supports the ability to separate the authentication and authorization of connection requests. The IAS proxy can forward the user authentication to an external RADIUS (Remote Authentication Dial-in User Service) server for authentication and perform its own authentication using a user account in the active directory. With Windows Server 2003, Standard Edition or higher, you can have up to 50 RADIUS servers and a maximum of two RADIUS server groups and unlimited users.

- ▶ Layer Two Tunneling Protocol (L2TP)

This is a more secure version of PPTP, Point-to-Point Tunneling Protocol. It is used for tunneling and address assignment.

- ▶ Point-to-Point Protocol over Ethernet (PPPoE)

This new feature allows you to make broadband connections to certain ISPs without the need for additional software. Now when you use the Demand-Dial wizard to select your connection type, such as Connect Using a Modem, the new option is there.

- ▶ Virtual Private Networking

Users can have VPN connections to the company network through a secure tunnel. Windows Server 2003, Web Edition can only accept one VPN connection at a time. Windows Server, Standard Edition can accept 1 000 concurrent VPN connections through

the ports. Windows Server 2003, Enterprise and Datacenter Editions support unlimited VPN connections. Your publicly connected VPN session will be secure and encrypted.

Performance features

Customers can expect improved performance when moving to Windows Server 2003 from Windows 2000, even when utilizing the same hardware. Also, taking into account the evolution of hardware, enterprises should be able to serve their clients and customers much more effectively since their servers can handle larger workloads, or handle existing workloads more effectively. One new key feature, WSRM, stands out as one of the most powerful new tools in Windows Server 2003.

- ▶ Windows System Resource Manager (WSRM)

WSRM allows administrators to set CPU and memory utilization for specific processes or applications.

A good example of use is as follows. If you were to consolidate three applications onto one server, you might find that each application requires more resources at different times. With WSRM, you could schedule a critical customer order entry program to get 70% of CPU resources during business hours while limiting a Human Resources application to 20%. Then, at night, you could allocate 80% of CPU resources to your batch processing applications while setting the customer order program and the Human Resources program to share remaining system resources.

Departments within an organization that fund server purchases could be shown accounting policies through WSRM that show their application being guaranteed a specific amount of system resources. So, if Purchasing paid for 50% of a new server purchase, they can be guaranteed 50% of the server resources through use of WSRM policy settings.

WSRM is discussed in more detail in the Redpaper *Implementing Windows System Resource Manager*, REDP3701.

- ▶ Hyper-Threading

Windows Server 2003 takes advantage of Hyper-Threading capable processors to provide a significant performance improvement. The xSeries 445 with Windows Server 2003 supports Hyper-Threading.

Hyper-Threading allows multi-threaded software applications running on Windows Server 2003 to execute threads in parallel. While Windows 2000 supported Hyper-Threading, the operating system was not designed to take advantage of it as Windows Server 2003 does. As more applications are written to take advantage of Hyper-Threading, customers will benefit further.

- ▶ NUMA (Non-uniform Memory Access) Architecture

Windows Server 2003 supports NUMA architecture. This architecture is built into the IBM xSeries 445 Server. NUMA is designed to reduce the single memory controller bottleneck that results from multiple CPUs sharing access to a single memory controller. In the NUMA design, systems scale by adding SMP expansion modules and by adding large low-latency fast access memory caches.

Each of the SMP Expansion Modules contains up to four CPUs and 16 memory DIMMs. Each module also contains a dedicated local memory controller and a 64 MB Xcel4 Level 4 cache. The additional fourth level of cache greatly improves performance for the four processors in the SMP Expansion Module because it is able to respond quickly to a majority of processor-to-memory requests, thereby reducing the load on the memory controller and speeding up average memory access times.

- ▶ Networking performance improvements:
 - Reduced CPU usage required for sending and receiving provides more CPU capacity for applications and services.
 - Task offload (checksum offload and large send offload) is now on by default. This can have a significant impact for systems equipped with Gigabit and 10 Gigabit Ethernet adapters, providing the same beneficial effect as above.
 - TCP/IP stack now auto-negotiates RFC 1323 TCP options for window scaling, timestamps, etc., so that there is a better use of net bandwidth.
 - Network stack now also sets the size of the default TCP receive window based on the speed of the underlying link. Clients, or attached servers, get the best throughput possible, not just whatever the average is of all clients.
 - Windows Server 2003 is the first server release that ships with the WinSock TransmitFile call. This call works with the operating system cache scheme, so it speeds up file transfers considerably if the file is already cached.
 - WinSock Direct is now included in all members of the Windows Server 2003 family.
 - Support for IPV6 is provided.

For more information about the TCP/IP improvements with Windows Server 2003, see the paper *Microsoft Windows Server 2003 TCP/IP Implementation Details*, available from:

<http://www.microsoft.com/technet/prodtechnol/windowsserver2003/plan/TCP03.asp>

2.1.2 Requirements

The requirements for Windows Server 2003, Datacenter Edition fall into several categories. The IBM xSeries 445 server meets or greatly exceeds these requirements.

Datacenter Edition must run only on systems specified in the Microsoft Hardware Compatibility List (HCL). The list of IBM hardware is available by searching on Manufacturer for "IBM" at:

<http://www.microsoft.com/windows/catalog/server/>

or directly (in the PDF of this Redpaper, the link is active):

<http://www.microsoft.com/windows/catalog/server/default.aspx?subID=22&xs1t=search&pgn=b9da458c-f7a0-43c0-a7a6-02127c20a282&maxrows=0&sortcol=server2003&sortdir=descending&qu=IBM&scope=2&btnSearch=Go>

Specific requirements by the operating system for hardware components are as follows. The requirements by the applications that are to run on the server are in addition to these requirements.

- ▶ Multiprocessor support
 - Minimum 8-way capable machine required.
 - Maximum 32-way with the 32-bit version and on the 64-bit version, a maximum of 64 processors in a single partition and license up to 128 processors.
- ▶ CPU speed
 - Minimum: 400 MHz for x86-based servers.
 - Minimum: 733 MHz for Itanium-based servers.
 - Recommended: 733 MHz or better.
- ▶ Memory
 - Minimum: 512 MB.

- Recommended: 1 GB or better.
 - Maximum: 64 GB for x86-based servers on the 32-bit version.
 - Maximum: 512 GB physical memory or up to 16 TB virtual memory for Itanium-based servers on the 64-bit version.
- ▶ Disk space
- 1.5 GB for x86-based (32-bit) servers.
 - 2 GB for Itanium-based (64-bit) servers.

Microsoft support of Datacenter Edition on IBM hardware is ensured since the Microsoft Hardware Compatibility List (HCL) lists IBM equipment that has passed rigorous testing (by IBM and Microsoft) and is certified by Microsoft. Additionally, on the IBM ServerProven site, IBM lists all servers that are certified for use with Windows Datacenter Editions. For details, see:

<http://www.pc.ibm.com/compat>

2.1.3 Comparing Datacenter to Enterprise and Standard Editions

Table 2-1 on page 21 provides a comparison between Windows Server 2003, Datacenter, Enterprise, Standard, and Web Edition 32-bit versions. Because the table displays all versions, you will better understand what separates the four versions and what is specific to Datacenter Edition.

Windows Server 2003, Datacenter Edition offers the ability to provide server clustering and scale in memory and processors beyond any other current or previous version of Windows Server. It is also part of a program designed to provide the highest level of availability and support.

Please note that while Enterprise Edition seems similar to Datacenter Edition in terms of feature sets and the ability to scale to eight cluster nodes (servers), there is a distinct difference in terms of the ability to scale in memory and processors, and with Datacenter there is a complete high availability program designed around it.

Enterprise Edition is limited to 32 GB of memory per node while Datacenter Edition can scale to 64 GB. Datacenter surpasses Enterprise Edition in SMP support since it can scale to 32 processors versus only eight for Enterprise Edition.

Table 2-1 Feature comparison of the 32-bit Windows Server 2003 family

Features	Datacenter Edition	Enterprise Edition	Standard Edition	Web Edition
Scalability				
Maximum RAM memory (GB)	64	32	4	2
Maximum CPU	32	8	4	2
Maximum Server Cluster Nodes	8	8	0	0
Datacenter Program	Yes	No	No	No
Hot-add memory	Yes	Yes	No	No
NUMA support	Yes	Yes	No	No
Directory Services				
Active Directory	Yes	Yes	Yes	Limited
Metadirectory support	Yes	Yes	No	No
Security Services				
Internet Connection Firewall	No	Yes	Yes	Yes
PKI, Certificate Services, Smart Cards	Yes	Yes	Limited	Limited
Terminal Services				
Remote Desktop for Administration	Yes	Yes	Yes	Yes
Terminal Server	Yes	Yes	Yes	No
Terminal Server Session Directory	Yes	Yes	No	No
Clustering Technologies				
Network Load Balancing	Yes	Yes	Yes	Yes
Cluster Service	Yes	Yes	No	No
Communications and Networking				
VPN Support	Yes	Yes	Yes	Limited
Internet Authentication Service (IAS)	Yes	Yes	Yes	No
Network Bridge	Yes	Yes	Yes	No
Internet Connection Sharing (ICS)	No	Yes	Yes	No
IPV6	Yes	Yes	Yes	Yes
File and Print Services				
Distributed File System (DFS™)	Yes	Yes	Yes	Yes
Encrypting File System (EFS)	Yes	Yes	Yes	Yes
Shadow Copy Restore	Yes	Yes	Yes	Yes
Removable and Remote Storage	Yes	Yes	Yes	No
Fax Service	Yes	Yes	Yes	No

Features	Datacenter Edition	Enterprise Edition	Standard Edition	Web Edition
Services for Macintosh	Yes	Yes	Yes	No
Management Services				
IntelliMirror	Yes	Yes	Yes	Limited
Group Policy Results	Yes	Yes	Yes	Limited
Windows Management Instrumentation (WMI) Command Line	Yes	Yes	Yes	Yes
Remote OS Installation	Yes	Yes	Yes	Yes
Remote Installation Services (RIS)	Yes	Yes	Yes	No
Microsoft .NET Application Services				
.NET Framework	Yes	Yes	Yes	Yes
Internet Information Services (IIS) 6.0	Yes	Yes	Yes	Yes
ASP.NET	Yes	Yes	Yes	Yes
Enterprise UDDI Services	Yes	Yes	Yes	No
Multimedia Services				
Windows Media Services	Yes	Yes	Yes	No

2.1.4 Features not available in Windows Server 2003, 64-bit editions

There are two versions of Windows Server 2003, 64-bit editions: Enterprise and Datacenter Editions. Some features included in many of the 32-bit versions are not included in the 64-bit versions. This is due to architectural differences or because they were not appropriate for the 64-bit environment.

The following features are not included in 64-bit versions:

- ▶ 16-bit, DOS-based subsystem
- ▶ ACPI 2.0 (ACPI 2.0 64-bit fixed tables are supported)
- ▶ ASP .NET State Service
- ▶ Compressed (zipped) folders
- ▶ DirectMusic
- ▶ DVD video playback support
- ▶ Enterprise Memory Architecture
- ▶ Fast User Switching
- ▶ Fax support
- ▶ Hot Add Memory
- ▶ IEEE 1394 audio support
- ▶ Internet Connection Sharing (ICS)
- ▶ Internet Connection Firewall (ICF)
- ▶ Internet Locator Service (ILS)

- ▶ Earlier transports:
 - IPX/SPX LAN and WAN
 - Client Service for NetWare
 - Services for Macintosh
 - NetBIOS
 - Open Shortest Path First (OSPF)
 - Simple Network Management Protocol (SNMP) over IPX/SPX
- ▶ NetMeeting
- ▶ Network Bridge
- ▶ Network Setup Wizard
- ▶ Recovery Console (installed as a Startup option): Recovery Console is accessible from the product CD
- ▶ Remote Assistance
- ▶ Server Appliance Kit (SAK)
- ▶ Speech recognition
- ▶ Themes
- ▶ Windows Media Player
- ▶ Windows Media Services
- ▶ Windows Product Activation

2.2 Why Datacenter?

To understand why you would choose Datacenter Edition over another version of Windows Server 2003, you need to grasp the requirements and expectations of your project. Datacenter Edition could be chosen because of the level of high availability that is required, as a result of server consolidation, or maybe because of the need to provide a more scalable server for a large database.

- ▶ Does your business requires the highest possible level of server availability?

Your business relies upon the high availability of specific servers. Downtime can be expensive as transactions cease and your company's reputation suffers. The Datacenter Program with Datacenter's high availability features, including the ability to cluster and provide failover, meet this need.
- ▶ Are you migrating critical database servers from other platforms?

Many customers have key database servers that they want to migrate from other platforms in an effort to reduce the number of platforms their staff needs to maintain. Datacenter clustering and scalability meet this need.

2.3 High availability

How is IBM able to offer you an implementation of Windows Server 2003, Datacenter Edition that ensures a high level of availability?

IBM Enterprise Services for Microsoft Technologies (ESMT) has the needed skill and experience with both enterprise computing and Microsoft technologies, including Datacenter Server, to assist you from the analysis and planning phases through development and implementation of your Datacenter solution.

IBM's ESMT Windows 2003 Server, Datacenter Edition offering provides custom integration services to help you:

- ▶ Define business and technical requirements.
- ▶ Evaluate the existing environment.
- ▶ Identify and plan necessary solution components and upgrades.
- ▶ Configure, test, and tune the hardware and operating system.
- ▶ Develop, test, and document custom software stacks.
- ▶ Develop, test, and document systems management procedures.
- ▶ Implement the solution in the production environment.

Microsoft Windows Server 2003, Datacenter Edition stands out as a solution for high availability. You are purchasing more than a piece of software. Datacenter is a program. Here are some examples of what makes it more than just software, and high availability a reality.

- ▶ Hardware

IBM has chosen to provide utmost focus on the xSeries 445 Server. To be certified by Microsoft for Windows Server 2003, Datacenter Edition, the hardware must pass rigorous testing to ensure it meets Microsoft's stringent standards for compatibility and receive a specific logo from Microsoft. This is where the hardware meets the kernel drivers. On the Microsoft Windows Server 2003, Datacenter Edition HCL, specific IBM configurations are listed. Microsoft allows the download of an information file that details the hardware components and kernel drivers that were tested together.

- ▶ Software

Applications developed by software vendors are not just certified by those vendors. Microsoft has selected VeriTest to test enterprise applications and certify that they are secure, reliable, and manageable.

- ▶ Drivers

Drivers are rigorously tested, certified, and their numbers restricted.

- ▶ Operating system

Windows Server 2003, Datacenter Edition delivers new features and enhancements that mark it as the most dependable, enterprise ready server operating system Microsoft has ever created. Not only has there been a broad enhancement to many of the Windows 2000 features in Windows Server 2003, there are new technologies such as the Windows System Resource Manager (WSRM) and common language runtime (CLR), which safeguards networks from poorly designed code.

- ▶ Services

IBM can provide skilled technicians to ensure quality from requirement definition through production implementation.

IBM can roll out a production Datacenter Edition implementation that will use the certified combination of hardware, hardware flashed to the correct certified level, corresponding correct certified drivers, and software certified by VeriTest, then provide 24-hour support, maintenance, and monitoring of your site.

Change control is a key process that can be delivered by IBM as a service to ensure greater system stability and a way for customers to ensure their Datacenter Edition implementation continues to be a qualified configuration.

- ▶ Support

Technicians from IBM with high-level skills in Datacenter Edition and xSeries hardware are available to assist customers 24 hours a day, every day, through a joint IBM/Microsoft support program.

2.4 Clustering

There are many types of clusters throughout the industry with names such as availability clusters, scalability clusters, parallel computing clusters, and so forth. A Datacenter Edition cluster is an availability cluster.

Scalability cluster topologies focus on scalability by increasing overall processing power through adding nodes (servers). They often load balance and serve front-end IP applications. The loss of a single node is of minimal importance. Whatever nodes are left will handle the existing workload.

Availability clusters focus more on high availability, often with fewer (generally more powerful) nodes, and primarily serve back-end applications such as large databases, ERP applications and messaging. If a node fails, the work fails over to a surviving node which now must run the applications of the node that failed.

Our focus is on the availability or failover cluster provided by Windows Server 2003, Datacenter Edition, 32-bit version as implemented on the IBM xSeries 445 Server.

A cluster is a collection of two to eight nodes providing resources to clients. The clients do not know nor need to know which server in the cluster is providing resources to them. These clusters use a shared-nothing cluster architecture where each node in the cluster controls its own disk resources. If a node in the cluster fails, a “failover” occurs, and another node, typically the “1” in the N+1 cluster, takes over the role of the failed node. Disk resources once owned by the failed node are now owned by the new active node.

Clients are connecting to a server, but that connection is to a virtual server which can be configured on the cluster side to run on different physical servers in the cluster. During a failover, the clustering software provides a layer that shields client machines from the underlying hardware and the need to perform any configuration changes. If the application that a client is accessing can save session state data, clients may notice a momentary delay in access to their application but then can continue working. If session state data is not available, the client can reconnect to his application without knowing that there is now a different physical server in use.

There are many cluster service components in Windows Server 2003 that provide cluster functionality, keeping the cluster healthy. These include:

- ▶ Checkpoint Manager
- ▶ Database Manager
- ▶ Event Log Replication Manager
- ▶ Failover Manager
- ▶ Global Update Manager
- ▶ Log Manager

Many of these components are interrelated and manage the cluster quorum resource. The quorum is a required physical set of disks that contain a configuration database for the entire cluster. The quorum contains information on the nodes in the cluster, their configurations and state.

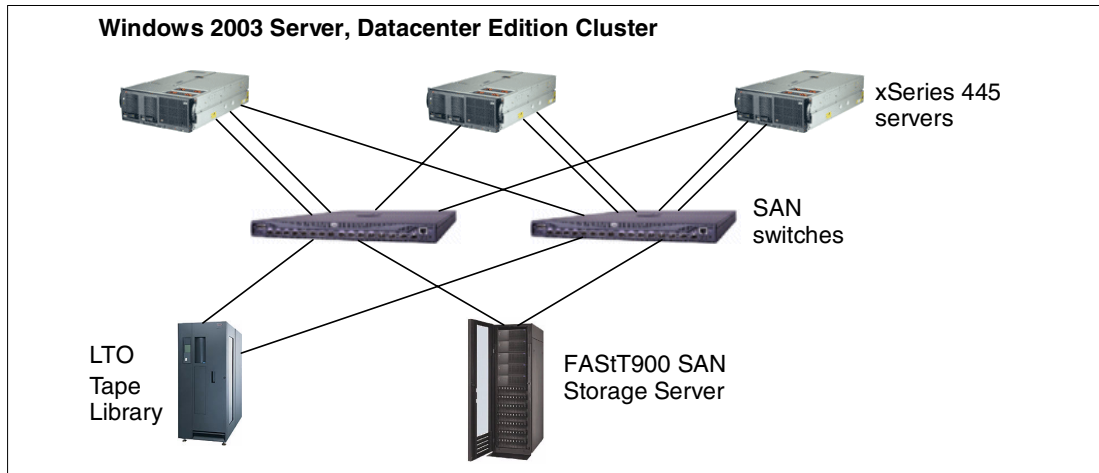


Figure 2-1 Typical Windows Server 2003, Datacenter Edition cluster

Windows Server 2003, Datacenter Edition uses Fibre Channel clustering to scale beyond the limitations of a SCSI-controlled disk. IBM implements a Datacenter Edition solution in a switched fabric environment.

A cluster is composed of two to eight IBM xSeries 445 servers, each running Windows Server 2003, Datacenter Edition.

Each IBM xSeries 445 Server contains multiple Fibre Channel PCI/PCI-X host bus adapter cards (HBAs). The first two HBAs are for SAN traffic. A third HBA is for dedicated tape backup traffic. The HBAs for SAN traffic are redundant and provide for failover.

Each server in the cluster has redundant connections to the SAN fibre switches. The two SAN switches provide redundancy. Each switch provides a connection to the IBM FAST900 SAN Storage Server (or other certified SAN storage servers) and to the LTO tape library.

The FAST900 SAN Storage Server with redundant RAID controllers, fans, and power supplies have one or more redundant paths to the EXP700 Storage Expansion Units.

Each EXP700 Storage Expansion Unit with redundant power and fans contains up to 14 drives. The size and speed of the drives should be sufficient to meet the requirements of the project. They have RAID implemented along with multiple hot-spare drives.

The LTO tape library is implemented as part of the solution. LTO, or Linear Tape Open, is a joint initiative of several vendors including IBM. LTO provides customers with investment protection since LTO has established open-format standards for high performance tape storage products for the midrange and network backup market. LTO-compliant media from any vendor can be read and written to in LTO drives from any compliant vendor.

Backups are most commonly provided by tape. In a cluster solution, a LAN-free backup is not supported. A LAN-free backup provides for backup across the SAN fabric without utilizing the Ethernet network.

There are advantages to running Datacenter in a cluster utilizing a higher number of nodes.

- ▶ An eight-node cluster offers more flexibility in application deployment since there are more possible failover scenarios and more choices as to where applications will reside within a clustered environment.
- ▶ Financially, it is a better solution because additional applications can be deployed without creating new clusters, and more clustered nodes could share a failover node(s).
- ▶ It allows for more centralized management of mission-critical servers.

2.5 Server consolidation

In today's competitive and challenging business environment, IT departments are focusing on leveraging their existing systems more effectively, making them more secure, and reducing their overhead expenses. Additionally, with the growth of the Internet, e-commerce, and automation, organizations are looking for ways to deliver new services and scale their existing ones, while realizing a higher rate of return on their investment.

To achieve these goals, IT departments are making greater use of server consolidation and are increasingly turning to Windows and Intel platforms for price and performance.

IT departments are focusing on optimizing the use of their IT infrastructure as a key way to control IT costs and enhance manageability. They are discovering that server consolidation, or right-sizing, is one of the best ways to achieve this optimization.

The driving force behind server consolidation stems from three common problems experienced by IT organizations today:

- ▶ Many organizations have up to ten times more servers than they had just a few years earlier, a situation that adds significant administrative overhead.
- ▶ Organizations often have as many as three different server platforms in their IT infrastructure, which results in a significant overhead burden in training alone.
- ▶ Many of the servers deployed today are under-utilized, meaning their capital costs represent an ongoing drain on the bottom line.

Server consolidation addresses each of these issues. To some, server consolidation may imply a return to old methods, where everything resided on a central server. But such a monolithic approach may fail to optimize the use of IT assets and could potentially lead to costly inefficiencies.

What server consolidation really means is optimizing your IT infrastructure by:

- ▶ Updating to more powerful servers, resulting in fewer servers, an approach that may reduce IT management costs, create a smaller data center footprint, and lower facility costs.
- ▶ Reducing the number of server platforms, which generally lowers management, maintenance, support, and training costs.
- ▶ Increasing interoperability between a variety of server platforms, hardware environments, and applications to help eliminate duplication and under-utilization of resources, a step which can reduce management and administration costs.

Taking these steps can help your organization achieve the degree of scalability, availability, and agility needed to meet your service delivery goals at the most affordable price.

IT departments can choose from four basic strategies for server consolidation:

- ▶ Centralization: relocating existing servers to fewer sites.
- ▶ Physical consolidation: replacing multiple smaller servers in the same architecture with fewer, more powerful servers.
- ▶ Data integration: physically combining a wide variety of data into a single repository.
- ▶ Application integration: migrating multiple applications to a new platform to take advantage of its technologies and benefits.

Whichever strategies you choose, you will find that server consolidation can help deliver:

- ▶ Extreme availability for mission-critical applications, such as ERP and CRM, through failover clustering, an approach that can provide outstanding availability.
- ▶ Higher performance, greater scalability, and more capacity.
- ▶ Faster response times by upgrading servers to higher performance systems with the latest server platforms.
- ▶ Improved access to information by consolidating data on fewer servers.
- ▶ Increased efficiency through the integration of existing data and application architectures.
- ▶ Lower maintenance and technical support costs by replacing servers that are close to the end of their life with newer systems that are under a warranty.
- ▶ A lower total cost of ownership because of centralized systems management and associated space and personnel savings.
- ▶ Improved disaster recovery capabilities, because there are fewer servers to restore.

The combination of the xSeries 445 server running Windows Server 2003, Datacenter Edition is ideal for server consolidation with solid, scalable high availability. In addition, with Windows System Resource Manager (WSRM), a part of Windows Server 2003, Datacenter Edition, enterprises can be sure critical applications can get the resources they need during peak usage periods.

2.5.1 IBM's role in server consolidation

IBM is committed to maximizing the value of server resources in the enterprise. To this end, IBM has developed expertise and tools to assist in server consolidation. Consolidating enterprise resources is a critical task for every organization, since by consolidating resources, many systems can be utilized more efficiently, maximizing resource utilization and reducing server count and administrative staff.

Not surprisingly, the process of server discovery and collecting individual server information can be very time consuming and becomes an obstacle in the process of consolidation analysis, planning, and execution. IBM is continually developing and testing an automated Consolidation Discovery and Analysis Toolset (CDAT) that is providing tremendous value in the process of server discovery and information gathering for Windows-based consolidation projects. CDAT is a stand-alone discovery toolset that automatically gathers certain server information.

CDAT is an excellent tool for generating an overview of enterprise-wide server infrastructure and discovering potential servers under various consolidation scenarios. It uses remote discovery methods to collect read-only information related to targeted systems and does not require installing new code on such systems. No write access to any of the targeted systems is required.

CDAT data is used in studies by IBM's highly skilled SCON (server consolidation) team headquartered in Beaverton, Oregon to design server consolidation solutions for customers.

Work on the CDAT toolset is performed at the IBM Center for Microsoft Technologies in Kirkland, Washington, near the Microsoft Redmond Campus. The facility is an integral part of the IBM Server Group with over 130 engineers, test and quality specialists, and enterprise solutions staff. The primary focus of the center is adapting Microsoft software to IBM products, as well as providing worldwide level-3 technical support for Microsoft products on IBM hardware. It is also home to the West Coast Executive Briefing Center and the Windows Solutions Lab where IBM hardware and Microsoft software products are demonstrated in a setting that shows current and future technologies.

2.5.2 Datacenter Edition clustering in a consolidated server environment

In an environment where the number of servers will be reduced through server consolidation, Windows Server 2003, Datacenter Edition can play a key role. Consider implementing a Datacenter cluster to provide availability in the event of failure of any of your consolidated servers, or other servers critical to your operation, such as a strategic database server that requires persistent uptime. Look at the performance options that can be leveraged using the new Windows System Resource Manager (WSRM) since applications can be allocated specific amounts of server resources at scheduled times.

Figure 2-2 on page 30 shows a current environment where less scalable servers are replaced by highly available servers in a Windows Server 2003, Datacenter Edition cluster:

- ▶ Six mail servers will be consolidated into two mail servers.
- ▶ Six file servers will be consolidated into two file servers.
- ▶ One mission critical database server with high memory and CPU requirements is replaced by one xSeries 445 server.
- ▶ Three database servers are migrated to one xSeries 445 server.




Figure 2-2 Before and after view of server consolidation and Datacenter Edition clustering implemented

The 32-bit version of Windows Server 2003, Datacenter Edition can support clusters of up to eight nodes, and individual servers with up to 32 processors and 64 GB of memory. While it is extremely scalable, another reason to select Datacenter as the right solution is business continuity.

How important is a specific server to business continuity? The IBM Datacenter Program offers services that complement a certified mix of software with the Windows Server 2003, Datacenter Edition operating system and the appropriate hardware.

2.6 The Datacenter High Availability Program

Customers who implement Datacenter are seeking the highest level of Windows scalability and availability, and therefore, serious consideration must be given to ensuring protection of their investment.

Microsoft created the Datacenter Program with its release of Windows 2000 Datacenter Server. With the release of Windows Server 2003, Datacenter Edition Microsoft has enhanced the program based on its experiences with customers and IBM.

Microsoft, qualified partners such as IBM, and customers must adhere to the standards of the Datacenter Program to ensure the highest level of support.

2.6.1 Datacenter test programs

There are several test programs that are a prerequisite for solutions to become Datacenter Edition qualified configurations. Qualified configurations have first gone through a series of

tests by Microsoft and IBM. To ensure the highest level of uptime, these tests are significantly more stringent than the standard stress tests performed on other operating systems, hardware, and drivers. Most tests include a mandatory 14-day no-fail period to stress individual applications, drivers, and hardware.

A qualified configuration must support a minimum of eight processors in any individual server. The OS must be Windows 2000 Datacenter Server or Windows Server 2003, Datacenter Edition. All components that make kernel level calls, software or hardware, are stress tested.

The required tests or programs are as follows.

- ▶ Datacenter Hardware Compatibility Test
- ▶ Enterprise Qualification Program
- ▶ Datacenter Driver Program

The sections below describe each of these.

Datacenter Hardware Compatibility Test (HCT)

The Hardware Compatibility Test (HCT) stress tests core infrastructure components on items such as:

- ▶ The network stack
- ▶ Disk subsystems
- ▶ File and print services
- ▶ Memory and processor usage

The stress test is a 14-day test of the complete configuration of proposed hardware/software/OS that simulates a similar load to what is expected in the production environment.

Retesting must occur if changes are made to hardware or system kernel components.

Table 2-2 Test comparison between Datacenter Editions of Windows 2000 and Windows Server 2003

Initial Testing	Windows 2000 Datacenter Server	Windows 2003 Server, Datacenter Edition
Initial Datacenter server configuration	14-day no-fail test	14-day no-fail test
Kernel mode hardware drivers	14-day no-fail test	<ul style="list-style-type: none"> ▶ 14 day no-fail test ▶ DC Driver Program
Kernel mode software (applications and drivers)	14-day no-fail test	<ul style="list-style-type: none"> ▶ 14 day no-fail test ▶ DC Certification
Applications	Windows compatible	Windows compatible
Storage subsystem (if kernel drivers)	14-day no-fail test	IBM or DIV Device retest and resubmission to WHQL
Changes to qualified configuration	7-day no-fail test for all changes	<ul style="list-style-type: none"> ▶ Application and driver retest, per certification or qualification specifications ▶ 14-day test for major system changes

Enterprise Qualification Program

The Enterprise Qualification Program (EQP) calls for stress testing Datacenter Edition solutions at the following levels:

- ▶ Device
- ▶ Utility
- ▶ Application
- ▶ System
- ▶ Solution

Datacenter Driver Program

The Datacenter Driver Program qualifies the hardware, software, and kernel mode drivers that reside on Datacenter Solutions through a 14-day no-fail test period. Independent Hardware Vendors (IHVs) qualify through the Windows Hardware Quality Labs (WHQL). Independent Software Vendors (ISVs) qualify through VeriTest Datacenter testing.

ISVs applications that do not touch the kernel or have kernel components are not required to be certified to be fully supported on Datacenter Edition.

2.6.2 Datacenter Certified Applications program

The Datacenter Certified Applications program is run by VeriTest (<http://veritest.com>). To certify an application, ISVs work with Microsoft to ensure the application can meet the highest standards of reliability, availability, security, and supportability. These certification standards apply to both Microsoft and third party applications.

Through IBM, a support arrangement with the ISVs creates a single point of contact for the Datacenter solution.

Not everything needs to be Datacenter certified to be used on a Datacenter system:

- ▶ Applications without kernel drivers do not require certification. Customer should check with IBM if there is a question. However, certification is always recommended.
- ▶ Inhouse customer applications do not need to be certified and may be used on a Datacenter system; such use is fully supported.

Datacenter level certification requirements are as follows:

- ▶ Applications can be installed and run on a 32 processor system.
- ▶ Applications must be stable enough to pass a rigorous stress test run on a 32-processor system.
- ▶ Applications must remain available after being subjected to crash and failover conditions occurring on an 8-way, four-node cluster configuration.
- ▶ Application vendors must be capable of providing live technical support to customers 24 hours a day, seven days a week.

Applications certified at VeriTest can be found at:

<http://cert.veritest.com/CfWreports/server/>

2.6.3 Datacenter service provider and vendor programs

Microsoft has two main programs for vendors and service providers: the Datacenter Service Provider program and the Datacenter Infrastructure Vendor program. Both programs are discussed next.

Qualified Providers and Datacenter Service Providers

IBM is a Gold Certified Microsoft Partner and Datacenter Service Provider (DSP). Qualified providers are DSPs, or Datacenter Service Providers. DSPs participate in the Datacenter High Availability Program and can be IBM, other service providers, or Microsoft.

DSPs manage partner notifications and escalation triggers for high severity issues, problem resolution, and reporting. DSPs are Gold Certified Microsoft Partners. If Microsoft is the DSP, they provide Datacenter services, resolve escalation of issues, and can escalate issues to partners.

The goal of the support structure is prompt issue resolution and continued improvement of the program. The support process involves problem resolution, notifications and triggers for support events, as well as reporting.

Datacenter Infrastructure Vendor program

The Datacenter Infrastructure Vendor (DIV) program enables storage and other peripheral vendors to participate directly in the Datacenter High Availability Program.

IBM and DIVs must establish support between each other. This allows IBM to act as a single point of contact for coordinating support delivery from the DIVs to the customer. The program allows a single coordinated support experience for the customer in which IBM, Microsoft, and DIV participate jointly.

2.6.4 Datacenter support and maintenance

Datacenter Edition support should involve use of a Datacenter Service Provider, an ongoing auditing process, change management procedures, and a maintenance program. Each of these is briefly discussed.

- ▶ High availability support

Once IBM, as the Datacenter Service Provider (DSP), has implemented your production Datacenter site, you need to protect the stability and high availability of your investment.

With Windows Server 2003, Datacenter Edition, the High Availability Resolution Queue (HARQ) replaces and enhances the Windows 2000 Datacenter Server Joint Support Queue.

- ▶ Pre-deployment audit

Before Datacenter Edition is deployed to production, the DSP will audit the configuration by running the Configuration Audit Test (CAT) tool delivered with Windows Server 2003, Datacenter Edition.

The CAT tool documents the initial installation. When run again, it verifies the system's configuration and documents any changes since the last time the CAT was run.

- ▶ Configuration audits

The Datacenter implementation should be audited regularly. The audit will check the system configuration for changes to the hardware, operating system, and kernel mode drivers. Audit tools should compare previous logs with current audit information.

- ▶ Change management

A change control service must be offered to customers. This important part of the Datacenter Program ensures a way for necessary updates to be applied and keeps the Datacenter implementation a Datacenter qualified configuration.

- ▶ Contracts and maintenance packs

IBM can provide contracts and maintenance packs from planning through implementation. Maintenance packs can go beyond production implementation to ensure proper change control and the integrity of your implementation.

For service packs, customers must pay maintenance fees (through a Software Update Subscription or Maintenance Update Subscription, as described in 3.6, “Services” on page 42) and will receive these service packs from IBM after IBM has retested the configuration.

All hotfixes (Quick Fix Engineering or QFEs) for Windows Server 2003 are available free of charge and can be applied to the Datacenter system; such a system change is fully supported by both Microsoft and IBM.

For more information on Windows Server 2003, Datacenter Edition or the Datacenter Program:

- ▶ Microsoft: <http://www.microsoft.com/windowsserver2003/>
- ▶ IBM product information: <http://www.pc.ibm.com/ww/eserver/xseries/windows/datacenter.html>



Datacenter and the x445

This chapter describes the relationship between the IBM x445 server and Microsoft Datacenter 2003. We discuss in depth why these two components, along with the service offerings, provide an industry leading Datacenter solution. This chapter is also designed to assist customers in understanding the engagement process involved when looking at the IBM Datacenter solution.

This chapter includes the following sections:

- ▶ 3.1, “Why Microsoft Datacenter from IBM?” on page 36
- ▶ 3.2, “Why the x445 and Datacenter Edition?” on page 36
- ▶ 3.3, “NUMA support in Windows Server 2003” on page 38
- ▶ 3.4, “How to engage IBM” on page 40
- ▶ 3.5, “Solution assurance” on page 41
- ▶ 3.6, “Services” on page 42
- ▶ 3.7, “HARQ and the JSQ” on page 44
- ▶ 3.8, “High Availability Guarantee Program” on page 44

3.1 Why Microsoft Datacenter from IBM?

Today, IBM is the world's largest information technology company, employing more than 310 000 employees in over 160 countries. With a heritage in enterprise computing and extensive experience in datacenter operations, IBM is able to provide customers with viable, robust and comprehensive solutions to suit all business needs.

As discussed in Chapter 2, "Features of Windows Server 2003, Datacenter Edition" on page 13, there are four editions of the Windows 2003 operating system.

- ▶ Windows Server 2003, Web Edition
- ▶ Windows Server 2003, Standard Edition
- ▶ Windows Server 2003, Enterprise Edition
- ▶ Windows Server 2003, Datacenter Edition

These provide customers with a variety of choices, depending on individual business needs. The Datacenter Edition has been designed to offer the highest level of reliability, scalability and manageability of all Microsoft Windows 2003 server offerings.

Datacenter Edition is only available through an Original Equipment Manufacturer (OEM) partner, such as IBM. When choosing Datacenter Edition, a customer has the ability to choose the best and most appropriate qualified OEM to provide all hardware, services and support. These components are essential to building and sustaining any high availability environment. OEMs such as IBM are part of the Windows Datacenter High Availability Program. This program is based on best practices gathered since the inception of Microsoft Datacenter Server in September 2000. For OEMs and service providers to be a member of this program, they must follow a stringent set of reporting and audit features which have been defined by Microsoft.

The IBM Datacenter Solution Program provides a comprehensive set of hardware, service, and support offerings for Microsoft Datacenter 2003 Server. IBM has the ability and expertise to provide our customers with all the components required when implementing this high availability, highly scalable solution.

Information about the IBM Datacenter Solution Program is available from:

<http://www.pc.ibm.com/ww/eserver/xseries/windows/datacenter.html>

3.2 Why the x445 and Datacenter Edition?

The x445 server is the ideal platform for a Datacenter Edition solution. The x445 server is designed using IBM Enterprise X-Architecture (EXA) which brings mainframe features and capabilities to the IBM Intel-based server range. A core feature of Enterprise X-Architecture is the chipset. The x445 contains the second generation of the XA-32 chipset. With an increase in the Xcel4 Server Accelerator cache to 64 MB per 4-way SMP Expansion Module, the x445 is able to increase performance by reducing the need to access main memory.

The x445 is supported for use with Datacenter Edition. It has passed all Microsoft requirements to become a fully certified Windows Server 2003, Datacenter Edition solution. This 4U server is the most rack-optimized, scalable server in the marketplace. The x445 server enables customers to become more responsive and resilient to the technical problems they face. The x445 server offers several levels of hardware protection for maintained high availability, and offers unparalleled scalability and ease of use.

The x445 uses predictive and proactive technology to deliver high availability solutions. Coupled with IBM's award winning IBM Director and other useful tools, the x445 is the most robust and autonomic Intel-based server in the marketplace.

IBM has worked with Microsoft to deliver the industry leading Datacenter solution. Below are examples of how the x445 and Datacenter Edition integrate together.

► **Datacenter Program**

Q: Can IBM provide an approved OEM Datacenter Solution Program?

How IBM delivers: The IBM Datacenter Solution Program is available through your local IBM sales representative. IBM offers a wide range of hardware, services and support.

IBM is an approved Microsoft Datacenter OEM partner. IBM is a member of the Microsoft Windows Datacenter High Availability Program.

► **NUMA support**

Q: I'm looking for a server today to meet my current requirements. The server also needs to offer me add memory, processors and more, while maintaining my initial investment and still providing high levels of performance.

How the x445 delivers: The IBM x445 server is built on NUMA-based architecture. This provides customers with a long term growth plan while maintaining excellent system performance. NUMA eliminates typical bottlenecks associated with SMP scalability.

How Datacenter Edition delivers: NUMA-based hardware is supported using Windows Datacenter 2003 Edition. NUMA support enables the operating system to run fully optimized on its underlying NUMA-based hardware. This support offers additional performance on NUMA-based hardware.

► **SMP scaling**

Q: I need to ensure that I buy enough headroom to meet my business needs for today and the future; I don't want to buy a large expensive SMP server today if I'm not going to use all of its capacity for another 12 months.

How the x445 delivers: The x445 server can scale from a 2-way to a 16-way SMP configuration, with 32-way SMP in the future. You can buy the capacity today, and add capacity when you need it. This can be done while maintaining your initial investment.

How Datacenter Edition delivers: Microsoft Datacenter 2003 supports large SMP environments. Microsoft Datacenter 2003 supports from 2 to 32 processors. Combine this SMP growth capacity with the additional NUMA support, and optimum system performance is assured.

► **Memory capacity**

Q: The x445 server needs to have the ability to scale its memory. If I decide to purchase Microsoft Datacenter, will I be able to accommodate an increase in memory demands?

How the x445 delivers: The x445 server is an extremely scalable server. Its memory capacity can scale up to 64 GB in an 8-way configuration, and even higher with future 32-way SMP support.

How Datacenter Edition delivers: Microsoft Datacenter 2003 supports up to 64 GB RAM on a 32-bit system.

► **Hot-add memory support**

Q: I'd like the ability to add more system memory when I need it, and I don't want to schedule system downtime.

How the x445 delivers: The x445 memory subsystem supports adding memory to the server while the system is still running. Certain restrictions apply, but a reboot is not required to take advantage of the newly installed memory.

How Datacenter Edition delivers: Windows 2003 Datacenter Edition supports the x445's ability to hot-add memory. This support is provided by loading the Static Resource Affinity Table (SRAT) during system boot. These changes have been made by Microsoft to allow the memory manager to support the dynamic addition of memory to a running system.

► **Hot-swap memory support**

Q: I want to reduce system downtime due to memory failures. If a memory module fails, I should be able to hot-replace it without shutting down my server.

How the x445 delivers: With the use of memory mirroring, if a faulty DIMM is present in the x445, the module can be replaced while the server is in operation. This feature is independent of the operating system.

► **Extra PCI-X slots**

Q: Each 8-way x445 chassis has six PCI-X slots. Can I add additional slots?

How the x445 delivers: The x445 supports the addition of an RXE-100 remote I/O enclosure which provides six or twelve additional PCI-X slots. One RXE-100 can be connected to a 8-way or 16-way configuration.

How Datacenter Edition delivers: Windows Server 2003 fully supports the use of the RXE-100.

3.3 NUMA support in Windows Server 2003

Figure 3-1 is a block diagram of the components of an 8-way x445 and shows the two SMP expansion modules connected through the 3.2 Gbps SMP Expansion Ports (using two SMP scalability cables). The block diagram shows the NUMA architecture where the system memory is divided up and shared between the processors.

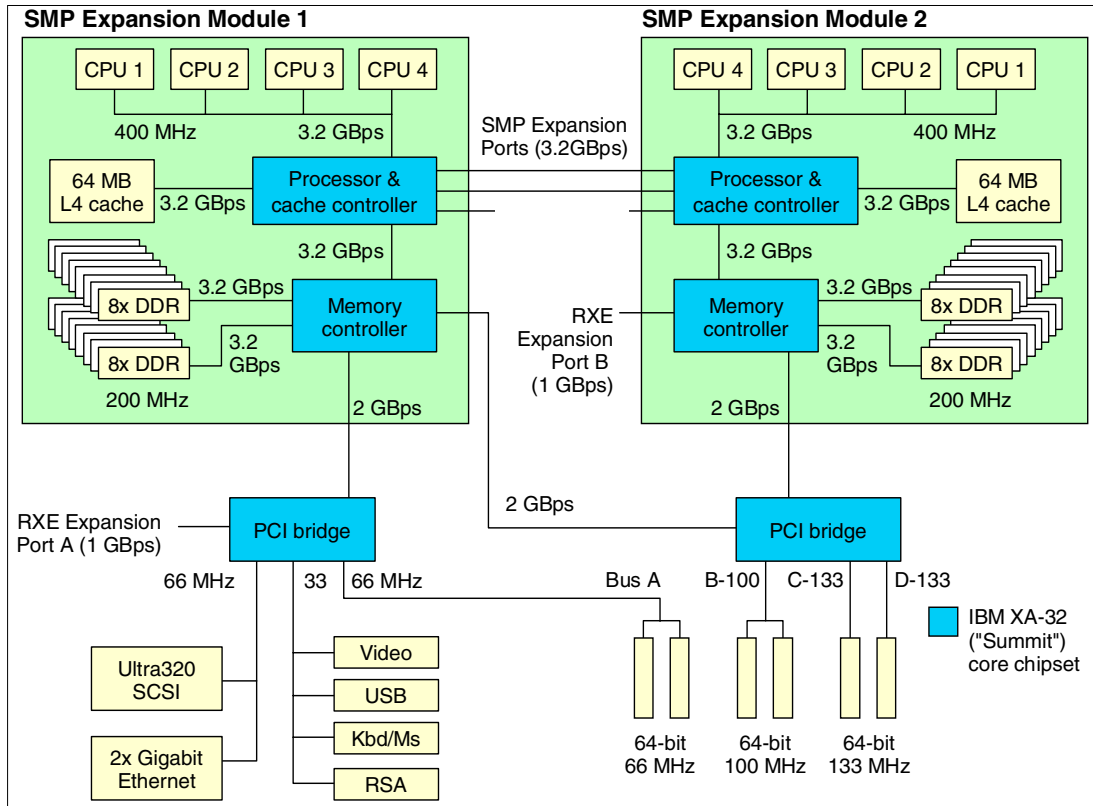


Figure 3-1 xSeries 445 system block diagram: two SMP Expansion Modules

To support the advanced multiple memory controller technology used in the x445, Microsoft has optimized Windows 2003 to take maximum advantage of its NUMA-based architecture. Windows 2003 can identify each SMP Expansion Module in the system, and without any change to the application, can automatically balance the workload between the multiple SMP expansion modules for optimal performance.

To identify each SMP Expansion Module, Windows 2003 uses a *Static Resource Affinity Table (SRAT)*. The SRAT defines the specific topology of the processors, installed memory, interrupts, and adapters local to each SMP Expansion Module. The operating system uses this information to affinitize each process and all of its associated components (threads, memory, I/O buffers, and interrupts) to a specific SMP Expansion Module.

Optimal performance and scaling is obtained by keeping all threads, memory, buffers, and interrupts for each process local to a specific SMP Expansion Module. When combined with the XceL4 cache, remote processor-to-processor communication is greatly reduced. Furthermore, all memory accesses by processors are issued to a local memory controller, all threads execute on local processors, and all interrupts are serviced by processors local to the specific interrupt request.

This optimization is done automatically by the operating system without any change to the application. Since each SMP Expansion Module has processors, a local memory controller, memory, and I/O, the net effect is to produce scale-out-like parallel processing inside a single x445 system.

This results in increased performance when compared to standard SMP designs.

3.3.1 Using SRAT to support hot-add memory

The SRAT is referenced and its information loaded into BIOS when the system is booted. The information used in the SRAT is then used to affinitize each process and all of its associated components to a specific SMP Expansion Module. When each process and associated components are affinitized, they are placed into a *proximity domain* which is then allocated a proximity domain ID. The SRAT uses proximity domains to store and identify the affinitized information. For more information on the SRAT, refer to:

<http://www.microsoft.com/whdc/hwdev/platform/proc/sratdwn.msp>

The ability to affinitize newly added memory during system operation (that is, memory which has been hot-added) is also required. This is needed to ensure optimum system performance, without the requirement to reboot the system after the successful installation of additional memory.

SRAT provides an association between a region of memory and an appropriate proximity domain ID (in which this region, or section of memory, should belong). By doing this, proper system performance can be assured without the need to reboot. Without this SRAT feature, hot-adding memory to a system would require a system reboot to properly optimize the NUMA architecture, thus defeating the purpose of the hot-adding capability.

Note: The memory slots which support hot-add have been associated with the appropriate proximity domain before physical installation takes place. This is built into the SRAT by IBM and loaded into BIOS during the system boot phase.

3.4 How to engage IBM

There are two ways for a customer to engage in the IBM Datacenter Solution Program.

- ▶ Large and enterprise customers who work with IBM directly should engage their IBM sales representative.
- ▶ For customers with relationships with IBM Business Partners, the business partner will engage IBM after qualifying the customer's requirement.

The IBM sales representative will engage IBM Global Services to develop a comprehensive services package to suit the customer's requirements.

The process used to successfully configure and implement a Datacenter Program solution is highly customized to suit the customer, but the basic process is as follows. The entire process is managed by the IBM sales representative and the IBM Datacenter project office.

Tip: The IBM Datacenter project office e-mail address is xhe@us.ibm.com, and is currently managed by Wayne Phipps.

1. Proposal phase
 - a. The customer engages the IBM sales representative or IBM EXAct Business Partner.
 - b. The sales representative (or Business Partner) uses Configurator to produce a Datacenter configuration, available from:
<http://www.pc.ibm.com/us/eserver/xseries/library/configtools.html>
 - c. The sales representative produces a rack configuration.

- d. The sales representative contacts a IBM Global Services contract specialist to produce a services plan.
 - e. The sales representative e-mails the entire solution package to the IBM Datacenter project office.
2. Quote phase
 - a. The sales representative obtains pricing.
 - b. The sales representative coordinates a customer Solution Assurance Review.
 - c. The sales representative e-mails Solution Assurance Review results to the project office.
 - d. The customer accepts the IBM offer and places the order.
 - e. The sales representative e-mails the purchase order and parts list to the project office.
 3. Order processing
 - a. The order is entered into the IBM ordering system.
 - b. The project office monitors progress.
 - c. The project office conducts weekly conference calls to coordinate progress.
 - d. The project office coordinates delivery of the solution components.
 4. Installation and on-site services
 - a. IBM Global Services integrate the solution into the customer's environment.
 - b. The skills transfer take place.
 - c. The solution is then brought online.
 - d. The IBM team hands over the solution to the customer or business partner.

Tip: For further information on the IBM Datacenter Solution sales process, IBM employees can review the Sales Guide, available via Notes®, on server D01DBM20, database x_dir\u20a1xserie1.nsf.

3.5 Solution assurance

Some level of Solution Assurance Review (SAR) should be performed on all IBM solutions. The level of SAR (Self, Peer, or Expert) should match the complexity of the solution. However, due to the complexity of a Datacenter solution, an expert SAR is always performed.

The IBM Datacenter Solution Program assures that all components, including hardware, software and services have been through the appropriate solution assurance process. IBM will perform a Solution Assurance Review as a standard part of a Datacenter Solution sale.

For further information on Solution Assurance in general, refer to the following Web sites:

- ▶ For IBM employees:

<http://w3.ibm.com/support/assure>

- ▶ For Business Partners:

<http://www.ibm.com/partnerworld/techsupport>

Select **Solution Assurance** from the links provided.

For geography-specific documents, refer to the following:

- ▶ For the Americas, reference document SA447.
- ▶ For EMEA, reference document SA424.
- ▶ For Asia Pacific, reference document SA441.

3.6 Services

When deploying Datacenter Edition, it is imperative that the solution be robust and reliable. With the initial installation of the hardware, strict guidelines must be put in place to maintain the integrity of the solution. This includes change control processes and the correct level of support and on-site service. As part of the IBM Datacenter Solution Program, IBM has the unique ability to provide all facets of the installation. This includes solution customization, planning, advanced support and maintenance.

IBM Global Services can deliver a wide variety of Datacenter service offerings. Their staff has diverse knowledge of all facets of Datacenter operations. The following is a list of all services available with the IBM Datacenter Solutions Program.

Note: These services are defined as either mandatory, strongly recommended or recommended, which specifies whether they should be included in any proposal to a customer.

- ▶ Product customization services

IBM helps reduce the installation time of the IBM xSeries servers involved in the Datacenter solution because these arrive pre-configured. Customers can be assured that the IBM servers have been set up using the highest standards in practice at our ISO-9001 certified facilities. Product customization is pre-delivery work performed on IBM hardware and includes the following subcomponents.

- IBM hardware setup and integration- mandatory

Using state-of-the-art manufacturing facilities, IBM will set up and install the servers with all requested hardware options such as processors, memory and hard disks. All options are tested to ensure the correct operability.

- IBM software imaging services- mandatory

These services include the initial creation, modification, verification, final testing and loading of the Microsoft Datacenter server operating system onto the x445 servers.

Imaging services can also optionally include:

- Asset tag services

IBM can install tamper-evident asset tags on xSeries servers before they are delivered. These tags assist customers in asset tracking and identification. The information on these tags can include the company name, machine type and model, serial and asset number.

- Operating systems and firmware personalization services

IBM can change the software and hardware settings to IT specifications before the equipment is delivered. These can include custom disk partitioning schemes, personalization of the operating system (computer name, workgroup, etc.) and custom BIOS settings.

More information on product customization services can be found at:

http://www.pc.ibm.com/us/eserver/xseries/systems_management/pcs.html

- ▶ System warranty upgrades - strongly recommended (geography-specific)

The x445 ships standard with a three year, 9x5 warranty (9 hours a day, 5 days a week, next business day response). Upgrades to this base warranty are available and are strongly recommended. This is to ensure that a customer will have the appropriate support in the event of a problem. Warranty upgrades are available. These include 24x7x4 (24 hours, 7 days, 4 hour response) and 24x7x2 (24 hours, 7 days, 2 hour response).

More information on IBM warranty upgrades can be found at:

<http://www.ibm.com/pc/support/site.wss/document.do?lnDocid=MIGR-4CMSC7>

- ▶ Microsoft Authorized Premier Support (MAPS) - strongly recommended

IBM provides cost-effective, quality software support for Microsoft Datacenter with MAPS. As a base offering for Datacenter, IBM teams up with Microsoft to provide world-class expertise which delivers the highest level of product support. Support is provided with a base of ten incidents, while including access to the Joint Support Queue.

The MAPS service consists of the following:

- Prompt problem resolution is offered for Microsoft technologies with 24x7 support.
- The customer receives an IBM technical account manager. This provides the customer with a single point of contact for Microsoft problem resolution.
- The customer receives a Microsoft designated support account manager. This provides the customer a single point of contact for Microsoft problem resolution.
- The support coverage for MAPS is allocated in blocks of ten incidents.
- Monthly status calls and reports take place.
- Escalation management is provided.
- There is unlimited access to Microsoft's premier online Web site.
- A one year subscription to Microsoft TechNet Plus is offered, providing monthly CD kits with software fixes, drivers and technical documentation.

- ▶ Advanced Support (strongly recommended)

This is the highest level of support available through IBM. The coverage includes 24x7, around the clock support. This support is designed for continuous, business-critical system operation. The HARQ, as discussed in 3.7, "HARQ and the JSQ" on page 44, is included in the advanced support offering.

Based on your requirements, IBM can provide many options to further enhance this unique offering. Some examples include:

- A dedicated on-site services specialist
- Remote system(s) monitoring your systems and networking environment
- A single point of entry for OEM products
- Customized escalation procedures

Advanced Support offers customers customized support to help them meet their company's business-critical requirements. It provides the proactive support needed to help achieve the highest possible availability.

More information can be found at:

<http://www.ibm.com/services/its/us/mus89d1.html>

- ▶ Software Update Subscription (strongly recommended)

The Software Update Subscription provides periodic updates to the Microsoft Windows Datacenter operating system, which you license for a period of one year. This subscription also includes IBM updates to firmware and device drivers certified by Microsoft for use

with the Datacenter Solution. IBM builds, tests, and provides the complete certified package of these components.

- ▶ Maintenance Update Subscription (strongly recommended)

The Maintenance Update Subscription is similar to the Software Update Subscription except that it does not include new versions of the operating system (just service packs and fixes).

- ▶ Planning and Implementation Services (strongly recommended)

This service provides planning and consultation services. This includes IBM Datacenter solution sizing and implementation planning and the services required, which are customized on an individual solution basis.

- ▶ Business Continuity (recommended)

These services include high availability (clustering) and disaster recovery services and planning.

- ▶ Systems Management (recommended)

These are consulting and implementation services for the optimal use of IBM Director and the xSeries Remote Supervisor Adapters. This also includes the overall systems management design.

- ▶ Education (recommended)

This service includes skills transfer for the xSeries and Microsoft offerings. This teaches the customers to better use the technology.

3.7 HARQ and the JSQ

As included in the previous Datacenter Program for Windows 2000, IBM and Microsoft employees managed the Joint Support Queue (JSQ). This queue was generated to provide a single point of contact for any problems occurring in the Datacenter environment for the customer.

The High Availability Resolution Queue (HARQ) has been created to replace the JSQ, and represents an expanded version of it. Current customers involved with the JSQ will be transitioned to the HARQ.

The HARQ is a Microsoft support program that provides 24x7x365 access to Microsoft support technicians. When a customer implements the IBM Datacenter Program, should a software-related problem occur, IBM will engage Microsoft using the HARQ for problem resolution assistance. IBM will act as an intermediary between the customer and Microsoft.

For more information about the HARQ, see:

<http://www.microsoft.com/windowsserver2003/datacenter/dcprogramfaq.msp>

3.8 High Availability Guarantee Program

IBM xSeries continues to improve hardware availability through the 99.99% Availability Guarantee Program. The program has recently been enhanced for more flexible server hardware configurations and larger clusters. IBM has expanded the bounded server configurations to include xSeries servers that have completed Microsoft Cluster Service certification as well as Options by IBM, as listed on the ServerProven Web site. This includes

the x445 server. All servers in the cluster must be running Microsoft Windows Datacenter 2003 edition.

This Datacenter solution wraps a robust suite of required services around clustered systems to help avoid unscheduled downtime. The services that provide the 99.99% availability include:

- ▶ Pre-sales solution assurance
- ▶ Installation services
- ▶ Setup for remote monitoring
- ▶ Warranty upgrades and maintenance options
- ▶ Project manager
- ▶ Weekly review of system logs

The 99.99% availability in the IBM Datacenter solution is measured from the operational availability of the system over a 12-month period of time. Exclusions from the high availability guarantee calculation are all system outages not directly caused by the specific system. Planned outages for preventative maintenance, upgrades, and changes are not calculated in the measurement.

Unplanned outages are defined as a component failure within the specified system which causes the system to not be accessed and/or to not process data. IBM will conduct a root-cause analysis and determine whether any unplanned outages were caused by the specific system failure, hence missing the high availability requirement.

x445 servers set up in the IBM Datacenter Solution Program are configured in a cluster with external storage enclosures, UPS power protection, monitors and keyboards.

Abbreviations and acronyms

ACPI	Advanced Configuration and Power Interface	NOS	network operating system
ADMT	Active Directory Migration Tool	NUMA	Non-Uniform Memory Architecture
ASP	Application Service Provider	OEM	original equipment manufacturer
CAT	Configuration Audit Test	OS	operating system
CDAT	Consolidation Discovery and Analysis Toolset	OSPF	Open Shortest Path First
CEC	Central Electronics Complex	PFA	Predictive Failure Analysis
CLR	common language runtime	PPTP	Point-to-Point Tunneling Protocol.
CMAK	Connection Manager Administration Kit	RADIUS	Remote Authentication Dial-in User Service
CRM	Customer Relationship Management	RAID	redundant array of independent disks
DC	Domain Controller	SAK	Server Appliance Kit
DDR	Double Data Rate	SAN	Storage Area Network
DIV	Datacenter Infrastructure Vendor	SAR	Solution Assurance Review
DSP	Datacenter Service Provider	SCON	server consolidation
ECC	error checking and correcting	SRAT	Static Resource Affinity Table
EFS	Encrypting File System	VDS	Virtual Disk Service
EQP	Enterprise Qualification Program	VPN	Virtual Private Network
ERP	Enterprise Resource Planning	VSCS	Volume Shadow Copy Service
ESMT	Enterprise Services for Microsoft Technologies	WHQL	Windows Hardware Quality Labs
EXA	Enterprise X-Architecture	WLBS	Windows Load Balancing Service
FRS	File Replication Service	WSRM	Windows System Resource Manager
HARQ	High Availability Resolution Queue		
HBA	host bus adapter		
HCL	Hardware Compatibility List		
HCT	Hardware Compatibility Test		
IAS	Internet Authentication Service		
ICF	Internet Connection Firewall		
ICS	Internet Connection Sharing		
ILS	Internet Locator Service		
ISP	Internet Service Provider		
ISV	Independent Software Vendor		
JSQ	Joint Support Queue		
LDAP	Lightweight Directory Access Protocol		
LED	light emitting diode		
LTO	Linear Tape Open		
MAPS	Microsoft Authorized Premier Support		
MMC	Microsoft Management Console		
MMS	Metadirectory Services Support		
MSCS	Microsoft Cluster Services		
NDIS	network driver interface specification		
NLB	network load balancing		

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this Redpaper.

IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Relevant documents include:

- ▶ *IBM @server xSeries 445 Solution Assurance Product Review Guide*, REDP3691 (IBM intranet access required; download from <http://w3.itso.ibm.com>)
- ▶ *IBM @server xSeries 440 Planning and Installation Guide*, SG24-6196
- ▶ *Implementing Systems Management Solutions using IBM Director*, SG24-6188

Online resources

These Web sites and URLs are also relevant as further information sources:

IBM Web pages

- ▶ IBM Datacenter Program information
<http://www.pc.ibm.com/ww/eserver/xseries/windows/datacenter.html>
- ▶ IBM infotips site (IBM intranet access required)
<http://w3.pc.ibm.com/helpcenter/infotips/>
- ▶ ServerProven and hardware compatibility
<http://www.pc.ibm.com/compat>
<http://www.pc.ibm.com/us/compat/serverproven>
<http://www.pc.ibm.com/ww/eserver/xseries/serverproven>
- ▶ Enterprise X-Architecture
<http://www.pc.ibm.com/us/eserver/xseries/xarchitecture/principals.html>
<http://www.pc.ibm.com/us/eserver/xseries/xarchitecture/enterprise>
- ▶ IBM Product Customization Services
http://www.pc.ibm.com/us/eserver/xseries/systems_management/pcs.html
- ▶ IBM warranty upgrades
<http://www.ibm.com/pc/support/site.wss/document.do?ln docid=MIGR-4CMSC7>

- ▶ Solution Assurance (IBM intranet or Business Partner access required)
<http://w3.ibm.com/support/assure>
<http://www.ibm.com/partnerworld/techsupport>
- ▶ Operational Support Services — Advanced Support
<http://www.ibm.com/services/its/us/mus89d1.html>

Microsoft Web pages

- ▶ Windows Server 2003 product information
<http://www.microsoft.com/windowsserver2003>
- ▶ Windows Datacenter High Availability FAQ
<http://www.microsoft.com/windowsserver2003/datacenter/dcprogramfaq.mspx>
- ▶ HCL Criteria for Datacenter Servers
<http://www.microsoft.com/whdc/hwdev/platform/server/datacenter/DCHCT.mspx>
- ▶ Static Resource Affinity Table
<http://www.microsoft.com/whdc/hwdev/platform/proc/sratdwn.mspx>
- ▶ Microsoft Hardware Compatibility List (HCL)
<http://www.microsoft.com/windows/catalog/server>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



Windows Server 2003, Datacenter Edition on the IBM *e*server xSeries 445



Describes the features of the xSeries 445

Introduces the capabilities of Datacenter Edition

Explains why Datacenter Edition and the x445 are an ideal match

This IBM Redpaper describes two flagship products from IBM and Microsoft and explains why these products are a perfect fit for one another. The IBM *e*server xSeries 445 is the next-generation scalable enterprise server, with up to sixteen processors and up to 128 GB of memory. Microsoft Windows Server 2003, Datacenter Edition is the high-end operating system for Intel-processor servers. These two products are ideal for enterprise customers who require an extremely large single-image business computing environment.

This Redpaper is divided into three chapters:

Chapter 1, "Features of the x445" on page 1 covers Enterprise X-Architecture and NUMA, the architectures upon which the x445 is based, and describes the hardware technology that makes up the various subsystems in the server.

Chapter 2, "Features of Windows Server 2003, Datacenter Edition" on page 13 covers the features of the operating system, why customers should use Datacenter Edition over other Windows operating systems, and where clustering, high-availability, and server consolidation are appropriate.

Chapter 3, "Datacenter and the x445" on page 35 examines the relationship between IBM, Microsoft and the customer when a Datacenter Edition-based solution is being implemented.

This Redpaper is aimed at customers, Business Partners and IBM employees who need to understand how Datacenter Edition and the xSeries 445 are an ideal combination for high capacity business computing environments.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:
ibm.com/redbooks**