IBM

IBM @server xSeries 445 Planning and Installation Guide

- Learn the technical details of the x445 models
- Discover key IBM Director management tools
- Prepare for and perform an installation

David Watts
Jure Arzensek
Antony Collins
Jian Yong Wang
Russ Weiner

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM @server xSeries 445 Planning and Installation Guide

October 2003

SG24-8870-00

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

First Edition (October 2003)

This edition applies to the IBM @server xSeries 445, machine type 8870.

Minor updates were made to this redbook. The last update was October 21, 2003

Contents

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:

@server™	Enterprise Storage Server®	PS/2®
@server™	FlashCopy®	Redbooks™
Eserver™	Informix®	Redbooks (logo) ™
eServer™	IBM®	RETAIN®
ibm.com®	Lotus Notes®	ServerProven®
iSeries™	Lotus®	ServeRAID™
pSeries®	Netfinity®	Summit®
xSeries®	NetBAY3	ThinkPad®
zSeries®	Notes®	Tivoli®
Chipkill™	OnForever™	TotalStorage®
Domino®	PowerPC 750™	WebSphere®
DB2®	PowerPC®	X-Architecture™
DFS™	Predictive Failure Analysis®	

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

The IBM® @server™ xSeries® 445 is the IBM flagship industry standard server and is the second generation implementation of the 32-bit IBM XA-32 chipset as part of the Enterprise X-Architecture™ strategy. The x445 provides new levels of high availability and price performance, and offers scalability from two-way to 32-way SMP, from 2 GB to 256 GB of memory, and up to 36 PCI slots, all in one single system image.

This IBM Redbook is a comprehensive resource on the technical aspects of the server, and is divided into six key subject areas:

- Chapter 1, “Technical description” introduces the server and its subsystems and describes the key features and how they work.
- Chapter 2, “Positioning” examines the types of applications that are used on the x445 server, including server consolidation, line-of-business applications, and infrastructure applications. It reviews the features that make the x445 such a powerful system.
- Chapter 3, “Planning” describes the aspects of planning to purchase and install the x445. It covers such topics as configuration, operating system specifics, scalability, and physical site planning.
- Chapter 4, “Windows Server 2003, Datacenter Edition” describes the IBM Datacenter program, and describes how Datacenter Edition takes advantage of the capabilities of the x445.
- Chapter 5, “Installation” goes through the process of installing Windows® Server 2003, Windows 2000, Red Hat and SuSE Linux, NetWare, and VMware ESX Server. It describes what BIOS and drivers updates are appropriate and when to install them.
- Chapter 6, “Management” describes how to use the key IBM Director extensions designed for the x445: Scalable Systems Manager, Active PCI Manager, and Application Workload Manager.

The team that wrote this redbook

This redbook 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 writer residencies and produces IBM Redbooks™ on hardware and software topics related to IBM xSeries systems and associated client platforms. He has authored more than 30 Redbooks and Redpapers; his most recent books include *Implementing Systems Management Solutions Using IBM Director*. He has a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM for more than 14 years. He is an IBM @server Certified Specialist for xSeries and an IBM Certified IT Specialist.

Jure Arzensek is an Advisory IT Specialist at IBM Slovenia, and responsible for the post-sales technical support for the xSeries, PCD and FAStT products in the Central and Eastern Europe (CEE) region. He has 11 years of experience with Intel®-based servers and operating systems. He holds a bachelor's degree in Computer Sciences from the University of Ljubljana. Besides the xSeries servers and operating systems, his areas of expertise include Microsoft® clustering and FAStT Storage Servers. He has co-authored two other Redbooks, *Fibre Array Storage Technology A FAStT Introduction*, SG24-6246-00, and *Netfinity Server Disk Subsystems*, SG24-2098-03.

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). He co-authored the Redpaper *Windows Server 2003, Datacenter Edition on the IBM @server xSeries 445*, REDP3700.

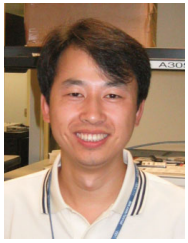

Jian Yong Wang is an Advisory IT Specialist working at IBM China, and he is a Senior xSeries Products Specialist on the Advanced Technical Support Team. He is responsible for country-wide technical support for xSeries servers and Level 2 technical support for sales teams. Jian Yong has six years of professional experience with xSeries, Netfinity® and PC Server systems including post-sales and pre-sales technical support at IBM. His areas of expertise includes high-end solutions for xSeries Server, Windows 2000 Datacenter Server, VMware, and clustering. He is an IBM @server Certified Systems Expert for xSeries, Microsoft Certified Professional Systems Engineer, MCP +Internet, and SCO Certified System Engineer.

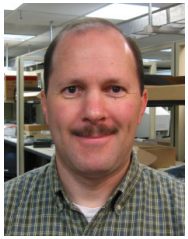

Russ Weiner is an xSeries Technical Support Specialist in Seattle, Washington. He has twenty years of experience in Information Technology as a programmer, network administrator, and in mainframe operations, the help desk environment, and data security. He has worked at IBM in the Advanced Technical Support organization for five years. His current area of focus is field technical support for xSeries servers. He co-authored the Redpaper *Windows Server 2003, Datacenter Edition on the IBM @server xSeries 445*, REDP3700.

Thanks to the following people for their contributions to this project:

Marc Baker, xSeries Universal 4-way Development, Raleigh
Ralph Begun, xSeries System Development, Raleigh
Donn Bullock, Global Brand Manager, xSeries 445, Raleigh
Mark Chapman, xSeries Marketing, Raleigh
Henry Chung, Technical Project Manager, Datacenter Development, Seattle
Art Fisher, Operating System Technology Support, Raleigh
Alan Fontaine, Integral Functions Team Lead, Raleigh
Conor Flynn, Scalable Systems Manager Development, Kirkland
Dave Hardman, VMware Global Marketing Manager, Raleigh
Ramona Hendren, xSeries Information Development, Raleigh
Ray Isler, WW Service & Support Planning, Raleigh
Bruce Jones, Manager Linux Technology Center xSeries, Beaverton
Dave Jensen, xSeries Mechanical Development, Raleigh
John Kranz, IBM Center for Microsoft Technologies, Seattle
Michael S. Lee, IBM Center for Microsoft Technologies, Seattle
Grace Lennil, IBM Center for Microsoft Technologies, Seattle
John McAbel, WW Database/Cluster Product Manager, Beaverton
Deborah McDonald, Scalable Systems Manager Development, Kirkland
Paul McGrath, xSeries Market Management Services & Software, Raleigh
Gregg McKnight, Distinguished Engineer, xSeries Performance, Raleigh
Franklin Nanney, WW Service & Support Planning, Raleigh
Michael L. Nelson, eServer™ Solutions Engineering, Raleigh
Kenny Padgett, WW Service & Support Planning, Raleigh
Michael Parris, WW Service & Support Strategy, , Raleigh
Kiron Rakkar, Manager of WebSphere® Beta Programs, Raleigh
Bob Stephens, xSeries Software Architecture, Kirkland
Gary Turner, Scalable Systems Manager Project Manager, Kirkland
Damon West, IBM eServer Technology Enablement Center, Raleigh
Bob Zuber, WW xSeries Product Marketing, Raleigh

Thanks to the following Novell employees:

Jean-Jacques Clar
Jack Hodge

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

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 Redbooks to be as helpful as possible. Send us your comments about this 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

1

Technical description

The next-generation scalable enterprise server, the x445, powered by Enterprise X-Architecture, sets the pace in the market for eight-way and 16-way servers, blending unprecedented XpandOnDemand processor and I/O scalability, #1 benchmark performance, and OnForever™ availability for mission-critical database, ERP, and CRM solutions.

The following are the key features of the x445:

- XA-32 second generation chipset.
- Models with the Intel Xeon MP processors, up to 2.8 GHz and 2 MB L3 cache, upgradable to four-way, eight-way and 16-way. Future plans also have the x445 expandable to 32-way.
- Entry-level models with 3.0 GHz Intel Xeon DP processors, upgradable to four-way.
- Active Memory with Memory ProteXion, memory mirroring, hot-swap and hot-add memory.
- XceL4 Server Accelerator Cache, now 64 MB per SMP Expansion Module (128 MB when two SMP Expansion Modules are installed in one x445).
- 2 GB DDR memory standard, expandable to 64 GB per eight-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 that can be shared between two x445s.
- Remote Supervisor Adapter standard. Plans are to make future models standard with the Remote Supervisor Adapter II-EXA adapter and to provide an option to upgrade existing x445s to the Remote Supervisor Adapter II-EXA.
- Serial port available via a cable and PCI-slot bracket assembly.

1.1 x445 models

The x445 models, shown in Figure 1-1, are available from July 2003. Multi-node configurations are described in 1.1.1, “16-way configurations” on page 3.

Table 1-1 Models available from July 2003

Model	Standard processors	Max SMP	Max SMP multi-node	L2 cache	L3 cache	Standard memory (DDR DIMMs)
8870-1RX	2x 2.0 GHz Xeon MP, 1 MB L3	8-way	16-way	512 KB	1 MB	2 GB (4x 512 MB)
8870-2RX	2x 2.5 GHz Xeon MP, 1 MB L3	8-way	16-way	512 KB	1 MB	2 GB (4x 512 MB)
8870-4RX	4x 2.8 GHz Xeon MP, 2 MB L3	8-way	16-way	512 KB	2 MB	2 GB (4x 512 MB)
8870-3RY	2x 3.0 GHz Xeon DP, 512 K L2	4-way	See note	512 KB	0 MB	2 GB (4x 512 MB)
8870-4RY	4x 3.0 GHz Xeon DP, 512 K L2	4-way	See note	512 KB	0 MB	2 GB (4x 512 MB)
Note: The Xeon DP models can be upgraded to use Xeon MP by replacing all existing CPUs. Two servers can then be connected together for a 16-way.						

The Xeon MP processor-based x445 models support processor configurations of two, four, eight and 16 (two eight-way servers) processors. The x445 models that have Xeon DP processors only support processor configurations of two or four DP processors, but these models can also be upgraded to use Xeon MP processors, if desired.

Restriction: xSeries 440 servers cannot be upgraded to x445 configurations. An x440 cannot be connected together with an x445 to form a 16-way configuration.

Figure 1-1 shows the available single-node configurations and the possible CPU and memory configurations.

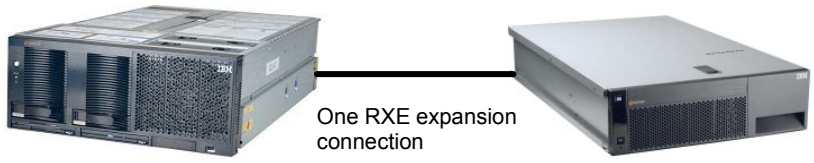

xSeries 445
- Two Xeon DP processors, 2-32 GB
- Four Xeon DP processors, 2-64 GB
- Two Xeon MP processors, 2-32 GB
- Four Xeon MP processors, 2-32 GB
- Eight Xeon MP processors, 2-64 GB

RXE-100
- 6 PCI-X slots
- 12 PCI-X slots

Figure 1-1 x445 configurations

The attachment of a single RXE-100 Remote Expansion Enclosure is also supported, as shown in Figure 1-1. Furthermore, the PCI-X slots in a single RXE-100 can be shared between two x445 configurations. See 3.3.3, “RXE-100 connectivity” on page 76 for more details.

The RXE-100 has six PCI-X slots standard, upgradable to 12 PCI-X slots, giving the customer up to a total of 12 PCI-X or 18 PCI-X slots, respectively.

Restriction: The x445 SMP Expansion Port cannot be used as a high-speed interconnect for clustering purposes. Gigabit Ethernet is the preferred industry-standard interconnect for clustering.

1.1.1 16-way configurations

Two x445 servers can be configured together to form a single 16-way system for use with Microsoft Windows Datacenter Server or VMware ESX Server only. This is achieved by connecting two eight-way x445s together. Other 16-way configurations (such as four four-way servers) are not supported.

Tip: The only multi-chassis configurations supported are a 16-way comprised of two identical eight-way servers, and, in the future, a 32-way comprised of four identical eight-way servers. No others (for example, two four-way servers to form an eight-way) are currently supported.

Once the two x445 chassis are configured in a 16-way configuration, the secondary server does not require disks or a RAID controller installed.

To form a 16-way system, you will need the x445 Two Chassis 16-way Configuration Kit, part number 02R2013. The kit contains the following items:

- Four 2.5 m copper-colored scalability cables
- One Ethernet cross-over cable to connect the chassis' system management adapters

The 16-way configuration kit will also be used as part of the 32-way upgrade, once the 32-way configuration is supported.

Note: Unlike the x440, no specific 16-way systems are orderable. You will need to order the necessary components (SMP Expansion Modules, processors, etc.) to build two eight-way systems and combine them with the 16-way configuration kit to form the 16-way system.

Like the single-node configurations, an RXE-100 may also be attached to a two-node 16-way configuration. The RXE-100 must have the second six-pack of PCI-X slots installed. This is shown in Figure 1-2.

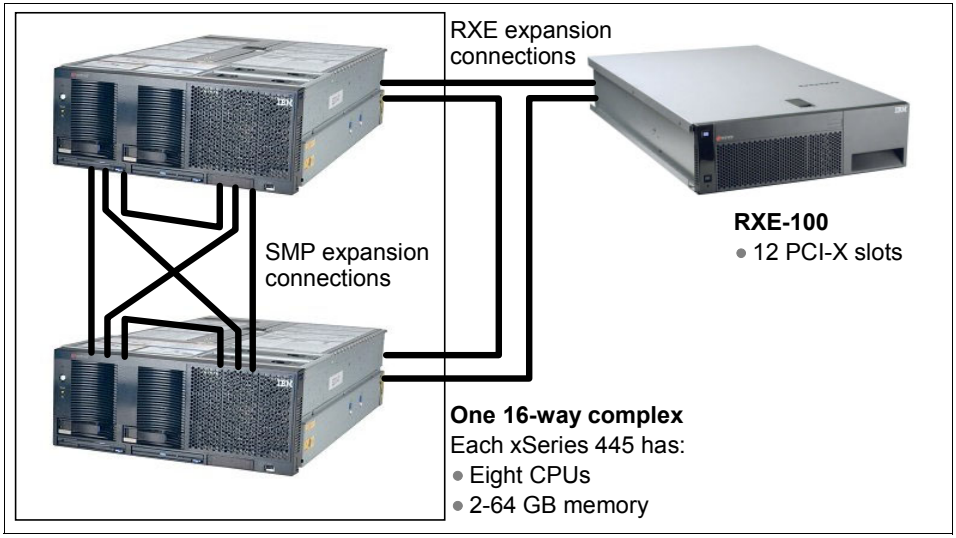

Figure 1-2 16-way configuration

A more detailed diagram of connectivity is shown in Figure 1-6 on page 10.

1.1.2 Windows Datacenter models

For users of Windows 2000 Datacenter Server and Windows Server 2003 Datacenter Edition, Table 1-2 shows the certified models that are available.

Table 1-2 Datacenter Server certified models

Model	Standard processors	Max SMP	Max SMP two nodes	Std disk & memory
8870-1AX	2x 2.0 GHz Xeon MP, 1 MB L3 cache	8-way	16-way	None
8870-2AX	2x 2.5 GHz Xeon MP, 1 MB L3 cache	8-way	16-way	None
8870-4AX	2x 2.8 GHz Xeon MP, 2 MB L3 cache	8-way	16-way	None

When you order these models you specify which operating system you want preloaded, as described in 3.6.1, “Windows Datacenter” on page 95.

1.1.3 32-way configurations

In late 2003, IBM plans to announce support for 32-way configurations, comprised of four eight-way x445 chassis connected together to form a single system. The concept of the connections between the four systems is shown in Figure 1-3. As you can see the 32-way is effectively cabled as two 16-way systems.

Like the 16-way configuration, the 32-way supports connectivity to the RXE-100 for additional PCI-X slots. One or two RXE-100 units are supported, however, only certain models of the RXE-100 are supported when connected to 32-way configurations.

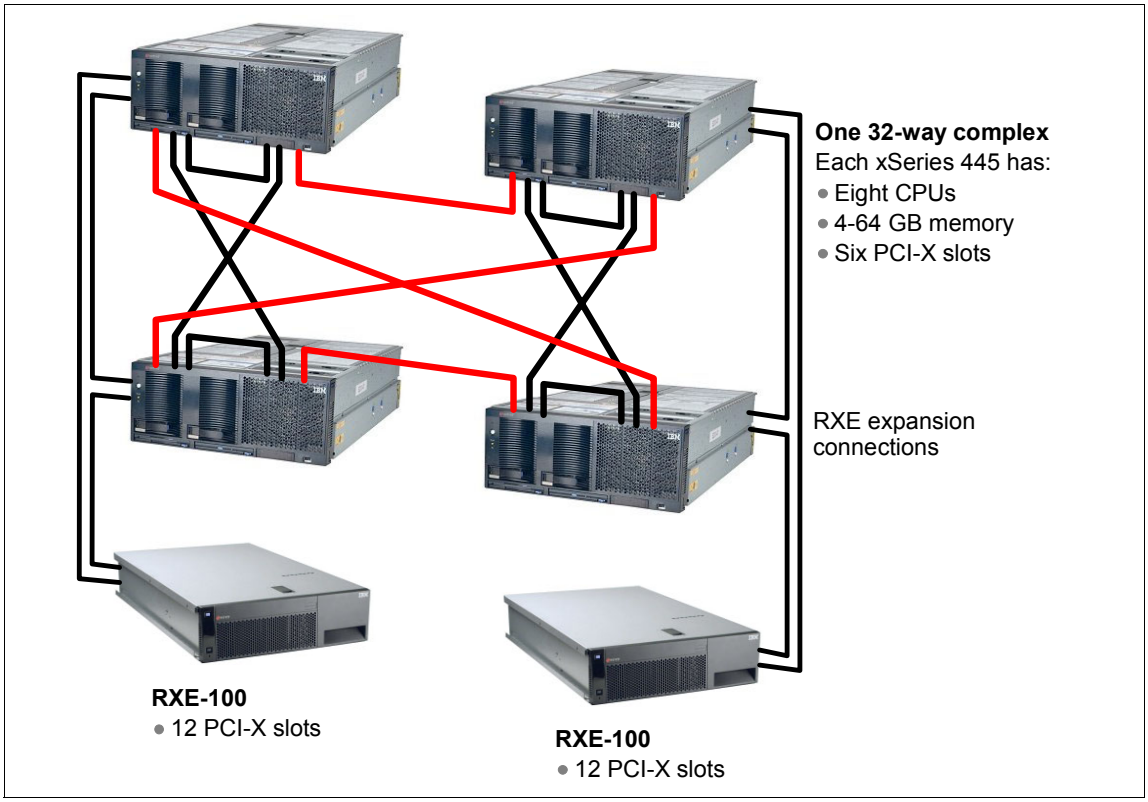

Figure 1-3 Conceptual layout of a 32-way x445 configuration with two RXE-100s

1.2 System partitioning

Partitioning is the ability to divide a system to support multiple operating system images simultaneously. The benefits of system partitioning include:

- Hardware consolidation
- Clustering, such as with DB2® or Oracle RAC
- Software migration and coexistence
- Version control
- Development, testing and maintenance
- Workload isolation
- Resource optimization around a particular application and operating system combination

- Independent backup and recovery on a partition basis

There are two types of system partitioning: physical partitioning (hardware-based, available with the introduction of Scalable Systems Management) and logical partitioning (software-based, enabled with VMware ESX Server):

- Physical partitioning

  This form of partitioning will be available with the introduction of Scalable Systems Management, a plug-in to IBM Director.

  Scalable Systems Manager is an extension of IBM Director that you can use to create and manage static hardware partitions on xSeries 445 servers. Scalable Systems Manager includes a server component and a console component that are installed on the corresponding IBM Director Server or IBM Director Console (there is no agent component). The console component of Scalable Systems Manager provides three tasks you can use:

  - A configuration task for viewing and creating static hardware partitions
  - A partition discovery task
  - A partition validation task

  As an administrator, you can use Scalable Systems Manager to configure a scalable partition for a specific offline server before starting the operating system for that server. After you power on a scalable partition created with Scalable Systems Manager, you can use IBM Director to perform administrative tasks for the resulting managed system such as event listening, security, and management console support.

  Scalable Systems Manager is described in detail in 6.2, “Scalable Systems Manager” on page 210.

- Logical partitioning

  Using logical partitioning, administrators can partition a multi-node complex at the individual processor level (with associated memory, I/O and other required resources) or even lower (that is, multiple partitions per processor) without shutting down and restarting the hardware and software.

  VMware ESX Server is the IBM-supported choice to enable logical partitioning. The x445 supports ESX Server 2.0. See 3.6.7, “VMware ESX 2.0 Server” on page 101 for more information.

1.3 IBM XA-32 second-generation chipset

The x445 utilizes the second-generation of the IBM XA-32 chipset (previously known by its code name “Summit®”). With an increase in XceL4 Server Accelerator Cache (now 64 MB per SMP Expansion Module) it maximizes CPU throughput by reducing the need for main memory access under demanding workloads, resulting in an overall enhancement to system performance.

The architecture consists of the following components:

- Xeon MP or Xeon DP processors
- XceL4 Server Accelerator Cache
- One or two SMP Expansion Modules (one standard)
- Two PCI-X host-bridge controllers

Figure 1-4 shows the various IBM XA-32 second-generation components in a four-way x445 configuration.

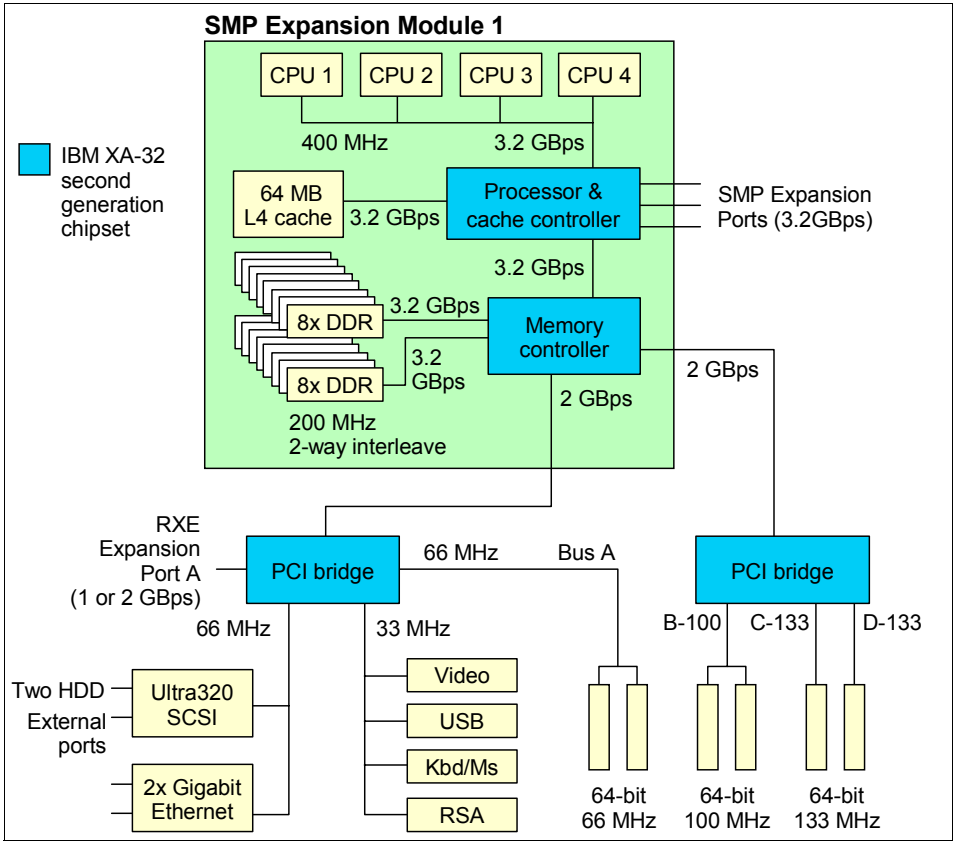

Figure 1-4 xSeries 445 system block diagram — one SMP Expansion Module

The component that contains the CPUs, processor/cache controller, memory controller, memory, and cache is called the *SMP Expansion Module* (or central electronics complex—CEC). The Xeon MP-based models of the x445 will ship with one SMP Expansion Module with two or four CPUs and a total of 2 GB of RAM. The Xeon DP-based models will have either two CPUs in one SMP Expansion Module or four CPUs in two SMP Expansion Modules and a total of 2 GB of RAM.

When two SMP Expansion Modules are installed, they are connected together using two 3.2 GBps SMP Expansion Ports. The third scalability port is not used in this single-node eight-way configuration. The block diagram with two SMP Expansion Modules is shown in Figure 1-5.

The two SMP Expansion Ports provide a cumulative bandwidth of 6.4 GBps between the SMP Expansion Modules.

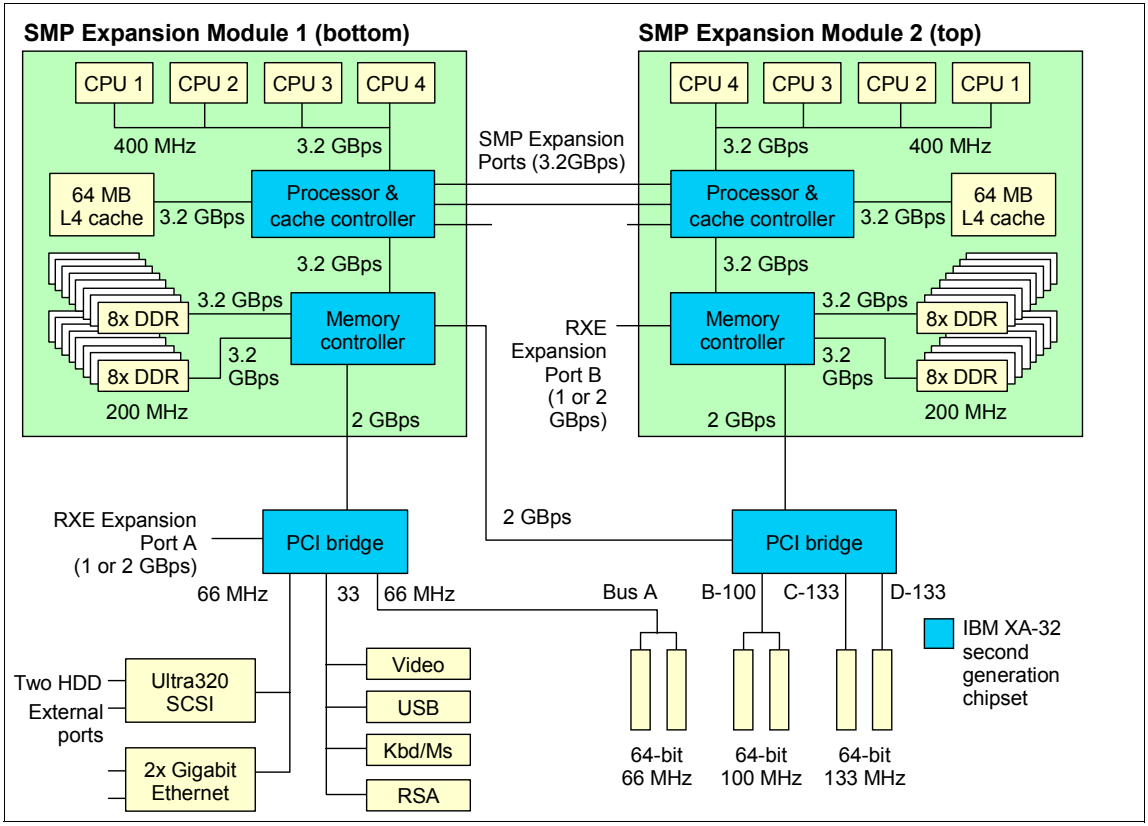

Figure 1-5 xSeries 445 system block diagram — two SMP Expansion Modules

The PCI bridge also has a bi-directional Remote Expansion I/O port (RXE port) for connectivity to the RXE-100 enclosure. This port is labeled “RXE Expansion Port A” in both Figure 1-4 on page 8 (four-way) and Figure 1-5 (eight-way). The RXE-100 provides up to an additional 12 PCI-X slots. When the second SMP Expansion Module is installed to form an eight-way system (Figure 1-5), the memory controller of the second SMP Expansion Module connects internally to the second RXE port, labeled “RXE Expansion Port B”.

The bandwidth of the connections between the x445 and the RXE-100 is 1 Gbps or 2 Gbps, depending on the model of the RXE-100:

- 8684-1RX supports 1 Gbps
- 8684-2RX supports 2 Gbps

When configured in a 16-way configuration, two x445 nodes are connected together using all three SMP Expansion Ports, as shown in Figure 1-6 on page 10.

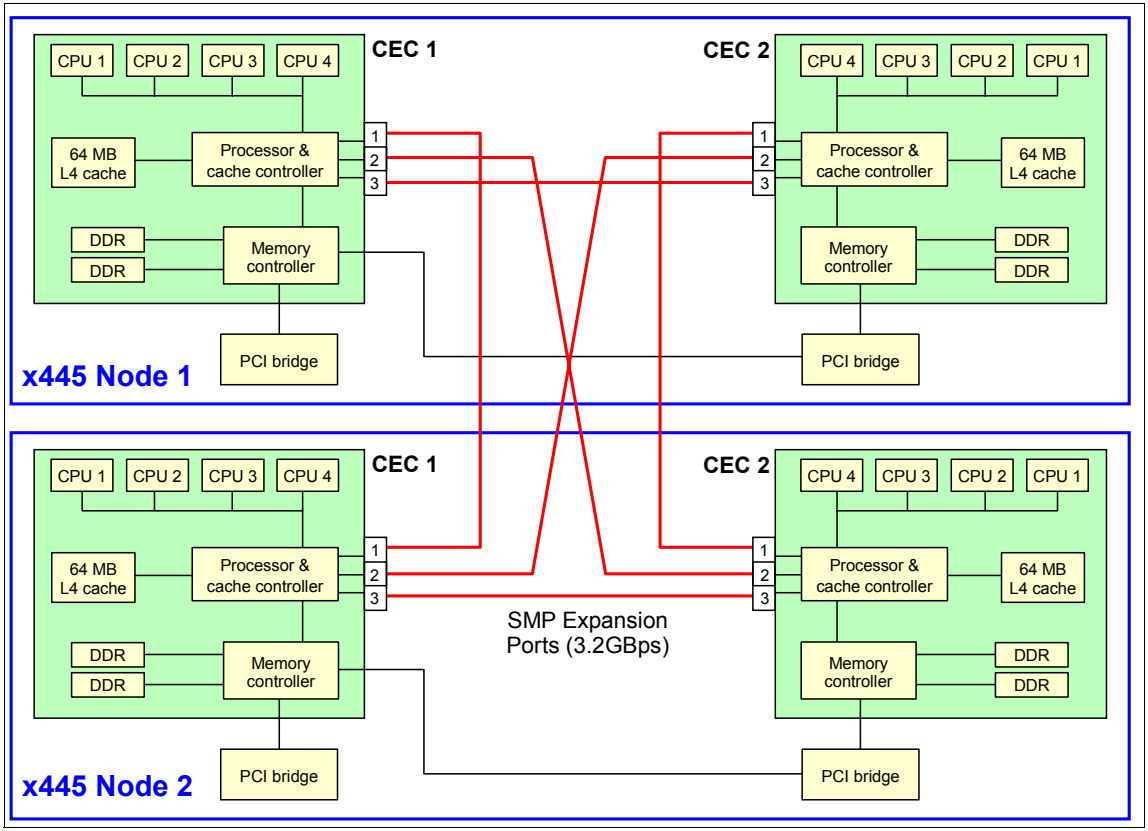

Figure 1-6 16-way configuration (two eight-way x445 servers)

The rear panel of the x445, indicating the location of the SMP Expansion Ports and RXE Expansion Ports, is shown in Figure 1-7.

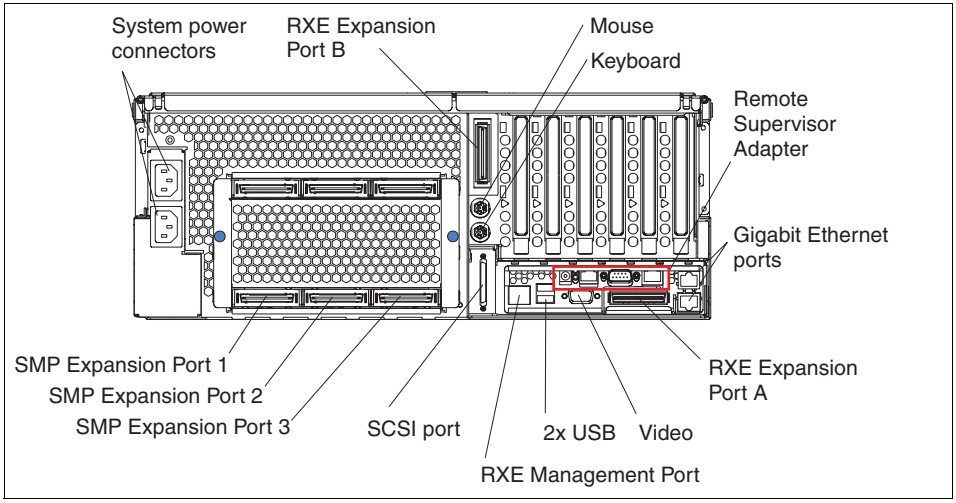

Figure 1-7 Rear view of the x445 with two SMP Expansion Modules installed

1.4 Processors

The x445 models use one of the following processors:

- Xeon Processor MP (“Gallatin”)
- Xeon Processor DP (“Prestonia”)

The Xeon MP models of the x445 come with two or four processors installed in the standard SMP Expansion Module. Up to four processors are supported in the standard module and, with the addition of a second SMP Expansion Module, up to eight processors can be installed in an x445.

The x445 entry-level systems can be ordered with either two Xeon DP processors in a single SMP Expansion Module or with four Xeon DP processors in two SMP Expansion Modules. There is no further upgrade beyond four Xeon DP processors, other than replacing them with Xeon MP processors.

See 3.1.1, “Processors” on page 58 for a further discussion about what you should consider before implementing an x445 solution.

1.4.1 Intel Xeon Processor MP

The Xeon Processor MP (code named “Gallatin”) uses the ZIF socket design instead of the Slot 2 cartridge design of the Pentium® III Xeon processors. This smaller form factor means that the x445 can have up to eight processors in a 4U chassis.

The Xeon MP processor has three levels of cache, all of which are on the processor die:

- Level 3 cache is equivalent to L2 cache on the Pentium III Xeon. The x445 processors contain either 1 MB or 2 MB L3 cache.
- Level 2 cache is equivalent to L1 cache on the Pentium III Xeon and is 512 KB in size in “Gallatin” processors. The L2 cache implements the Advanced Transfer Cache technology, which means L2-to-processor transfers occur across a 256-bit bus in only one clock cycle.
- A new level 1 execution trace cache, 12 KB in size, is “closest” to the processor and is used to store micro-operations (that is, decoded executable machine instructions); it serves those to the processor at rated speed. This additional level of cache saves decode time on cache hits. There is an additional 8 KB data cache for data related to those instructions, which is clocked at twice the speed of the processor’s clock rate

For more information, see the Intel Technology Overview white paper, available from:

http://www.intel.com/eBusiness/pdf/prod/server/xeon/ds020903.pdf

Intel has also introduced a number of features associated with its NetBurst micro-architecture. These are available in the x445, including:

- 400 MHz frontside bus

  The older Pentium III Xeon processor had a 100 MHz frontside bus that equates a burst throughput of 800 MBps. With protocols such as TCP/IP, this has been shown to be a bottleneck in high-throughput situations. The Xeon Processor MP improves on this by using two 100 MHz clocks, out of phase with each other by 90° and using both edges of each clock to transmit data. This is shown in Figure 1-8 on page 12.

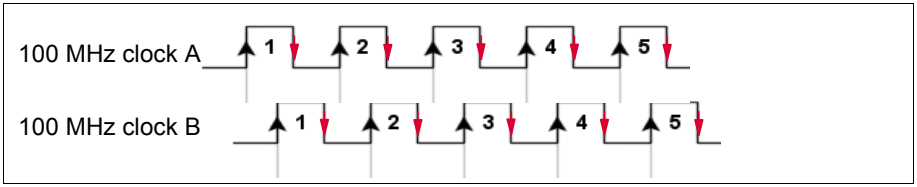

Figure 1-8 Quad-pumped frontside bus

This increases the performance of the frontside bus without the difficulty of high-speed clock signal integrity issues. The end result is an effective burst throughput of 3.2 GBps, which can have a substantial impact, especially on TCP/IP-based LAN traffic.

- Hyper-Threading

  Hyper-Threading technology enables a single physical processor to execute two separate code streams (threads) concurrently. To the operating system, a processor with Hyper-Threading appears as two *logical* processors, each of which has its own architectural state - that is, its own data, segment and control registers, and its own advanced programmable interrupt controller (APIC).

  Each logical processor can be individually halted, interrupted, or directed to execute a specified thread, independently from the other logical processor on the chip. Unlike a traditional two-way SMP configuration that uses two separate physical processors, the logical processors share the execution resources of the processor core, which include the execution engine, the caches, the system bus interface, and the firmware.

  Note: As discussed in 3.1.1, “Processors” on page 58, Hyper-Threading is disabled by default on the x445, but can be easily enabled in the BIOS.

  Hyper-Threading technology is designed to improve server performance by exploiting the multi-threading capability of operating systems, such as Windows 2003 and Linux, and server applications, in such a way as to increase the use of the on-chip execution resources available on these processors.

  Fewer or slower processors usually yield the best gains from Hyper-Threading because there is a greater likelihood that the software can spawn sufficient numbers of threads to keep both paths busy. The following performance gains are likely:

  – Two physical processors: 15-25% performance gain
  – Four physical processors: 1-13% gain
  – Eight physical processors: 0-5% gain

  Figure 1-9 on page 14 shows that two physical processors will outperform one processor with Hyper-Threading enabled.

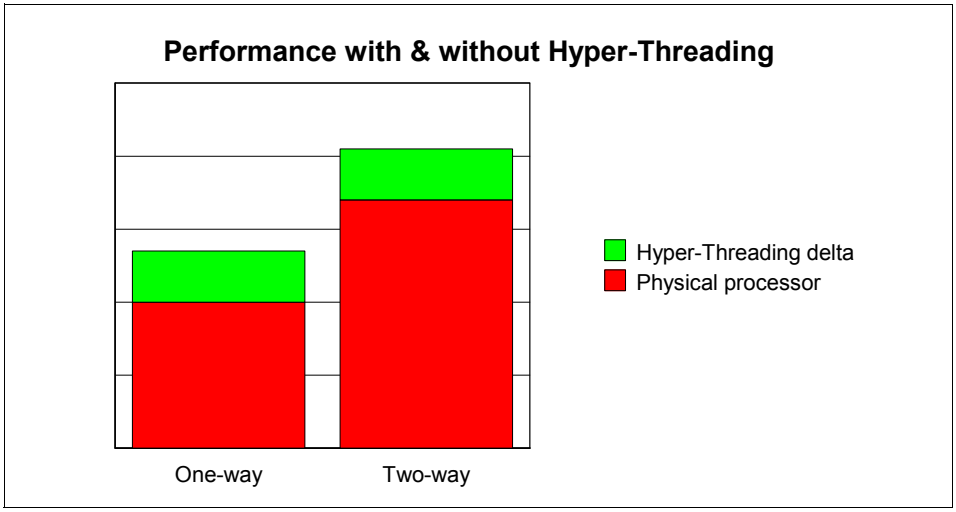

Figure 1-9 Comparing processor performance

Tests have shown that software often limits SMP scalability, but customers should expect improved results as software matures. Best-case applications today are:

- Databases
- Java™
- Web servers
- E-mail

Note: Microsoft licensing of the Windows operating systems is by number of processors.

- For Windows 2000 Server, the logical processors presented by Hyper-Threading affect the count of processors for licensing. If Hyper-Threading is enabled, the operating system will need to be licensed for twice the number of physical processors to take full advantage of the processors' capabilities. Windows 2000 first counts physical processors and, if the license permits more processors, then logical processors will be counted.
- Windows Server 2003, on the other hand, recognizes the difference between physical and logical processors, and licensing only counts physical processors. Hyper-Threading has no affect on processor count for licensing purposes.

See 3.6, "Operating system considerations" on page 92 for details.

For more information about Hyper-Threading, see the following URLs:

http://www.intel.com/technology/hyperthread/

http://www.microsoft.com/windows2000/server/evaluation/performance/reports/hyperthread.asp

- Advanced Dynamic Execution

  The older Pentium III Xeon processor had a 10-stage pipeline. However, the large number of transistors in each pipeline stage meant that the processor was limited to speeds under 1 GHz due to latency in the pipeline.

  The Xeon Processor MP has a 20-stage pipeline, which can hold up to 126 concurrent instructions in flight and up to 48 reads and 24 writes active in the pipeline. The lower complexity of each stage also means that future clock speed increases are possible.

  It is important to note, however, that the longer pipeline means that it now takes more clock cycles to execute the same instruction when compared to the Pentium III Xeon.

The next generations of operating systems will likely improve performance of the Xeon MP processor as they take advantage of the NetBurst architecture. These include Windows Server 2003 and the Linux 2.5/2.6 kernels.

For more information about the features of the Xeon Processor MP, go to:

http://www.intel.com/design/Xeon/xeonmp/prodbref/

The x445 models that have Xeon MP standard will ship with pairs of one of the following:

- Xeon Processor MP 2.0 GHz 1 MB L3 Cache (“Gallatin”)
- Xeon Processor MP 2.5 GHz 1 MB L3 Cache (“Gallatin”)
- Xeon Processor MP 2.8 GHz 2 MB L3 Cache (“Gallatin”)

The Xeon MP x445 models support the following configurations:

- Two-way SMP
- Four-way SMP
- Eight-way SMP (with the addition of a second SMP Expansion Module)
- 16-way SMP (achieved by using two eight-way systems only)

Other combinations (such as one, three, twelve, etc.) are not supported. However, 32-way (4x eight-ways) will be supported in the future.

Certain hardware and operating system restrictions may affect the number of CPUs you are able to use in the x445. See 3.1.1, “Processors” on page 58 and 3.6, “Operating system considerations” on page 92 for details.

1.4.2 Intel Xeon Processor DP

The Xeon DP is similar to the Xeon MP and is also based on the Intel NetBurst microarchitecture. The Xeon DP was designed by Intel to be suitable only in uniprocessor and two-way SMP processor systems. However, with the use of the IBM XA-32 second-generation chipset, the x445 can have up to four Xeon DP processors installed. The Xeon DP models of the x445 models use 3.0 GHz processors.

Tip: The 3.0 GHz Xeon DP used in the x445 is not the same as the 3.06 GHz Xeon DP that is used in the x235, x335, and x345.

The key differences between the processors are listed in Table 1-3 on page 16.

Table 1-3 Differences between the Xeon DP and the Xeon MP

Feature	Xeon Processor DP	Xeon Processor MP
Maximum CPUs per SMP Expansion Module	Two	Four
Maximum CPUs per x445 node	Four	Eight
Supported in multi-node configurations	No	Yes
Core frequency (x445 models)	3.0 GHz	2.0, 2.5, 2.8 GHz
Level 2 cache	512 KB	512 KB
Level 3 cache	None	1 MB, or 2 MB

For more information about the features of the Xeon Processor DP, go to:

http://www.intel.com/design/xeon/prodbref

Tip: To upgrade a two-way x445 Xeon DP server to four-way, add SMP Expansion Module 02R1871, as described in 1.5, “SMP Expansion Module” on page 16. This module is populated with two 3.0 GHz Xeon DP processors.

1.5 SMP Expansion Module

The SMP Expansion Module is the central electronics complex that contains the processors, memory, XceL4 system cache, and respective controllers for these components. All x445 models come standard with one SMP Expansion Module. Each SMP Expansion Module contains slots for up to four Xeon MP processors (or two Xeon DP processors) and 16 DIMMs.

Each SMP Expansion Module contains:

- Four processor connectors
- 64 MB of XceL4 Server Accelerator Cache
- 16 memory slots, supporting up to 32 GB of ECC DDR SDRAM
- Memory controller that supports:
  - Data flow between the processor, XceL4 system cache, and memory, and to the two PCI-X host-bridge controllers
  - High-speed port to exploit external PCI-X expansion, connecting an optional RXE-100 Remote Expansion Enclosure
  - Chipkill™ ECC memory function
  - Hot-swap and hot-add memory

There are two SMP Expansion Module part numbers for x445 models:

- 02R1870 is used in Xeon MP models. It is “unpopulated,” which means it does not contain any processors or memory. Any of the supported Xeon MP processors can be installed in it, but all installed processors must be identical in speed and cache size.
- 02R1871 is used in Xeon DP models. It contains two 3.0 GHz Xeon DP processors and VRMs and is used to upgrade a two-way Xeon DP x445 to a four-way configuration.

Tip: 02R1871 is also compatible with Xeon MP processors. If you want to upgrade your Xeon DP-based x445 to use Xeon MP processors, you can simply replace the processors and VRMs with supported Xeon MP processors.

The SMP Expansion Module is installed from the top of the server and mounts to the side of the centerplane using two levers on the top, as shown in Figure 1-10. These same levers are used to open the SMP Expansion Module when adding additional processors or memory to that SMP Expansion Module.

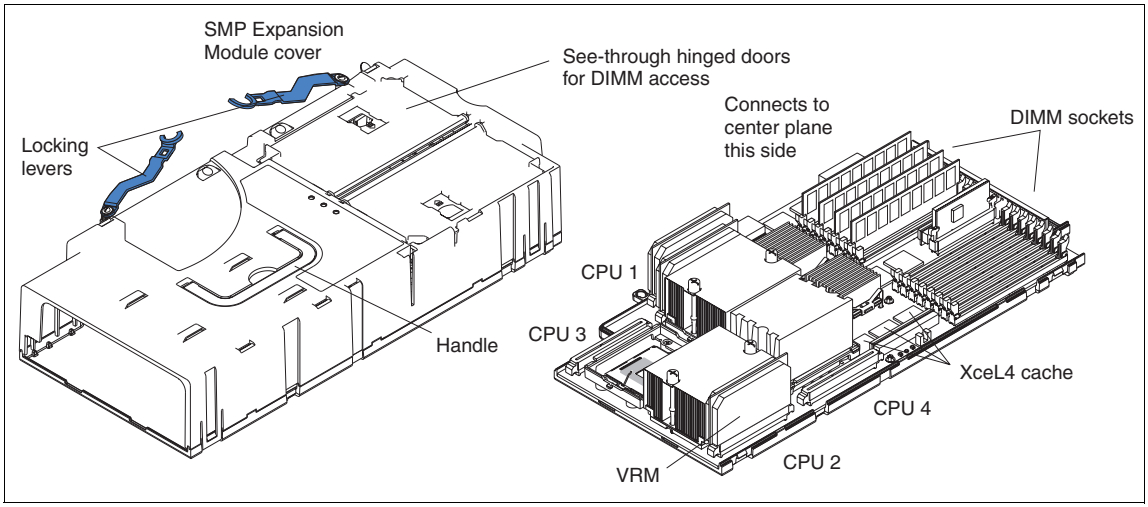

Figure 1-10 SMP Expansion Module (with Xeon MP processors installed)

By adding a second optional SMP Expansion Module, customers can grow the base four-way Xeon MP to an eight-way SMP system, or the base two-way Xeon DP to a four-way SMP system, while still retaining the rack density of the 4U form factor. The second SMP Expansion Module also contains an additional 16 DIMM slots to take the memory up to a maximum of 64 GB (using 2 GB DIMMs) and an additional 64 MB of Level 4 system cache for a maximum of 128 MB per chassis.

When two SMP Expansion Modules are installed, they are connected together using two 3.2 GBps SMP Expansion Ports (also known as scalability ports). Using two connections improves throughput beyond that of one connection and provides load balancing. The third scalability port is not used in this single-node eight-way configuration.

Each SMP Expansion Module is also equipped with the following LEDs for light path diagnostics:

- Individual DIMMs
- Individual CPUs
- Each VRM
- SMP Expansion Module board

Integrated into each SMP Expansion Module is 64 MB of high-speed level 4 cache. This XceL4 Server Accelerator Cache provides the necessary extra level of cache to alleviate the bottlenecks caused by memory contention across eight CPUs.

The XceL4 Cache memory is 200 MHz DDR memory and is faster than standard memory because it is directly connected to the memory controller and does not have the additional latency associated with the large fan-out necessary to support the 16 DIMM slots.

1.6 System memory

All models of the x445 have 2 GB standard, implemented as four 512 MB PC266 (also known as PC2100) ECC DDR DIMMs. There are 16 DIMM sockets (two ports of eight) in each of the two SMP Expansion Modules, for a total of 32 sockets. Using 2 GB DIMMs means that each x445 eight-way can have up to 64 GB RAM.

The memory is two-way interleaved (meaning that memory DIMMs are installed in pairs) with two ports to the memory controller supporting up to 6.4 GBps data transfers. See 3.1.2, “Memory” on page 60 for a further discussion of how memory is implemented in the x445 and what you should consider before an x445 installation.

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.

  It works somewhat like hot-spare disk sectors in the Windows NTFS file system, where if the operating system detects bad sectors on disk, it will write the data to spare sectors set aside for that purpose. Think of Memory ProteXion as providing hot-spare bits. The error correction is handled by the memory controller, so there is no operating system overhead or support requirement

  Currently, other industry-standard servers use 8 bits of the 72-bit data packets for ECC functions and the remaining 64 bits for data. However, the x445 needs only 6 bits to perform the same ECC functions, thus leaving 2 bits free. In the event that a chip failure on the DIMM is detected by memory scrubbing, the memory controller can re-route data around that failed chip through the spare bits (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.

  Memory scrubbing and Memory ProteXion work in conjunction with each other and do not require memory mirroring to be enabled to work properly.

  When a bit error is detected, memory scrubbing determines if the error is recoverable or not. If it is recoverable, Memory ProteXion is enabled and the data that was stored in the damaged locations 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 replaced. If hot-swap is enabled in the BIOS then no rebooting would be required and the new DIMM would 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 (see Figure 3-1 on page 61). 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.

  Certain restrictions exist with respect to placement and size of memory DIMMs when memory mirroring is enabled. See 3.1.2, “Memory” on page 60 for details.

- Hot-swap and hot-add memory

  There are two configurations where you can add or replace memory while the server is still running:

  – Hot-swap, where you can replace failed DIMMs of the same type, size, and clock speed without turning off the server.

  – Hot-add, where you can add new DIMMs without turning off the server, thereby increasing the amount of RAM available to the operating system. This feature is currently only supported by Windows Server 2003, Enterprise Edition and Datacenter Edition.

There are specific configuration rules and restrictions when implementing hot-add and hot-swap. See 3.1.2, “Memory” on page 60 for more information.

- Chipkill memory

  Chipkill is integrated into the XA-32 second-generation chipset and works with industry-standard ECC 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 16 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.

  Chipkill memory comes into play only if a server encounters so many errors in a short span of time that Memory ProteXion can’t handle them all. This should be a rare occurrence, but if it does happen you are still protected. Chipkill memory provides correction for up to four bits per DIMM (eight bits per memory controller), whether on a single chip or on multiple.

  If a memory 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.

  In essence, each DIMM acts as a separate memory array. If any one chip fails, it affects only a single bit from a byte of data, because the other bits are stored on other, working chips. The controller is then able to reconstruct the "missing" bit from the failed chip and continue working as usual.

  Like Memory ProteXion, Chipkill support is provided in the memory controller and implemented using standard ECC DIMMs, so it is transparent to the operating system.

In addition, to maintain the highest levels of system availability, if a memory error is detected during POST or memory configuration, the server can automatically disable the failing memory bank and continue operating with reduced memory capacity. You can manually re-enable the memory bank after the problem is corrected, via the Setup menu in the BIOS.

Memory mirroring, hot-swap, hot-add, Chipkill, and Memory ProteXion provide multiple levels of redundancy to the memory subsystem. Combining Chipkill with Memory ProteXion enables up to two memory chip failures per memory port (8 DIMMs) on the x445. A16-way system (two eight-ways) with its eight memory ports could sustain up to 16 memory chip failures (if there were two failures per memory controller).

1.7 PCI subsystem

As shown in Figure 1-4 on page 8, there are six PCI-X slots internal to the x445. These each support 3.3 V, 32-bit or 64-bit, PCI or PCI-X adapters. The slots support adapters of different speeds:

- Two 133 MHz slots, each on separate PCI-X buses, which accept adapters with speeds 33-133 MHz
- Two 133 MHz slots on the same bus, which accept two adapters with speeds from 33-100 MHz (two 133 MHz adapter will each operate at 100 MHz), or one adapter at 133 MHz if the second slot is left empty.
- Two 66 MHz slots, which accept adapters with speeds of 33 or 66 MHz

See 3.1.3, “Active PCI-X subsystem” on page 65 for details on what adapters are supported and in what combinations.

The PCI subsystem also supplies these I/O devices:

- LSI Logic dual Ultra320 SCSI with integrated RAID-1
- Broadcom dual port 5704 10/100/1000 Ethernet
- Remote Supervisor Adapter installed in dedicated slot
- SVGA (ATI Rage XL, 8 MB)
- EIDE interface
- Serial port (accessible using the PCI-slot bracket/cable assembly). If a PCI slot is not available for this assembly, use a USB-to-serial converter. See 3.3.4, “Serial connectivity” on page 84 for details.
- Three USB ports (one on front panel, two on rear)

Note: There is no parallel port on the x445.

With the addition of an RXE-100 Remote Expansion Enclosure, you can connect an additional six or 12 PCI-X adapters to the x445. See 3.3.3, “RXE-100 connectivity” on page 76 for details.

1.8 Remote Supervisor Adapter II support

The x445 models have an Remote Supervisor Adapter standard which provides in-band and out-of-band systems management. See 3.3.2, “Remote Supervisor Adapter connectivity” on page 75 for more information.

A follow-on to the Remote Supervisor Adapter, the Remote Supervisor Adapter II, offers significantly improved remote control features, letting the administrator remotely access and control the x445 via the service processor. The Remote Supervisor Adapter II lets you remotely access the power-on messages, Setup utility, and, depending on the operating system, full graphical, operating system access with mouse and keyboard support.

In November 2003, IBM announced an option, part number 13N0382, to allow existing x445 models to be upgraded to the Remote Supervisor Adapter II, replacing the existing Remote Supervisor Adapter. This option must be installed by an IBM service technician or authorized business partner.

1.9 Redundancy

The x445 has the following redundancy features to maintain high availability:

- Four hot-swap multi-speed fans

  With four hot-swap redundant fans, the x445 has adequate cooling for each of its major component areas. There are two fans located at the front of the server that direct air through the SMP Expansion Modules. These fans are accessible from the top of the server without having to open the system panels. In the event of a fan failure, the other fan will speed up to continue to provide adequate cooling until the fan can be hot-swapped by the IT administrator.

  The other two fans are located just behind the power supplies and provide cooling for the I/O devices. Similar to the SMP Expansion Module fans, these fans will speed up in the event that one should fail to compensate for the reduction in air flow. In general, failed fans should be replaced within 24 hours following failure.

  Important: Due to airflow requirements, fans should not be removed for longer than two minutes. The fan compartments need to be fully populated even if the fan is defective. Therefore, remove a defective fan only when a new fan is available for immediate replacement.

- Two hot-swap power supplies with separate power cords.

  Note: For large configurations, redundancy is achieved only when connected to a 220 V power supply. See 3.9, “Power considerations” on page 110 for details.

- Two hot-swap hard disk drive bays. Together with the integrated LSI Logic 1030 dual Ultra320 SCSI, they can be configured to create a RAID-1 disk array for the operating system.
- In eight-way systems, the two SMP Expansion Modules are connected together using two SMP Expansion Ports, as shown in Figure 1-5 on page 9. These two connections provide load balancing and performance. Furthermore, built-in Copper Diagnostics improves ease of use and eliminates single points of failure.
- The memory subsystem has a number of redundancy features, including memory mirroring, as described in 1.6, “System memory” on page 19.

The layout of the front panel of the x445, showing the location of the drive bays, power supplies, and fans, is shown in Figure 1-11.

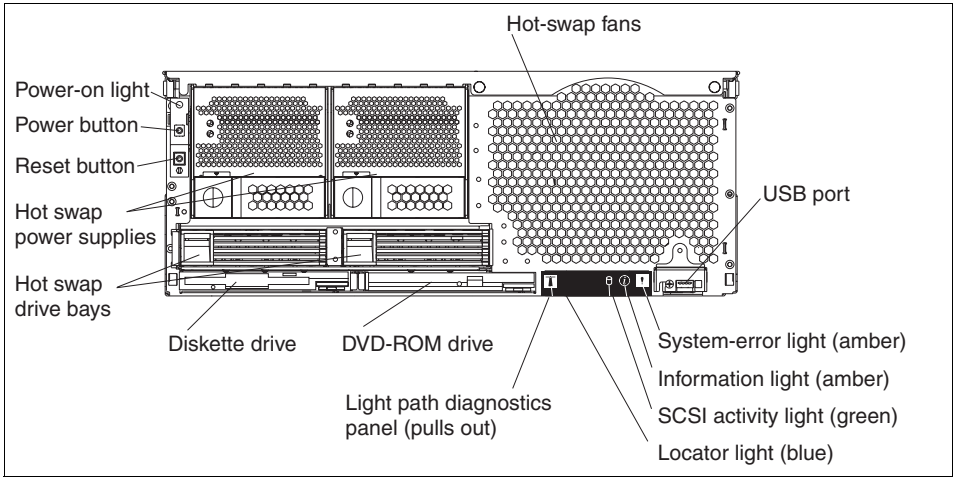

Figure 1-11 Front panel of the x445

1.10 Light path diagnostics

To limit the need to slide the server out of the rack to diagnose problems, a new light path diagnostics panel has been added to the front of the x445. This panel can be ejected from the server to view all light path diagnostics-monitored server subsystems. In the event that maintenance is then required, the customer can slide the server out from the rack and using the LEDs, find the failed or failing component.

As illustrated in Figure 1-12 on page 25, light path diagnostics is able to monitor and report on the health of CPUs, main memory, hard disk drives, PCI-X and PCI slots, fans, power supplies, VRMs, and the internal system temperature.

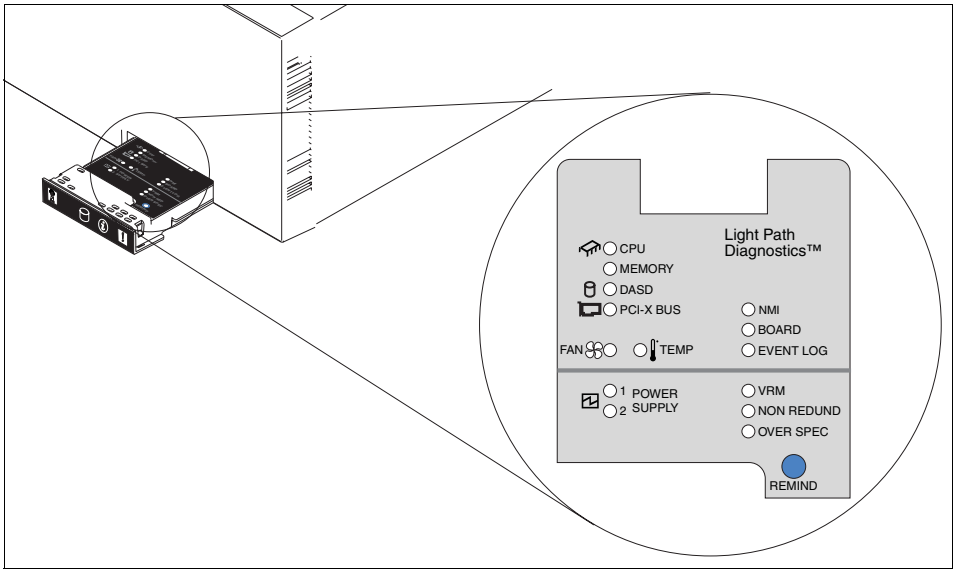

Figure 1-12 Light path diagnostics panel on the x445

The light path diagnostics on the x445 has three levels:

1. Level 1 is the pop-out panel as shown in Figure 1-12.
2. For further investigation, there are light path diagnostics LEDs visible through the top of the server. This requires the server to be slid out of the rack.
3. For the third level of diagnostics, LEDs on the planar indicates the component causing the error.

The pop-out panel (Figure 1-12) also has a Remind button. This places the front panel system-error LED into remind mode, which means it flashes briefly every 2 seconds. By pressing the button, you acknowledge the failure but indicate that

you will not take immediate action. If a new failure occurs, the system-error LED will turn on again. The system-error LED remains in the Remind mode until one of the following situations occurs:

- All known problems are resolved.
- The system is restarted.
- A new problem occurs, at which time it then is illuminated continuously.

2

Positioning

In this chapter we discuss topics that help you understand how the x445 can be useful to your business and what is the best configuration to use. The topics covered are:

2.1 xSeries 445 application solutions

The IBM @server xSeries 445 is the next-generation scalable enterprise server based upon the latest Enterprise X-Architecture technology. As the flagship of the high-performance xSeries server brand, the x445 sets a new standard in the marketplace for modular scalability, industry-leading performance, mission-critical availability, and remote management. The x445 builds upon the success of the x440 and extends that leadership with new capabilities including improved partitioning, reduced latencies, increased memory availability, and scalability up to 32-way.

The versatile x445 is ideal for customers running mission-critical applications. There are four key solutions where the x445 will be an extremely attractive platform as either an application server or a highly scalable database server:

- Server consolidation

  With Enterprise X-Architecture, the x445 server offers exceptional flexibility, availability and scalability to handle customer requirements for consolidating distributed workloads onto a single powerful and highly available platform. Most of the Intel opportunities for server consolidation are tied to the consolidation of highly distributed infrastructure applications such as database and messaging/collaboration.

  Recent advances in many of the application software products now make it possible for customers to consolidate onto larger servers with numerous high-availability capabilities. It is important to remember that server consolidation is not a discrete workload, but that it applies to all departmental applications and computing resources that could be centralized, regardless of geographical location.

- Enterprise applications

  Enterprise applications work with the most critical data of a business, so it is a requirement that these applications be highly available and secure. Enterprise applications include enterprise resource planning (ERP), supply chain management (SCM), customer relationship management (CRM) and business intelligence (BI) solutions.

  Most of these applications today utilize a Web-based infrastructure with interfaces to suppliers, customers, and internal company employees. There are three general architectures utilized by enterprise solutions:

  – Four-tier architecture (often referred to as an Internet architecture) with client systems, Web servers, application servers and database servers.

  – Three-tier architecture, which includes client systems, Web/application servers and database servers.

  – Two-tier architecture, which includes client systems and database servers.

The x445 provides an ideal solution as an application server and/or a database server. In addition, several of these enterprise applications will be deployed in a heterogeneous environment where there could be a pSeries® or zSeries® as the back-end database server, making the x445 an ideal application server.

- Infrastructure applications

  These are often thought of as the backbone that ties many enterprise applications together. Some categories within infrastructure applications are database applications, messaging/collaboration and e-business. The two key areas for x445 opportunities present themselves within database and messaging/collaboration applications. In particular, database applications provide the scalability to take full advantage of the x445 server architecture, making the x445 an attractive server platform for enterprise applications.

- Clustering

  A cluster is two or more interconnected servers (sometimes called nodes) that create a solution to provide higher availability, higher scalability or both. The advantage of using a high availability cluster is so that if one node fails, another node in the cluster can assume the workload of the failed node, and users see no interruption of access. The advantages of clustering servers for scalability include increased application performance and a greater number of users that can be supported.

2.1.1 Server consolidation

Server consolidation is a process of centralizing business computing workloads to reduce cost, complexity, network traffic, management overhead and, in general, to simplify the existing IT infrastructure and provide a foundation for new solution investment and implementation.

Server consolidation is discussed in detail in 2.4, “Server consolidation” on page 45.

Server consolidation solutions can be divided into two groups: those where no more than four-way SMP is needed, and those that will take advantage of more CPUs.

- Four-way configurations

  The four-way configurations would most likely be good candidates for traditional messaging/collaboration environments such as Microsoft Exchange and Lotus® Domino®. These applications do not scale well beyond a four-way SMP configuration. It is an optimal platform for customers who intend to migrate from Exchange 5.5 to Exchange 2000 using the new features of Exchange 2000 such as the support for more databases. Many

customers have distributed Exchange and Lotus Domino sites, which is costly and difficult to manage. Here, the x445 can be a very attractive platform to consolidate distributed sites into a central site.

Many ISPs are running different Internet applications and mail systems on several servers. In most cases, they run applications on several servers to get better I/O. The four-way x445 server connected to an RXE-100 fulfills this requirement and ISPs can continue servicing the customers by consolidating to an x445 server.

Although many applications such as file, print, and terminal servers do not scale well beyond two processors, the four-way x445 can be a good platform on which to consolidate those distributed applications. For example, using VMware ESX Server, many file and print servers that are distributed around the enterprise can be consolidated to a four-way x445 server, reducing the TCO.

Using logical partitioning with four-way configurations can produce a one-box cluster solution for small-to-medium-sized businesses (SMB) that need to protect their mission-critical applications and files. With ESX Server, the 4-way server could act like four or more uni-processor servers, each running a different software load. Similarly, developers could test multiple versions of applications, or test one version on multiple operating systems simultaneously on separate virtual servers.

In addition, a four-way x445 can be a good platform for light ERP solutions such as Navision.

- Eight-way and 16-way configurations (32-way in near future):

  The eight-way and 16-way x445 is ideal for customers who want to consolidate their enterprise applications (ERP, CRM, and SCM) or roll out new enterprise applications. These configurations offer computing power, high availability, and reliability, which are the main requirements when running enterprise applications. The goal is to help customers to control their expenses while establishing an environment that is easier to manage because of fewer nodes.

  The eight-way and 16-way configurations are solid platforms to be used for consolidating database applications such as DB2, SQL Server, and Oracle. For instance, a single database that spans multiple servers can be consolidated to an eight-way x445 server or multiple databases on multiple servers can be consolidated to a 16-way complex.

  Many customers have multiple databases distributed on multiple sites and they are planning to migrate to new database versions. This could be a very costly and time-intensive process. The migration process needs to be well planned and tested without any interruption of the business process. The eight-way or 16-way can be an optimal platform for these customers. For

example, you can consolidate the distributed databases on multiple sites to a 16-way x445. Using logical partitioning on x445, you can build, test and deploy many virtual databases on one physical server.

The main reasons to consolidate database applications are:

– Migration from older database versions to new versions getting the advantages relating to availability, reliability and performance.

– Support for more databases. For instance, SQL Server 2000 can support up to 32,767 open databases.

– Reducing the management costs of distributed database sites by consolidating to an easy-to-manage central site.

In addition, using logical partitioning with eight-way and 16-way configurations can produce a powerful server solution that is capable of hosting multiple applications.

2.1.2 Enterprise applications

Because enterprise applications such as ERP, SCM, CRM and BI work with the most critical data of a business, x445 with its high-availability features is an ideal server for these applications.

- Enterprise Resource Planning

  Enterprise Resource Planning (ERP) is an industry term for the broad set of activities supported by multi-module application software that helps a manufacturer or other business manage the important parts of its business, including product planning, parts purchasing, maintaining inventories, interacting with suppliers, providing customer service, and tracking orders. ERP can also include application modules for the finance and human resources aspects of a business. Typically, an ERP system uses or is integrated with a relational database system.

  The key operation areas of the x445 for ERP applications are:

  – As an application server and as a database server with two-way servers such as the x335 acting as Web servers.

  – As an application server front-end to a pSeries or zSeries database server, due to the fact that ERP applications involve integration across heterogeneous environments.

  – Using partitionable x445 servers to deploy ERP applications within a single large-scale server, which could be an attractive solution for SMB customers offering them new levels of manageability as it relates to their ERP implementation.

Key server attributes for ERP applications are availability, scalability, and performance. The x445, with its Enterprise X-Architecture technology such as XpandOnDemand capability, Active Memory, and XceL4 server accelerator cache, is a robust basis to build and implement successful ERP solutions.

Key ERP software vendors include SAP, Oracle, PeopleSoft, Microsoft, JD Edwards, Baan, and Navision.

- Supply Chain Management

  Supply Chain Management (SCM) is the oversight of materials, information, and finances as they move in a process from supplier to manufacturer to wholesaler to retailer to consumer. Supply Chain Management involves coordinating and integrating these flows both within and among companies.

  The x445 is a preferred platform for SCM management applications. The x445 offers a range of leading technologies that will help to deliver the uptime required for business-critical applications at the lowest price/performance ratio. The x445 covers all high-availability features for customers looking for servers to power their SCM solutions. Also, the x445 can be considered as an application server or in a heterogeneous environment as a front-end to a pSeries or zSeries database server.

  Key SCM software vendors include: i2 Technologies, SAP, International Business Systems (IBS), JD Edwards, and PeopleSoft.

- Customer Relationship Management

  Customer Relationship Management (CRM) is an information-industry term for methodologies, software, and usually Internet capabilities that help an enterprise manage customer relationships in an organized way.

  With the Intel Xeon Processor MP and the IBM XceL4 cache, the x445 provides a performance-based foundation upon which customers can build and deploy CRM solutions. The x445 will most likely be implemented as an application server and/or a database server. In addition, the x445's partitioning capabilities will help to build a partitioned CRM environment, allowing customers to maximize server utilization while simplifying overall management of the deployment.

  Key CRM software vendors include: Siebel Systems, Baan, Onyx, PeopleSoft, and SAP.

- Business Intelligence

  Business Intelligence (BI) is a broad category of applications and technologies for gathering, storing, analyzing, and providing access to data to help enterprise users make better business decisions. BI applications include the activities of decision-support systems, query and reporting, online analytical processing (OLAP), statistical analysis, forecasting, and data mining.

The recent move of BI solutions into smaller enterprises has led to the strong positioning of Windows on Intel processor-based servers within this market. The x445 brings scalability and performance to handle compute-intensive BI applications. The highlights of the x445 are its XceL4 cache, which will help speed up data-intensive BI applications that help companies to increase the productivity of their employees.

Key BI software vendors include SAS, Cognos, Business Objects, Hyperion, and Crystal Decisions.

2.1.3 Infrastructure applications

Some of the infrastructure applications are database, messaging/collaboration, and e-business applications. The x445 can be recommended for these three areas as follows:

- Database applications:

  Four-way and eight-way configurations can be used as database servers, and application servers or combination database and application servers providing an extremely scalable platform with room to scale to additional nodes. These configurations require an external storage enclosure or SAN, depending on the size of the database, which is driven by the number of users.

  The 16-way and 32-way configuration can deliver a highly reliable and capable platform for customers who need to run multiple instances of databases that can scale beyond eight processors.

  Key database software vendors include IBM (DB2), Microsoft (SQL Server), and Oracle.

- Messaging/collaboration:

  The four-way x445 with its high-availability features is a good platform for messaging/collaboration applications. Even though there are some scalability limits for Microsoft Exchange 2000 (which does not scale well above four processors), the x445 can be seen as an ideal server for Exchange 2000 deployments.

  Another possible operation area for the x445 in the messaging/collaboration arena is the utilization of partitioning, allowing customers to maximize server resources while improving overall manageability.

  Key messaging/collaboration software vendors include: Lotus (Domino) and Microsoft (Exchange).

- e-business:

  e-business is the use of Internet technologies to improve and transform key business processes.

This includes Web-enabling core processes to strengthen customer service operations, streamlining supply chains and reaching existing and new customers. In order to achieve these goals, e-business requires a highly scalable, reliable, and secure server platform.

The x445 is a strong candidate for an application integration server that integrates the back-end data with the servers containing end-user or client programs. This involves data transformation, process flow, and other capabilities, thus allowing companies to integrate applications and other data sources. These types of servers benefit from the processing power offered by the x445.

Key e-business software vendors include: IBM (WebSphere) and BEA.

2.1.4 Clustering

Clustering can be implemented at different levels of the system, including hardware, operating systems, middleware, systems management and applications. The more layers that incorporate clustering technology, the more reliable, scalable and manageable the cluster.

We will introduce two types of server clustering as below:

- Clustering with physical systems

  This clustering configuration uses two, three, four, up to eight x445 servers, each acting as high availability or high performance cluster nodes.

  For example, with Microsoft Windows 2000 or Windows Server 2003, you can configure a two-node (the most prevalent cluster configuration today) or four-node cluster with SQL Server leveraging SQL Server Failover Clustering and the integrated gigabit Ethernet adapter of the x445. With Windows Server 2003, you will be able to partition up to eight-node configurations with either Windows Server 2003 Enterprise Edition or Datacenter Edition while still leveraging industry-standard Gigabit Ethernet. With Red Hat or SuSE Linux, you can configure a two-node cluster with DB2 or Oracle leveraging SteelEye LifeKeeper. Other multi-node solutions include Oracle 9i Real Application Cluster, Veritas Cluster Server, PolyServe, and Novell Cluster Server.

- Clustering with virtual machine in VMware ESX Server

  In addition, the x445 in conjunction with VMware offers clustering which can be seen as another key solution for server consolidation. For example, a two-node IIS cluster and a two-node file server cluster can be consolidated into a single x445 server. This helps customers to save costs, facilitate cluster management, and improve cluster performance.

The x445 with its high-availability features is an optimal platform to protect mission-critical applications. The x445 offers three types of clustering with virtual machine for server consolidation purposes:

– Cluster in a physical machine

  As shown in Figure 2-1, this provides simple clustering to deal with software crashes or administrative errors. The cluster consists of multiple virtual machines (VMs) on a single physical machine. It supports shared disks without any shared SCSI hardware. It supports the heartbeat network without any extra network adapters.

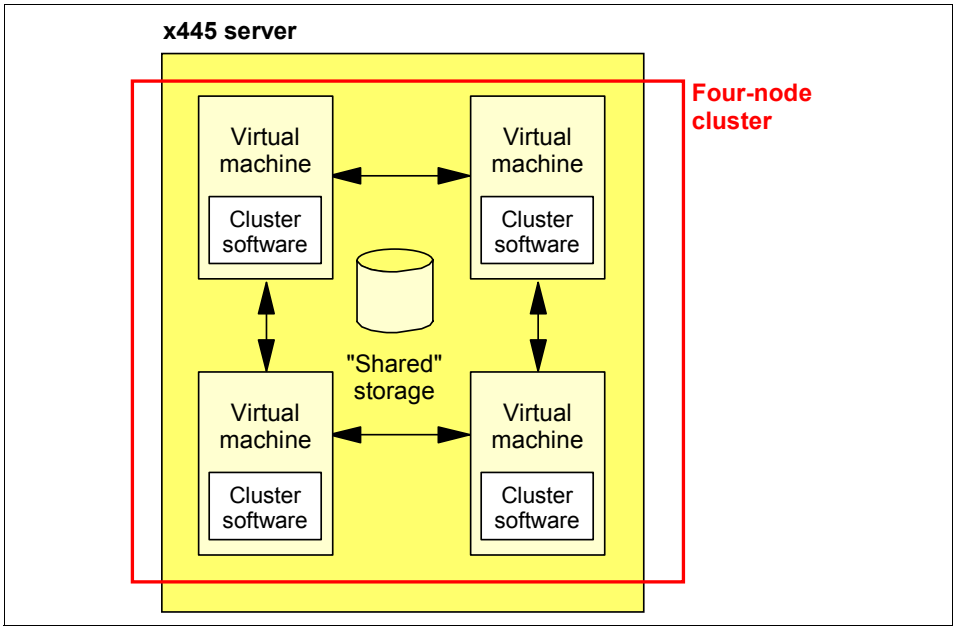

Figure 2-1 Four-node cluster on a single x445 running VMware ESX Server

  This configuration is ideal in circumstances where the primary focus is to consolidate hardware while still providing high availability and performance to the applications. While this configuration permits a single point of failure (the physical server), the x445 minimizes these risks with the advanced Enterprise X-Architecture. Advantages for this type of cluster include reduced hardware administration, since only one server needs to be maintained.

– Cluster across physical machines

  This type of cluster also uses virtual machines, but with this configuration, the virtual cluster nodes are on two separate x445 servers. The virtual disks are stored on real shared disks, so all virtual machines can access

them. Using this type of cluster, you can protect your mission-critical applications in a cost-effective way. For example, you can set up a cluster to protect your Web server applications and you can configure a second cluster to protect your file server.

You can consolidate two clusters of two machines each to two physical machines with two virtual machines each. This provides protection from both hardware and software failures.

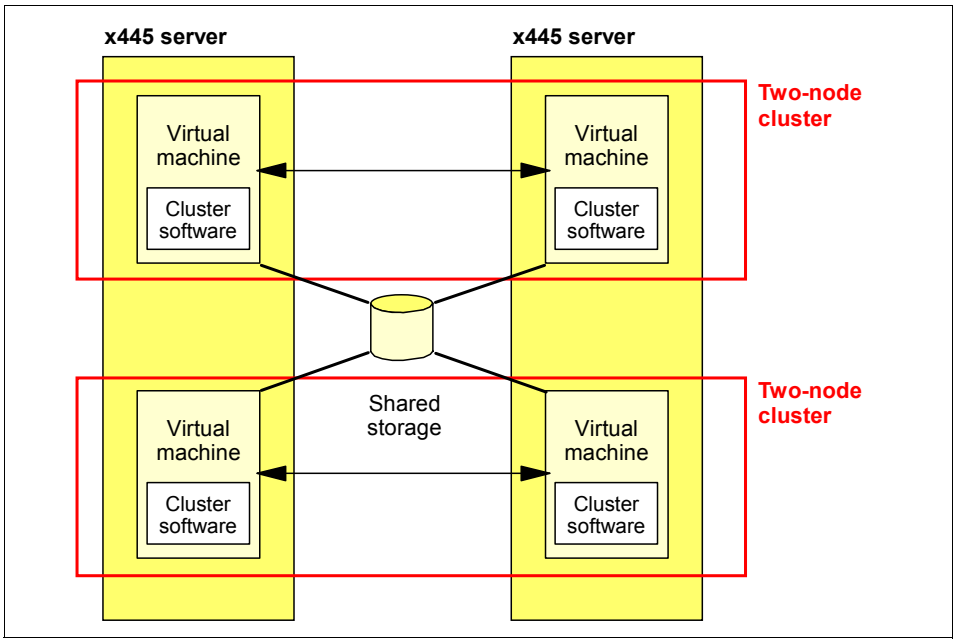

Figure 2-2 Clusters on two x445 servers running VMware ESX Server

A dual-chassis eight-way configuration can be used as clustered database servers and/or application servers in an ERP/CRM/SCM environment, delivering high performance, high availability, and reliability, which are key requirements of enterprise applications. This configuration requires an external storage enclosure or SAN, depending on the size of the database, which is driven by the number of users.

– Cost-effective standby host

Another option would be to cluster several physical servers with several virtual machines as shown in Figure 2-3. This configuration provides for cost-effective standby systems, and provides protection from hardware as well as software failures.

This type of implementation is ideal in situations where you are seeking savings from consolidation, but the availability of the application is of such

a critical nature that minimizing downtime is key. By clustering the primary servers with virtual servers, the hardware costs can be significantly less than traditional clusters, and the system is tolerant of hardware failures.

Advantages of cost-effective standby hosts include:

- Reduction in hardware costs — you are not buying two of everything for each cluster
- Reduced hardware administration — fewer servers to maintain
- Hardware failure tolerance — if the physical server suffers a hardware failure, the virtual server can continue to provide service until the primary is returned to an available state

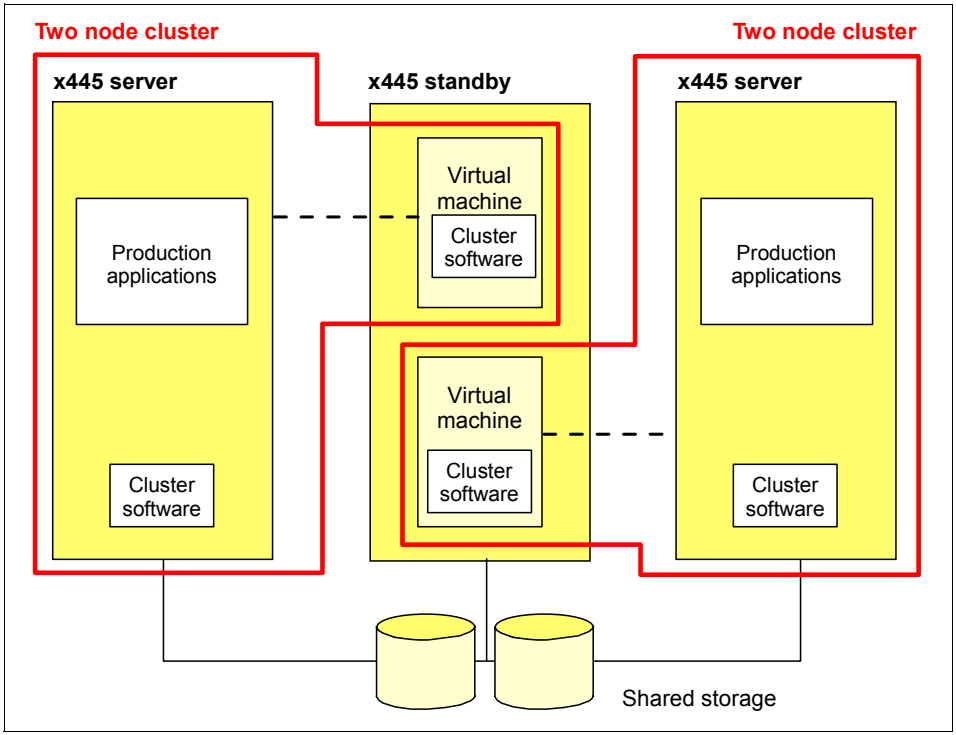

Figure 2-3 Cost-effective standby host

2.2 Why choose the x445?

We have introduced the new technologies of x445 in Chapter 1, “Technical description” on page 1. In this section, we describe the major benefits that you’ll get from these new technologies and the top reasons to choose the x445.

2.2.1 Benefits from new technologies

There are many features of the x445 that provide benefits to customers:

- **XpandOnDemand scalability**
  - Buy what you need today and, with XpandOnDemand, you can expand your system tomorrow
  - XpandOnDemand provides a simple route to higher performance using industry-standard hardware. Each x445 can scale from two-way to eight-way in a single chassis. Grow to 16-way simply by connecting two eight-way chassis or up to 32-way by connecting four chassis. This capability enables you to add processing and memory capacity as your needs change.
- **IBM XA-32 second-generation chipset**
  - The IBM XA-32 second-generation chipset delivers lower latency data communications while powering the latest Intel Xeon MP processors up to 2.8 GHz for higher performance and scalability up to 16-way and beyond.
  - Delivering the industry's only 2-way to 4-way SMP with the Intel Xeon DP 3.0 GHz processor for outstanding price/performance.
- **Next-generation Intel Xeon processors**
  - Extreme scalability up to 16-way today with the Intel Xeon MP processor capable of up to 2.8 GHz with 2 MB L3 cache.
  - Scalability and price/performance with the industry's only four-way leveraging the Intel Xeon DP 3.0 GHz processor.
  - Faster execution with Hyper-Threading technology based upon the Intel NetBurst micro-architecture that supports execution of multiple programs or threads within a single processor.
- **XceL4 Server Accelerator Cache**
  - XceL4 Server Accelerator Cache improves overall performance by reducing memory latency by up to 12% over the previous x440.
  - 64 MB of level-4 system cache per SMP Expansion Module (up to 256 MB per 16-way) maximizes CPU throughput by reducing the need for main memory access under demanding workloads resulting in overall enhancement to system performance.
- **Active Memory**
  - Chipkill and Memory ProteXion, the lower cost, high availability alternative to RAID memory, is standard at no additional cost for the detection and correction of single-bit and multi-bit memory errors.

- New OS-independent, hot-swap memory enables the replacement of failed DIMMs while the system still running to maximize availability of the mission-critical server.
- New hot-add memory supported with Microsoft Windows Server 2003 powers the dynamic addition of main memory to increase performance.
- Memory mirroring dramatically reduces unscheduled downtime for memory failures through high availability redundant memory banks.
- IBM-developed Memory ProteXion uses spare bits for redundancy that re-route data around a failed chip on a DIMM in the event of a failure.
- The x445 offers third-generation Chipkill memory for advanced ECC using off-the-shelf industry-standard DIMMs.
- With 16 DIMM slots per SMP Expansion Module, the x445 can be expanded up to 64 GB per chassis, the maximum supported by any IA-32 operating system.
- Two-way memory interleaving results in easier upgrades in pairs of DIMMs.

- **Active PCI-X and Remote I/O**
  - Six Active PCI-X slots spread across four PCI buses.
  - Optional RXE-100 Remote Expansion Enclosure for an additional 12 slots. 32-way configurations can have two RXE-100s attached.
  - Allows you to hot-add and hot-swap PCI and PCI-X adapters on the fly.
  - Active PCI-X delivers two times (up to 1 GBps) data throughput versus the fastest PCI using 64-bit 133 MHz I/O slots.
  - The second-generation of the XA-32 chipset increase the aggregate I/O bandwidth by 40% over the pervious generation x440, supporting simultaneous bi-directional aggregate I/O bandwidth of 7.0 GB per second per eight-way Xeon MP or four-way Xeon DP, making x445 is one of the highest performance Intel-based servers in the world.
- **Hot-swap Ultra320 SCSI hard disk drives with Integrated RAID**
  - The x445 has support for up to two hot-swap Ultra320 SCSI hard disk drives with the ability to mirror the operating system and swap files use the integrated RAID-1 controller.
  - The x445 supports up to 292 GB of internal storage and terabytes of external storage with FAStT Fibre Channel storage enclosures and the IBM Enterprise Storage System.
- **Light path diagnostics**
  - Light path diagnostics technology provides a lighted path to failed or failing components to expedite hardware repairs and reduce service time.

  - Three-level diagnostics allow you to more easily pinpoint the component and its location starting with the innovative drop-down light path panel.

- **Integrated Remote Supervisor Adapter**
  - Provides around-the-clock remote management capabilities standard that other companies require you to purchase optional.
  - Increase server availability by continuously monitoring your system and notifying you of potential system failures before they occur.
- **Copper Diagnostics**
  - Autonomic computing capabilities added to the x445 to improve ease of use and eliminate single points of failure.
  - Built into the server's BIOS requiring no operating system support.
  - Auto-detection: the ability to detect and report a correctly cabled multi-chassis configuration prior to boot for optimized performance.
  - Cable failover & redundancy: eliminates single point of failure by supporting failover of a single cable to any remaining cable in the event of a cable pull or cable failure. System transitions all traffic to the remaining cable without requiring a re-boot.
- **IBM Director**
  - Provides comprehensive systems management including remote supervision, alerting, and inventory with heterogeneous support across all xSeries servers.
  - New tools for partitioning including the Scalable Systems Manager for provisioning and managing multiple x445 partitions and Application Workload Manager for server consolidation of multiple applications on a single system image

2.2.2 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 2-1 for a comparison of the x445 with x440.

Restriction: x440 server cannot be upgraded to x445. An x440 cannot be connected together with an x445 to form a 16-way system.

Table 2-1 Major differences between x445 and x440

Feature	x440 server	x445 server
IBM chipset	First-generation XA-32 chipset	Second-generation XA-32 chipset
SMP scalability	Support configurations: ► Two-way, four-way, and eight-way using one x440 ► 16-way SMP using two x440 chassis	Support configurations: ► Two-way, four-way, and eight-way using one x445 ► 16-way SMP using two x445 chassis ► Future 32-way support using four x445 chassis
Intel processors	Supported processors: ► Xeon MP 1.5 GHz, 1 MB L3 ► Xeon MP 1.9 GHz, 1 MB L3 ► Xeon MP 2.0 GHz, 2 MB L3 ► Xeon DP 2.4 GHz, 512 MB L2 (up to four-way SMP only)	Supported processors: ► Xeon MP 2.0 GHz, 1 MB L3 ► Xeon MP 2.5 GHz, 1 MB L3 ► Xeon MP 2.8 GHz, 2 MB L3 ► Xeon DP 3.0 GHz, 512 MB L2 (up to four-way SMP only)
XceL4 Server Accelerator Cache	► 32 MB XceL4 per four-way SMP Expansion Module ► 64 MB XceL4 cache in eight-way configuration ► 128 MB XceL4 cache in 16-way configuration	► 64 MB XceL4 per SMP Expansion Module ► 128 MB XceL4 cache in eight-way configuration ► 256 MB XceL4 cache in 16-way configuration ► 512 MB XceL4 cache in 32-way configuration
Memory	► 512 MB and 1 GB DIMMs supported; 2 GB DIMMs by special bid only ► 16 GB RAM max per SMP Expansion Module ► Up to 32 GB RAM (64 GB in a 16-way)	► 512 MB, 1 GB, and 2 GB DIMMs supported ► 32 GB max per SMP Expansion Module ► Up to 64 GB (128 GB in a 16-way, 256 GB in a 32-way)
Type of memory used	► PC133 SDRAM memory ► Four-way interleaving	► 200 MHz DDR Memory ► Two-way interleaving
Memory features	Active Memory includes: ► Chipkill memory ► Memory mirroring ► Memory ProteXion (redundant bit steering)	Active Memory includes: ► Hot-add memory ► Hot-swap memory ► Chipkill memory ► Memory mirroring ► Memory ProteXion (redundant bit steering)

Feature	x440 server	x445 server
Internal hard drive capacity	► Two Ultra160 hard drive bays ► Drives up to 73 GB supported (146 GB total)	► Two Ultra320 hard drive bays ► Drives up to 146 GB supported (292 GB total)
SCSI/RAID	► Adaptec Ultra160 dual channel SCSI ► Support for an optional ServeRAID™ adapter	► LSI 1030 dual channel Ultra320 SCSI ► Integrated RAID-1 (drive mirroring) for internal drive bays ► Support for an optional ServeRAID adapter
Gigabit Ethernet	► Single port ► Broadcom BCM5700 10/100/1000 BASE-T	► Dual port ► Broadcom BCM5704 10/100/1000 BASE-T
PCI-X capacity	Six Active PCI-X (hot-swap) adapter slots, all 64-bit ► Two support adapters running at 33 or 66 MHz ► Two support adapters running at 33, 66, or 100 MHz, or one adapter running at 133 MHz and the other slot vacant ► Two support adapters running at 33, 66, 100, or 133 MHz	Six Active PCI-X (hot-swap) adapter slots, all 64-bit ► Two support adapters running at 33 or 66 MHz ► Two support adapters running at 33, 66, or 100 MHz, or one adapter running at 133 MHz and the other slot vacant ► Two support adapters running at 33, 66, 100, or 133 MHz
RXE-100 support	Adds six or 12 PCI-X slots to the server: ► Supported attachment to one x440 only ► Sharing an RXE among two x440 servers not supported	Adds six or 12 PCI-X slots to the server: ► Supported sharing between two x445s ► Can be shared between 2x four-way, 2x eight-way, and 2x 16-way x445 servers ► Maximum one RXE-100 for each 2-16-way configuration, or two RXE-100s for each 32-way configuration
Video controller	S3 Savage 4, 8 MB	ATI Rage XL, 8 MB
CD-ROM	24x CD-ROM	8x DVD-ROM
Serial port	► No stand-alone serial port included ► Can use the serial port of the Remote Supervisor Adapter	► Includes a cable to route a serial port to the rear of the server, occupying a PCI slot ► Can also use the serial port of the Remote Supervisor Adapter
Media Bay	24x-10x CD-ROM	24x-8x DVD

Feature	x440 server	x445 server
Service processor	Remote Supervisor Adapter	Remote Supervisor Adapter Future: Remote Supervisor Adapter II-EXA
Power supply	1050 W supply, hot-swap, full redundancy, 10 mm power distribution card	1050 W supply, hot-swap, full redundancy, 12 mm power distribution card (improved distribution)
Mechanical	depth 27.5”	depth 27.5” (not 28.1” as listed in the announcement letter)
Warranty	1 year, 24x7, same business day response, parts & labor	3 year, 9x5, next business day response, parts & labor

2.3 The benefits of system partitioning

System partitioning is virtualization of system resources, including processor, memory, I/O, and storage, so that all concurrent users appear to have access to the system, although each user is actually segmented and protected from the actions of other users. If one virtual partition freezes up, it would not affect the others.

System partitioning offers the ability to divide a system so that it can simultaneously support multiple operating system images. Among the benefits of system partitioning are:

- Server hardware consolidation
- High availability
- Software migration and coexistence
- Version control
- Development
- Testing and maintenance
- Better protection from viruses and software crashes
- Workload isolation
- Independent backup and recovery on a partition basis

System resources, including processor, memory, I/O and storage are virtualized so that all concurrent programs appear to have complete access to the system. If one virtual partition were to lock up, it would not affect the others.

Here are just a few of the ways that system partitioning can help you to improve IT efficiency:

- Server hardware consolidation — Consolidate many underused, underpowered, and unnecessary servers into a few productive ones. Reduce the number of current servers and buy fewer servers in the future.

- Increased server utilization — Divide a processor into multiple partitions rather than wasting an entire processor on one low-throughput application.
- Simplified server management — Manage fewer servers centrally versus many of them individually in multiple locations. Have fewer servers, cables, operating systems, and applications to deal with.
- Low-cost clustering/failover — Create clusters of partitions among hardware nodes. Have several different servers fail over to multiple partitions in one server.
- Simplified application deployment — Once you have tested and qualified a specific hardware platform for use with a particular operating system and application combination, you can deploy software images on multiple partitions, rather than having to requalify the software on another hardware platform.

Two types of system partitioning are:

- Physical partitioning

  With physical partitioning, a single server consisting of two nodes, such as the x445, can run multiple instances of an operating system in separate partitions. It can also run multiple versions of an operating system or even different types of operating systems.

  This means that a server can continue to run an operating system in one node while you install and test another version of that operating system, or a different operating system entirely in another node on that server without having to take the entire server offline.

- Logical partitioning

  VMware ESX Server, running on the x445 brings mainframe-type virtualized dynamic partitioning capabilities to the industry-standard server market, enabling you to partition the physical hardware into virtual machines capable of acting as independent and secure servers. ESX Server is a virtualization layer that sits between the operating system-application stack and the physical hardware turning the x445 into a pool of logical computing resources that can then be dynamically allocated to any operating system or application at varying levels of granularity.

  This technology dramatically reduces the cost and complexity of delivering such enterprise applications as Web serving, application serving, database, and e-mail.

  In addition, because ESX Server can support multiple similar or dissimilar operating systems, it gives the administrator the opportunity and flexibility to optimize the application by operating system in secure, partitioned virtual machines. For example, while some applications run better on Windows, other applications are more suitable or are only available on Linux.

If you intend to consolidate servers, system partitioning offers many benefits:

- Multiple operating systems previously run on multiple servers could all be running simultaneously on one server in one location.
- System partitioning enables you to set up different cluster types. Clustering delivers high availability, because multiple servers can be connected together with one server backing up the other. In the event that one of the servers requires maintenance or service, the second server can support the users and workload while corrective action is performed and the offline server is brought back online.
- Scalable clusters provide customers with industry-leading scalability at a system level, as well as load balancing to maximize performance and the support received by users accessing the system.

2.4 Server consolidation

Server consolidation means combining the functions performed by many servers into a fewer number of servers to reduce cost, complexity, network traffic, and management overhead, and to increase the efficiency of systems management, security, and resource utilization.

Server consolidation is complex, and needs a methodical approach because of the nature of the problem:

- Large numbers of servers are involved.
- Servers from different vendors, of different sizes, with different configurations.
- Software ranges from used and well-known to local and poorly understood.
- Business services being provided will vary greatly in volume and type.
- Consolidation may provide essential business functionality that must be protected from disruption.
- Consolidation must take place without delivering limits on an organization's future ability to adjust the size, scope, and direction of its business initiatives.

2.4.1 Types of server consolidation

One of the most important things to remember is that there are no "off-the-shelf" solutions for server consolidation. Every organization requires a unique solution that will match its unique infrastructure and business model.

There are four general types of server consolidation, offering a wide range of business value through varying degrees of solution complexity and investment.

There are four types of server consolidation:

- Centralization
- Physical consolidation
- Data integration
- Application integration

These are summarized in Table 2-2 and described in detail below.

Table 2-2 Server consolidation strategies

Type of Consolidation	Definition	Potential Benefit
Centralization	Relocate to fewer sites	Reduction in administration costs Increased reliability and availability Lower operation costs Improved security and management
Physical Consolidation	Replace with fewer, larger servers	Reduced hardware and software costs Improved processor utilization Reduced facilities costs (space, power, A/C) Lower operations costs Improved manageability
Data Integration	Combine data from multiple sources into a single repository	Reduced storage management costs Improved resource utilization Reduction in administration costs Improved backup/recovery capabilities Enhanced data access and integrity
Application Integration	Consolidation of multiple applications onto one server platform	Reduction in administration costs Increased reliability and availability Reduced facilities costs (space, power, A/C) Lower operation costs Scalability

- Centralization

  Server consolidation means different things to different people. As shown in Figure 2-4, in its simplest form, servers are physically moved to a common location. Because this simplifies access for the IT staff, it helps reduce operations support costs, improve security, and ensure uniform systems management. This is an important predecessor to future consolidation activities.

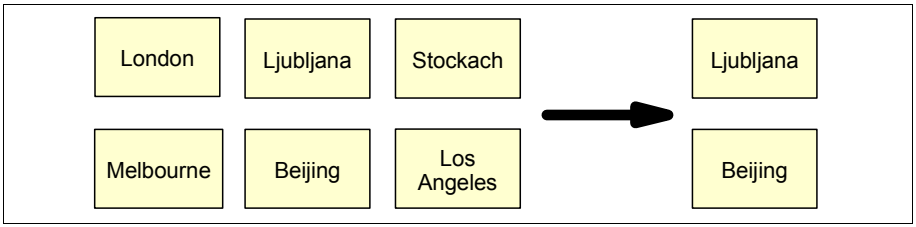

Figure 2-4 Centralization

Centralization involves relocating existing servers to fewer sites, for example, taking 20 servers scattered over three floors in your building and moving them to a single server room, or moving 200 servers originally installed across 20 locations to three data centers.

– **Relocating existing servers to one or fewer IT sites**

  Centralization, or data center consolidation, may be a first step for an organization after a merger. After a merger, the resulting entity does not want to attempt merging applications; however, they will collocate their systems as a first step.

  For both servers and storage systems, two subcategories of centralization are defined:

  - Virtual centralization, which is mainly made through the network
  - Physical centralization, where hardware is physically moved to different locations

  Centralization is often the initial step a company takes toward controlling costs through consolidation. It's also generally the first step taken toward rationalizing the architecture after a merger or acquisition.

  By simply relocating existing servers to fewer numbers of IT sites, economies of scale of operation can provide simplified management and cost improvement.

– **Virtual centralization or remote management**

  You can begin centralization in small steps. With virtual centralization or remote management, physically dispersed servers or storage systems are logically centralized and controlled through the network. Hardware remains physically distributed, but is brought under a common umbrella of systems management and network management tools. Operations costs can therefore be reduced, and system availability can be improved.

– **Physical centralization or server relocation**

  Existing servers or storage systems are physically relocated to one or fewer IT sites. Because this simplifies access for the IT staff, it helps reduce operations support costs, improves security, and ensures uniform

systems management. This is a step in the right direction, but the payback is relatively low. However, it is an important predecessor to future consolidation activities.

- Physical consolidation

  Physical consolidation is the replacement or reduction of some number of smaller systems with fewer and more powerful systems. This consolidation does have advantages:

  – It improves availability because there are fewer points of failure.
  – It can reduce the cost and complexity of system communications.
  – It simplifies operations.

  With its Enterprise X-Architecture enabled features, the x445 server offers flexibility, availability, and scalability to handle customer requirements for consolidating distributed workloads onto a single powerful and highly available platform to achieve total cost of ownership (TCO) savings.

  – **Reducing the number of servers by replacing many small servers with fewer large servers**

    Physical consolidation may be implemented on a site, department, or enterprise basis. For example, many x220 file/print servers can be consolidated onto newer, much faster, more reliable x445 servers, or older servers with high hardware maintenance costs can be consolidated or replaced by newer, much faster, cheaper-to-maintain x445 servers.

  – **Physical server consolidation**

    The number of separate hardware platforms and operating system instances within a consolidation site may vary considerably by customer. Typically, some reduction in the number of distinct servers is accomplished when gathering distributed systems into a central installation or when a number of small servers are replaced with larger servers of the same platform. Based on the enterprise's platform, four physical server consolidation cases can be considered.

    - Case 1: Small servers from one platform to server(s) on the same platform
    - Case 2: Small servers from different platforms to servers on different platforms (platform source and target are the same)
    - Case 3: Small servers from one platform to server(s) on a different platform
    - Case 4: Small servers from different platforms to server(s) on a different platform (platforms' source and target are not the same)

    Cases 1 and 2 are physical server consolidation, and there is no logical work to do. For cases 3 and 4, a platform migration has to be planned, and

applications and data have to be ported from one platform to another. The objective of the physical server consolidation phase is not to share applications or data but to have an application that was running on one platform run on a new platform. Therefore, this operation has to be differentiated from application or data integration.

Physical consolidation can be divided into two subcategories, namely server consolidation and storage consolidation.

This can take place within the same architecture -- for example, several two-way servers replaced with one 16-way x445 server or many uniprocessor servers moving to several multiprocessor x445 servers.

With x445 capabilities such as system partitioning, you can migrate and consolidate workloads across systems for improvements in systems management and resource utilization.

This approach is typically appropriate for implementations of key packaged applications such as SAP, PeopleSoft, and Siebel, where minimal integration with other applications and data is required. LAN file/print servers using Windows 2000 or Novell NetWare solutions represent another opportunity area for consolidation activities and savings.

Storage consolidation is combining data from different sources (same or disparate types) into a single repository and format. This means that storage is viewed as an enterprise resource, where centralized disk space is used to supply the storage for the servers of the enterprise.

Additional benefits can be gained through data integration and application integration. While these are often more complex projects that require extensive analysis, planning and implementation, they can provide significant return-on-investment.

- Data integration

  Data Integration involves physically combining data from different sources across the enterprise into a single repository and format. The result is that the merged data can reside on fewer servers and more centralized and consistent storage devices, greatly lowering the total costs.

  When all corporate data resides on the same system, consolidation allows high levels of security and data integrity that are nearly impossible to achieve in a distributed environment. In addition, data sharing throughout the enterprise is vastly simplified.

  The data can be file data such as Windows 2000, Novell, or Linux consolidated to a single network operating system. Also, multiple types of databases, such as DB2, Informix®, Oracle, Sybase, etc., can be converged to fewer database architectures.

In many client/server infrastructures, centralizing LAN data can bring dramatic improvements in data transfer speed. New enhancements in communications hardware will expand the high-speed connectivity options to server platforms of all types.

There can be two kinds of data integration:

- Data integration from several servers and consolidated on a unique repository
- Data integration from several repositories in one server and consolidated on a unique repository

Depending on the type of application integration selected, data integration can be performed separately or together with application integration.

- Application integration

  Application integration is the combining of multiple, similar applications, such as Web servers, onto one consolidated server.

  Application integration is also the combining of different application workload types within a single server/system and migrating an application or data to a new platform in order to collocate the application and data.

  It reduces administration, operation, and facilities costs and increases reliability and availability.

  The main objective of application integration is to migrate applications from one or several locations to a single location. Based on the consolidation platform, this migration can take different forms:

  - The migration may not bring any additional costs beyond that of relocating the application on a new server.
  - The migration may imply that application programs have to be recompiled in order to run on the new platform.
  - The migration may imply that application programs have to be redesigned and rewritten in order to run on the consolidation platform. As for physical server consolidation, application integration has several cases.
  - Application integration is combining different application workload types within a single server or system.
  - Distributed systems do not run identical applications and system software and have to be integrated into a consolidation server running a different operating system.

From another point of view, consolidation takes one of three basic approaches:

- Logical

  Logical consolidation brings all server resources to the same level so that they can be viewed logically as a single unified environment.

  In logical consolidation, actual systems are still distributed, while administrative procedures and processes are standardized company-wide. This kind of consolidation is relatively easy and safe to implement, but it carries the least potential for significant returns. Cost savings come from better asset management and opportunities to deploy high-quality, consistent administrative practices across the enterprise.

- Physical

  Physical consolidation does pretty much what it says: systems are relocated to a single server site. The number of servers you have to manage remains the same, and cost savings come from better staff utilization, higher service levels, simplified backups and restores, and better asset management and security.

- Rational

  In combination, or rational, consolidation, the company's distributed applications and services are combined onto fewer servers. It is a considerably more complex undertaking, but the potential rewards are greater. Cost savings range from 25-75% here resulting from better asset utilization and elimination of unnecessary systems, reduced staffing, lower maintenance costs, and fewer operating environments to support.

2.4.2 Why consolidate servers?

IT managers are feeling pressure to reduce costs, maintain or improve service levels, and maintain or improve the availability of systems that become ever more critical to daily operations.

Users want new applications that are delayed or inadequate because of IT infrastructure. IT needs to provide a cost-effective and reliable service, which is made difficult by constantly changing applications.

Many organizations are realizing that, as the number of servers increases, the cost and operational complexity are also propagating. In many cases, there are concerns whether multiple distributed servers can provide the application availability, hours of service, responsiveness, and ability to grow with the requirements of the business. These characteristics are being increasingly demanded by business applications. To reduce these costs, many customers are attempting to consolidate their servers into a more manageable central location.

The main objectives of server consolidation are:

- Recentralizing servers
- Merging workloads onto a single large server
- Consolidate architecture
- Optimize the IT infrastructure

2.4.3 Benefits from server consolidation

The main benefits of server consolidation are:

- **Single point of control**

  Rapidly growing firms, especially those growing through mergers and acquisitions, frequently felt that disparate distributed systems were so unwieldy to manage that they were losing control, which could constrain further corporate growth.

  A single point of control allows enterprises to:

  - Reduce or eliminate department operational costs
  - Reduce some software licenses
  - Reduce number of systems, disk storage costs
  - Reduce maintenance charges
  - Avoid multiple copies of the same application on distributed systems
  - Reduce owner operational costs
  - Offer better availability of service
  - Improve systems management
  - Have better version control management
  - Have better software distribution
  - Reduce risk and increase security

- **Giving users better services**

  With a consolidated infrastructure, end users can count on round-the-clock service, seven days a week. The response time is much better than with an overly distributed environment, and the data is more easily accessible while being highly protected. The control procedures are simpler, while security becomes even higher. And information sharing is improved, giving end users increased data consistency. The availability of service is improved mainly due to a reduction in the time needed to communicate between clients and servers in a single location.

- **Regaining flexibility**

  The standardization of procedures, releases, and servers also makes it easier to install new application software, for example, Internet and intranet, electronic commerce, and so on. In today's fast moving environment, computing resource consolidation enables a trouble-free upgrade of the information system and less costly adaptation to organization or environment

changes. Enterprises can react more quickly to market changes, since storage is readily available and can easily be reallocated.

- **Avoid floor space constraints**

  While a small server may be easily fit into a closet, as compute demands increase, enterprises find that suitable floor space is hard to find for proliferating small servers. The solution is a central site outfitted with appropriate power, cooling, access to communications links, and so on, and populated with more powerful systems, each giving more performance in the same footprint.

- **Reduction of the Total Cost of Ownership (TCO)**

  There are several costs associated with server consolidation, including:

  - Hardware costs — new servers and infrastructure, upgrades
  - Software costs — fewer software licenses are required with fewer servers
  - Disruption costs — migration, change management

- **Manageability and availability**

  Server consolidation can help you improve manageability and availability of IT systems in the following ways:

  - Enterprise management - Integrated operations allows for consistent management of all facilities and IT services.
  - Consistent performance - Providing consistent response time at peak load periods is very important.
  - Dependability - Commonly cited problems of distributed environments include frequency of outages and excessive requirements for manual intervention by the IT staff.

  In addition, it provides the following benefits:

  - It is easier to enforce consistent user policies in a consolidated environment.
  - Fewer servers lead to a simpler network structure that is easier to manage.
  - Reorganization following mergers or acquisitions is easier in a well-controlled environment.
  - Consolidation encourages standardization of tools, processes, and technologies to provide a stable and consistent application platform.

  Server consolidation can help you improve data access and protection in the following ways:

  - Network technology - The growth of networking and network speeds is enabling the centralization of IT networks today and will continue and expand into the future.

- Fragmentation and duplication of data - This is a core issue in most organizations with large numbers of distributed servers.
- Physical security - Consolidation of servers in a central data center can restrict unwanted access and ensure a more secure environment.
- Integrity, local backup and recovery - Enterprises are concerned about the dangers of business disruption, customer lawsuits, and regulatory action in the event of severe data loss, and they need to implement effective disaster recovery procedures.

Server consolidation can help you leverage existing investments in the following ways:

- Expand existing servers - Add new capabilities to the existing installation rather than to deploy new dedicated servers.
- Optimization of capacity utilization - In order to manage performance and have a level of acceptable, consistent response times, enterprises typically run at 50-60% utilization. Excess or underutilized capacity on one server cannot be shared with workloads of other servers in a distributed environment. Consolidation can help the customer to use much of the excess capacity for more efficient utilization of resources.
- Optimization of skilled resources - Under the distributed alternative, systems management responsibilities are often only part-time, extra-duty assignments such that a critical skill level is rarely achieved. Furthermore, since other departments may employ disparate architectures and applications, there is little opportunity to benefit from the experiences of others.

- **Scalability and workload growth**

  Server consolidation can help you handle scalability and workload growth issues in the following ways:

  - True scalability - Server consolidation provides the ability to deal with peak usage without crashing or seriously degrading performance. It also provides an upgrade path without degradation in response, excessively complex forms of database partitioning, or other problems.
  - Granular upgrades - Server consolidation provides the ability to quickly grow the number of users, the number of applications, or the size of an application when needed, without major disruptions to the current production environment.

- **Service level**

  Most companies spend much of their IT budget for services. They need services for hardware, software, and infrastructure maintenance. Server

consolidation can help you to reduce the increasing service costs in the following ways:

- Delivery of a specified service level is costly if servers are uncontrolled.
- Management of servers as corporate assets is easier when they are centralized.
- Application deployment is quicker and easier when not spread over a large number of servers.
- Staff time freed from server maintenance tasks can be used for activities more directly related to business needs.

- **Business continuity**

  Almost all enterprises need to run their business without interruption. Business interruption can be very costly and it influences the productivity of your business. Server consolidation can help you to run your business without interruption in the following ways:

  - Consolidating IT resources can help you ensure that critical business information and processes are accessible and shared across the enterprise.
  - Implementing critical new solutions that may enable a competitive edge is easier.

3

Planning

In this chapter we discuss the topics you need to consider before you finalize the configuration of your x445 system and before you begin implementing it. The topics covered are:

3.1 System hardware

The x445 provides a scalable and flexible hardware platform. There are a number of important aspects of the system hardware to consider when planning your configuration. These are discussed in this section.

3.1.1 Processors

There are currently two processor types available with the x445 system:

- Xeon DP models come with either two Xeon DP 3.0 GHz 512 KB L2 cache "Prestonia" processors in a single SMP Expansion Module or with four processors in two SMP Expansion Modules. You cannot have more than four Xeon DP processors, however you can replace all of them with Xeon MP processors, and then you can install up to eight Xeon MP processors in the one x445.
- Xeon MP models come with two or four Xeon MP "Gallatin" processors installed in the standard SMP Expansion Module. Up to four Xeon MP processors are supported in the standard SMP Expansion Module. Using the optional second SMP Expansion Module, part number 02R1870, you can install up to eight processors in an x445.

The following processors are available as options:

- Xeon Processor MP 2.0 GHz 1 MB L3 cache, 02R2062
- Xeon Processor MP 2.5 GHz 1 MB L3 cache, 02R2063
- Xeon Processor MP 2.8 GHz 2 MB L3 cache, 02R2064

Note: The Xeon DP processors are not available separately. Instead, to upgrade a two-way Xeon DP-based x445 to four-way, use the populated SMP Expansion Module 02R1871.

Only the configurations with two, four, or eight processors are supported in an x445 chassis. Therefore, you should not use any other number of processors (1, 3, 5, 6 or 7).

Key configuration rules:

- All processors must be the same type, speed, and L2/L3 cache size.
- For servers with Xeon MP processors:
  - The standard SMP Expansion Module must have four processors installed before the second one can be installed and used.

  - Use part number 02R1870 for the second SMP Expansion Module. This part number is not populated with any processors or memory.
  - The second SMP Expansion Module is supported only when it is fully populated with four Xeon MP processors. Consequently, if you install it, the system must have eight processors after the installation.
  - All Xeon MP processors must be identical for 16-way configurations.
- For servers with Xeon DP processors:
  - Each SMP Expansion Module must have two processors installed and these processors must be installed in positions 1 and 4 (see Figure 1-10 on page 18).
  - Use the part number 02R1871 for the second SMP Expansion Module. This part number includes two 3.0 GHz Xeon DP processors.
  - You can upgrade a Xeon DP model to have Xeon MP processors, but all Xeon DP processors must be removed. You cannot mix Xeon MP and Xeon DP processors in the same x445 system.

The Xeon MP and Xeon DP SMP Expansion Modules are identical. Therefore, if one fails, you could take a module from another system (and therefore downgrade it to a single SMP Expansion Module unit) and install it in the failed system. Ensure the CPU and memory configuration rules are still maintained.

16-way configurations

You can use a pair of 8-way x445 systems to create a 16-way configuration. Any of the base x445 models can be used, however you have to upgrade them to eight identical processors. The following operating systems support currently support 16-way configurations:

- Microsoft Windows 2000 Datacenter Server
- Windows Server 2003 Datacenter Edition
- VMware ESX Server 2.0

In addition, if you plan to run Windows Datacenter Server, there are also Datacenter-specific models, as shown in Table 1-2 on page 5.

The 16-way configuration requires 16 identical processors and an xSeries Two Chassis 445 16-way Configuration Kit, part number 02R2013. This kit contains four 2.5m copper-colored scalability cables and an Ethernet crossover cable to connect the Remote Supervisor Adapters together. However, if a Remote Supervisor Adapter connection to the external LAN is required (for use with out-of-band management), you will have to use an Ethernet switch or a hub. In this case we recommend to have your management LAN separate from your data LAN.

Notes for 16-way configurations:

- VMware ESX Server 2.0 is supported on the x445. The operating system is not included with the server and must be ordered separately. See 3.6.7, “VMware ESX 2.0 Server” on page 101 for more information.
- VMware ESX Server 1.5.x is not supported.
- Windows 2000 Datacenter Server and Windows Server 2003 Datacenter Edition are supported on the standard models as well as specific Datacenter models. The operating system is not included with the server and must be ordered separately. Datacenter can also be preloaded by IBM at the factory. See 3.6.1, “Windows Datacenter” on page 95 for more information.
- If you use any USB devices, you have to connect them to the primary chassis (which you specify in BIOS when setting up the 16-way configuration). These devices are not supported on the secondary node.

Hyper-Threading

As described in 1.4.1, “Intel Xeon Processor MP” on page 12, the processors used in the x445 include the Hyper-Threading technology. Hyper-Threading makes a single CPU appear to the operating system as two logical processors, which can receive and process two data/instruction streams simultaneously.

Hyper-Threading is disabled by default on the x445. Fewer or slower processors usually yield the best gains from Hyper-Threading, because there is a greater likelihood that the software can spawn sufficient numbers of threads to keep both paths busy.

Hyper-Threading can be enabled in BIOS if needed. Customers should test their specific application sets with Hyper-Threading turned on and turned off to gauge which environment provides the best performance for their specific needs.

3.1.2 Memory

The x445 standard models come with either 0 or 2 GB (four 512MB modules) of memory. You can expand the memory size up to 32 GB with one SMP module, or up to 64 GB with two SMP modules.

Each SMP Expansion Module contains 16 DIMM sockets divided across two ports and four banks within each port. A bank is a pair of DIMMs required for two-way interleaving.

- Port 1:
  - Bank 1 = DIMM connectors 1 and 3
  - Bank 2 = DIMM connectors 2 and 4
  - Bank 3 = DIMM connectors 5 and 7

  - Bank 4 = DIMM connectors 6 and 8

- Port 2:
  - Bank 1 = DIMM connectors 9 and 11
  - Bank 2 = DIMM connectors 10 and 12
  - Bank 3 = DIMM connectors 13 and 15
  - Bank 4 = DIMM connectors 14 and 16

Physically, the banks occupy alternating sockets, as shown in Figure 3-1 on page 61.

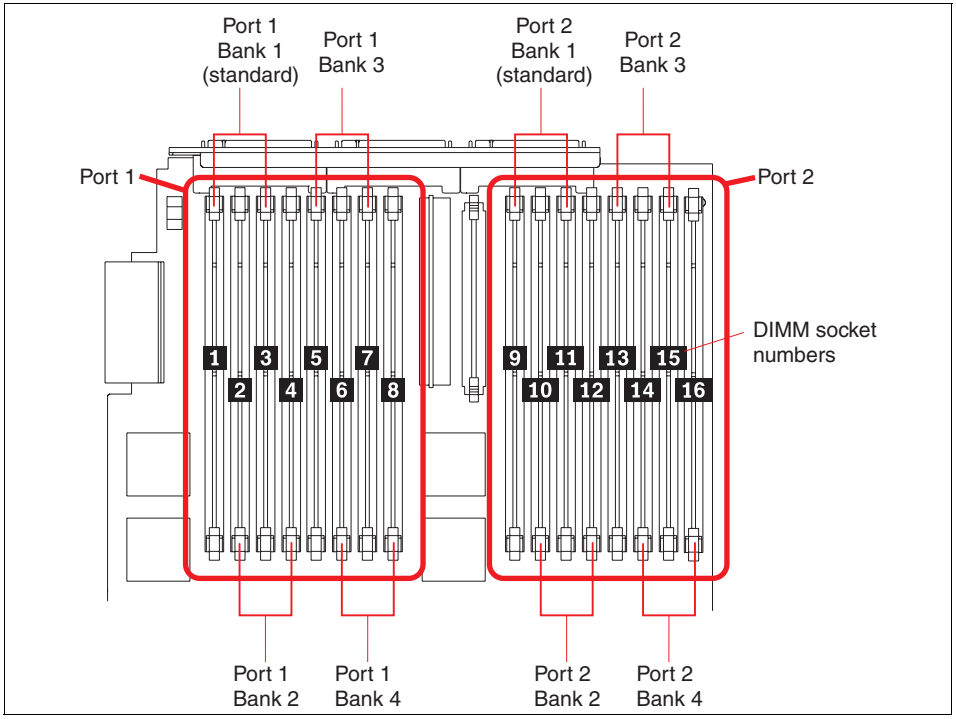

Figure 3-1 DIMM sockets on the x445 SMP Expansion Module

Memory configuration rules:

- Because the x445 uses two-way memory interleaving, memory modules must be installed in pairs, so that a particular bank is either fully populated with two DIMM modules, or completely empty. It will not work if you only install a single DIMM into one of the slots. The following sizes are supported:
  - 512 MB DIMMs (part number 33L5038)
  - 1GB DIMMs (part number 33L5039)
  - 2GB DIMM (part number 33L5040)

- You can use different capacity DIMMs within the same SMP Expansion Module, but both DIMMs in each bank must be the same size.
- If you want to install more than 32 GB of memory, you must install two SMP Expansion Modules and sufficient processors (eight MPs or four DPs). This is because more than 32 GB of memory requires the second SMP Expansion Module, which in turn must be fully populated with four Xeon MPs or two Xeon DPs.
- The models with standard 2 GB of memory contain four 512 MB DIMMs. These DIMMs are installed in port 1, bank 1 and port 2, bank 1 as shown in Figure 3-1 on page 61 (marked as “standard”).
- If you want to install the full 32 GB in an SMP Expansion Module, you will need to remove the existing 512 MB DIMMs and fully populate the module with 2 GB DIMMs.
- For performance reasons, you should install DIMMs in the bank order shown in Table 3-1.

Table 3-1 Recommended DIMM installation order

Sequence	Port	Bank	Slot numbers
1	1	1	1, 3
2	2	1	9, 11
3	1	2	2, 4
4	2	2	10, 12
5	1	3	5, 7
6	2	3	13, 15
7	1	4	6, 8
8	2	4	14, 16

Active memory considerations

The x445 supports two new features: hot-swap and hot-add memory. These technologies are mutually exclusive: you use a setting in the BIOS to specify which one you want to enable. This requires a shutdown, so you need to plan carefully and in advance in order to avoid unnecessary system downtime. The following sections describe these features.

Hot-add memory

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 Edition and Datacenter Edition.

The prerequisites for hot-add are as follows:

- Memory mirroring must be disabled in the Configuration/Setup utility in the BIOS.
- Either port 1 or port 2 must be completely empty. The empty port is where you can add DIMMs. The x445 models have four 512 MB DIMMs standard, two DIMMs in each port. Consequently, to enable hot-add, you must move two DIMMs so that all four are in one port.

  This prerequisite is because the power to the port you plan to add DIMMs to must be off before you add the DIMMs. Once any DIMMs are added and the door closed, power is applied.
- If the server has two SMP Expansion Modules installed, you can only hot-add memory into the top one, because the bottom module memory slots are physically not accessible. You also have to be aware that the bottom SMP Expansion Module must have memory installed (at least two DIMMs) for the server to boot.

The rules for hot-add are as follows:

- You can hot-add two, four, six or eight DIMMs. Once you add DIMMs and close the memory door, you cannot add more memory without powering off the server.
- Both DIMMs in a pair hot-added to a certain bank (see Table 3-1 on page 62) must be the same size.

Hot-swap memory

The hot-swap memory feature allows you to replace failed DIMMs without turning off the server. Hot-swap memory is operating system independent.

The rules for hot-swap are as follows:

- Memory mirroring must be enabled in the Configuration/Setup utility in the BIOS.
- The configuration rules for memory mirroring apply here, specifically that the DIMMs in partner banks across the two ports must be identical. See “Memory mirroring” on page 64 for details.
- If the server has two SMP Expansion Modules installed, you can only hot-swap DIMMs in the top SMP Expansion Module, as the bottom one is physically not accessible.

- If the server has two SMP Expansion Modules installed, you cannot move all memory to the top SMP Expansion Module. The bottom SMP Expansion Module must have memory installed (at least two DIMMs) for the server to boot.
- When replacing a failed DIMM, you do not have to replace the other DIMM within the bank.
- Ensure the new DIMM is the same type, size and clock speed as the other DIMM within the bank.

For optimum performance, balance the amount of memory between the two ports. See Table 3-1 on page 62 for the recommended DIMM installation order.

Memory mirroring

The x445 supports memory mirroring for higher system availability. The feature is roughly equivalent to RAID-1 in disk arrays: memory is accessed across the two ports and one port is mirrored to the other (see Figure 3-1 on page 61). If 8 GB is installed for example, then the operating system sees 4 GB once memory mirroring is enabled (it is disabled in the BIOS by default).

Memory mirroring is operating system independent, since all mirroring activities are handled by the hardware.

When memory mirroring is enabled, the data is written to memory modules in both ports. One copy is kept in the port 1 DIMMs, while a second copy is kept in the matching port 2 DIMMs. During the execution of the read command, the data is read from the DIMM with the least amount of reported memory errors through memory scrubbing, which is enabled with memory mirroring.

If memory scrubbing determines the DIMM is damaged beyond use, read and write operations are redirected to the remaining good DIMM. Memory scrubbing then reports the damaged DIMM and the light path diagnostics display the error. If memory mirroring is enabled, the mirrored copy of the data from the damaged DIMM is used until the system is powered down and the DIMM replaced. If hot-swap is enabled, the DIMM can be replaced immediately without a system power down. After the damaged DIMM is replaced, memory mirroring then copies the mirrored data into the new DIMM.

Key configuration rules relating to memory mirroring:

- Memory mirroring must be enabled in the BIOS (it is disabled by default).
- Both ports in an SMP Expansion Module must have the same total amount of memory. Partial mirroring is not supported.
- When using memory mirroring, the DIMMs in partner banks must be the same size and clock speed. For example, if the DIMMs in port 1, bank 1 are

512 MB, then the DIMMs in port 2, bank 1 (the partner bank) must also be 512 MB. See Table 3-1 on page 62 for the DIMM slots that make up each bank. All DIMMs do *not* have to be identical, just the DIMMs in partner banks.

- SMP Expansion Modules are individually configured for memory mirroring in the BIOS. This means that as well as full memory mirroring, you can also enable memory mirroring only in one SMP Expansion Module, although we do not recommend this.
- Memory mirroring does not work across the SMP Expansion Modules. For example, you cannot set up four 512 MB DIMMs in the bottom SMP Expansion Module to be mirrored by four 512 MB DIMMs in the top SMP Expansion Module.

Memory performance

From a performance perspective, you should attempt to balance memory between the SMP Expansion Modules. This is more important than maximizing memory bandwidth to a module. Make sure each module has the same amount of memory. Then, if possible, make sure each module has the same number of DIMMs installed.

Then, also attempt to balance memory between ports of the one SMP Expansion Module. For performance reasons, consider the following:

- When installing four DIMMs, install two in port 1, bank 1 (sockets 1 and 3) and two in port 2, bank 1 (sockets 9, 11). See Table 3-1 on page 62.
- When installing DIMMs, try to balance the amount of RAM between the two ports. For example, if you have 12 DIMMs (eight 512 MB DIMMs and four 1 GB DIMMs), install all eight 512 MB DIMMs (4 GB) in one port and four 1 GB DIMMs (also 4 GB) in the other port. This will give you better performance than mixing four 512 MB DIMMs and four 1 GB DIMMs (6 GB total) in one port and four 512 MB DIMMs (2 GB) in the other port.

3.1.3 Active PCI-X subsystem

As shown in Figure 3-2 on page 66, the x445 contains six internal PCI-X slots.

These six slots are implemented using four PCI buses, labeled A-D:

- Bus A (slots 1 and 2): Two 64-bit 66 MHz slots
- Bus B (slots 3 and 4): Two 64-bit 100 MHz slots (133 MHz if only one slot is occupied)
- Bus C (slot 5): One 64-bit 133MHz slot
- Bus D (slot 6): One 64-bit 133MHz slot

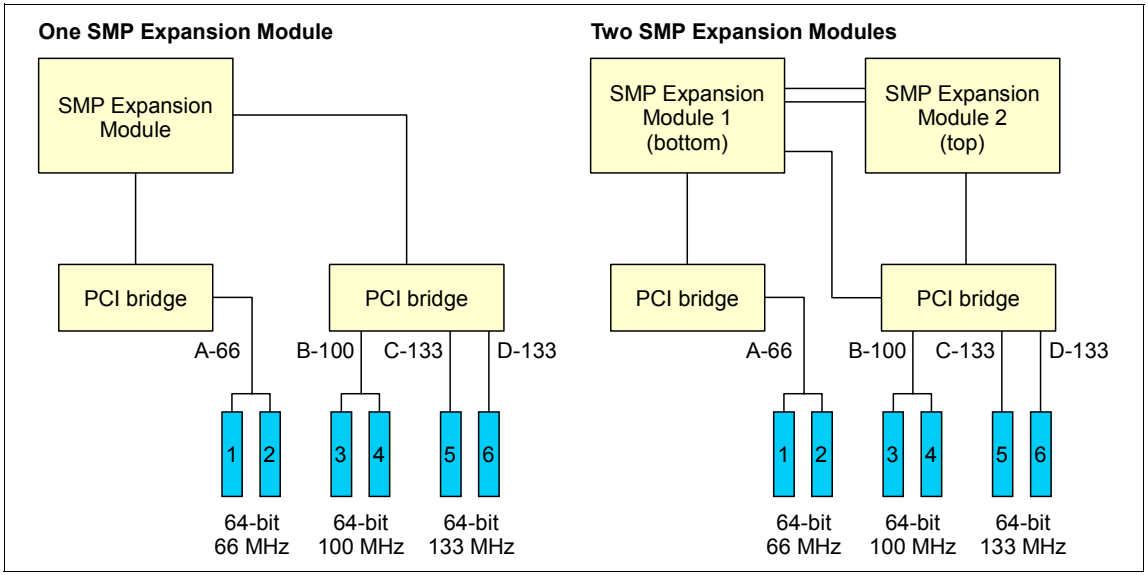

Figure 3-2 x445 PCI-X slots with one and two SMP Expansion Modules

Notice the connectivity of the PCI bridges when two SMP Expansion Modules are installed. PCI-X slots 1 and 2 are connected to the PCI bridge in the bottom SMP Expansion Module (module 1) — that is, they are “local” to that module — and slots 3-6 are connected to the top SMP Expansion Module. This has performance implications with NUMA-aware applications and operating systems, because the placement of adapters will determine which CPUs and memory are “local”.

In a 16-way configuration, all internal PCI slots are “local” to the bottom SMP Expansion Module (module 1) of each server. This is controlled by the server’s BIOS. To provide “local” PCI slots to the top SMP Expansion Modules in each server (for performance reasons), connect an RXE-100.

The x445 PCI-X slots can accept adapters rated at speeds ranging from 33 MHz to 133 MHz. When deciding how to populate the slots with the adapters, consider the following:

- Each adapter and each bus have a maximum rated speed.
- Adapters in a particular bus will operate at the slowest of three speeds:
  - The rated speed of adapter 1
  - The rated speed of adapter 2 (if the bus has two slots)
  - The rated speed of the bus
- Bus B supports one adapter at up to 133 MHz or two adapters at up to 100 MHz.

- 32-bit adapters can be installed in any of the slots and will run in 32-bit mode. 32-bit and 64-bit adapters can coexist in 64-bit slots in the same bus. The 32-bit adapters will run in 32-bit mode, and the 64-bit adapters will run in 64-bit mode.

Tip: Take the time to understand these rules and to select the best slots for your adapters. Incorrect choices can result in a reduction of PCI adapter performance. Consider using the Active PCI Manager, described in 6.1, “Active PCI Manager” on page 190.

As extreme configuration examples, you could configure either of the following:

- Six 33 MHz PCI adapters, all operating at 33 MHz.
- Six 133 MHz PCI-X adapters, with two operating at 133 MHz (buses C and D), two at 100 MHz (bus B), and two at 66 MHz (bus A).

Table 3-2 summarizes the supported adapter speeds. Take into account the speed reductions when two 133 MHz adapters are installed in the bus B.

Table 3-2 Supported adapter speeds

Slot	Bus	Width (bits)	
1	A	32 or 64	33 or 66
2	A	32 or 64	33 or 66
3	B	32 or 64	33, 66 or 100 (or 133 if slot 4 is empty)
4	B	32 or 64	33, 66 or 100 (or 133 if slot 3 is empty)
5	C	32 or 64	33, 66, 100 or 133
6	D	32 or 64	33, 66, 100 or 133

The physical location of these slots in the server is shown in Figure 3-3 on page 68.

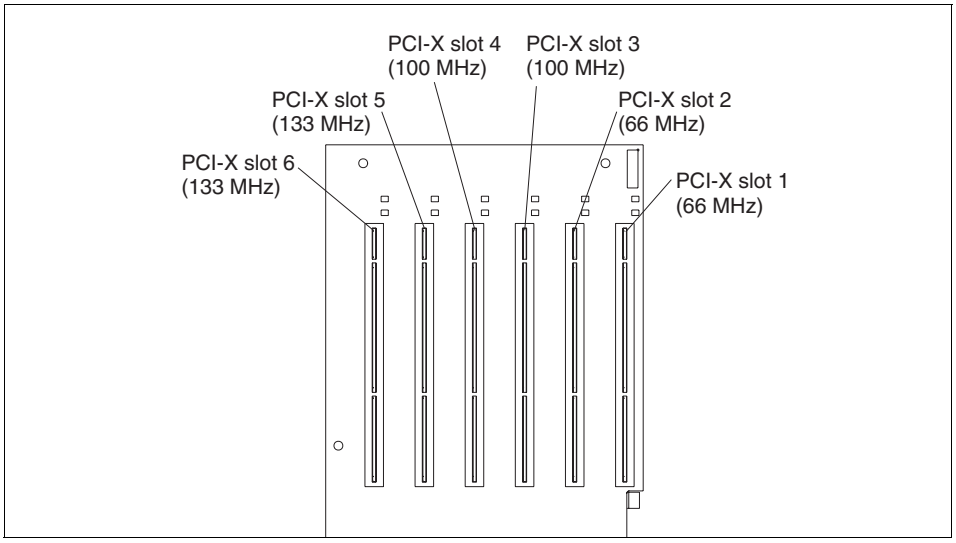

Figure 3-3 Physical location of slots

Other configuration information:

- Video adapters are not supported.
- The PCI slots support 3.3 V adapters only. 5 V adapters are not supported.
- The x445 ships with a serial cable that, when installed, occupies slot 1. See 3.3.4, “Serial connectivity” on page 84.
- The ServeRAID 4H adapter is not supported for internal drives because it is too high to fit in the 4U server when a cable is attached to its internal connector.
- The x445 comes with an additional pre-installed cable to enable the ServeRAID adapter to connect to the internal drives.

  Note: The pre-installed cable for the ServeRAID adapter is disconnected at *both* ends. To use it, disconnect the smaller SCSI cable from the hard drive backplane. Then connect the ServeRAID cable to the hard drive backplane and to the ServeRAID card itself.

- Some long adapters have extension handles or brackets installed. Before installing the adapter, you must remove the extension handle or bracket.
- The system scans PCI-X devices to assign system resources. It attempts to start the first device found. The search order is:

  a. DVD-ROM

b. Disk drives
c. Integrated SCSI controller
d. x445 PCI-X slots (in the order 1, 2, 3, 4, 5, 6)
e. Integrated dual Gigabit Ethernet controller

If an RXE-100 is attached, the order is:

a. DVD-ROM
b. Disk drives
c. Integrated SCSI controller
d. x445 PCI-X slots (1, 2, 3, 4, 5, 6)
e. RXE-100 slots (A1, A2, A3, A4, A5, A6, B1, B 2, B3, B4, B5, B6)
f. Integrated dual Gigabit Ethernet controller

PCI-X hot-add considerations

You have to be careful when hot-adding an adapter to a particular slot, if the other slot on the same bus is already populated. In such a case, the bus speed will already be set to a certain frequency, determined by the adapter installed. If this is higher than what the adapter being hot-added supports, the operation will not succeed. The speed on the bus cannot be decreased dynamically to suit the new adapter, and the new adapter will not be able to operate at speed which exceeds its capability. You should therefore plan in advance which adapters you might need to hot-add at a later time, and which slots they should go into.

This is also true for hot-adding a 133 MHz adapter into a slot 3 or 4, when the other slot on the same bus already contains another 133 MHz adapter. In such a case, the bus would operate at 133 MHz with one adapter, but the speed would need to be decreased to 100 MHz with two adapters. As this cannot be done dynamically, such hot-add operation will not work.

PCI-X hot-adding and hot-swapping is supported on the following operating systems:

- Windows 2000
- Windows Server 2003
- Red Hat Linux 2.1 AS
- SuSE Linux SLES 8.0
- NetWare 6

3.1.4 Integrated LSI SCSI controller

The x445 includes an integrated LSI Logic Ultra320 SCSI controller with two SCSI channels. One channel is attached to the internal two-drive hot-swap backplane, and the other to the external 0.8 mm VHDCI 68-pin SCSI connector at the back of the chassis.

The SCSI controller supports RAID-1 (mirroring) across the disk drives, and we suggest you to use it for protection of the operating system on the internal drives. You configure mirroring by accessing the LSI Logic Configuration Utility program. Press <Ctrl-C> at boot time, when prompted to do so, and then set the mirrored pair of drives up.

The SCSI controller supports hot spare drives as well. But since the x445 contains only two internal disk drive bays, you cannot use this option.

3.1.5 Integrated Broadcom dual-port Ethernet controller

The x445 includes a dual-port Broadcom BCM5704 10/100/1000 Ethernet controller on the PCI 64-bit 66 MHz bus. The BCM5704 has the following features:

- It supports full and half-duplex performance at all speeds (10/100/1000 Mbps, auto negotiated).
- It includes integrated on-chip memory for buffering data transmissions to ensure maximum network performance.
- It includes dual onboard RISC processors for advanced packet parsing.

The Broadcom controller includes software support for failover, load balancing, VLAN support and comprehensive diagnostics.

Category 5 or better Ethernet cabling is required with RJ-45 connectors. If you plan to implement a Gigabit Ethernet connection, ensure your network infrastructure is capable of the necessary throughput to match the server's I/O capacity.

Ethernet controller teaming considerations

The BASP (Broadcom Advanced Server Program) software provides load balancing, fault tolerance and VLAN features on Windows Server 2003, Windows 2000, Windows NT, Novell NetWare and Linux. The utility allows you to create teams of up to eight NIC interfaces. Each interface in the team can be set as primary or standby. Primary interfaces take part in the load balancing scenarios, while standby interfaces only become active when none of the primary interfaces can be used anymore.

Various Ethernet controllers can participate in a load-balanced, fault-tolerant team:

- Integrated BCM5704 controllers
- Broadcom NetXtreme Gigabit Ethernet adapters
- 3Com 10/100 server adapters
- Intel 10/100 server adapters

- Intel 1000Base-SX server adapters
- Intel PRO/1000 server adapters

3.2 Sizing information

Proper sizing of a solution is very important, and the most obvious reasons are as follows:

- Performance

  The performance of the servers must be adequate both with the typical workload and also during the peaks. This means we need to plan for the extra performance capacity above the usual workload. The components that affect performance are the CPUs, amount of installed RAM, network throughput, and the access to disk drives. These subsystems must operate in a balanced manner, so that none of them is significantly limiting performance of the others.

- Storage space

  The storage space sizing is extremely important: if a system runs out of the disk space, its operation will be severely impacted or even completely disabled. So it is essential that we plan for enough disk space to allow for the growth of the storage requirements.

- Expansion possibilities

  As it is quite possible that the capacity required at some point in the future will exceed even the carefully sized solution, it is very important to be able to expand the system. The x445 has very good expansion possibilities in terms of:

  – CPU (up to 16-way now, up to 32-way in the future)
  – Memory (up to 64 GB per 8-way system)
  – Storage (using the SAN solutions)
  – PCI-X slots (using the RXE-100 enclosure).

Be aware that proper sizing can only be achieved when it is based on quality and realistic input data. This usually requires implementation of real workload simulations, rather than relying on a simpler rules-of-thumb approach. Careful study of the workload will enable you to estimate the peak requirements correctly. A well-sized system will be able to withstand the peak, not just the average workload. Sizing based purely on averages of workload will most certainly cause slow performance at peaks.

It is also important to plan for adequate response time. This is especially meaningful on the servers running interactive applications, while servers running mostly batch jobs might not be that sensitive to this criteria. When users have to

wait too long for the application response, their productivity will decrease and the business process will suffer.

3.3 Cabling and connectivity

You need to consider a number of factors when cabling the x445 server:

- SMP Expansion Module connectivity
- Remote Supervisor Adapter connectivity
- RXE-100 connectivity
- Serial connectivity

We will discuss each of these in this section.

The rear panel of the x445 showing the locations of cable connectors is shown in Figure 3-4. For port locations on the Remote Supervisor Adapter, refer to Figure 3-8 on page 75.

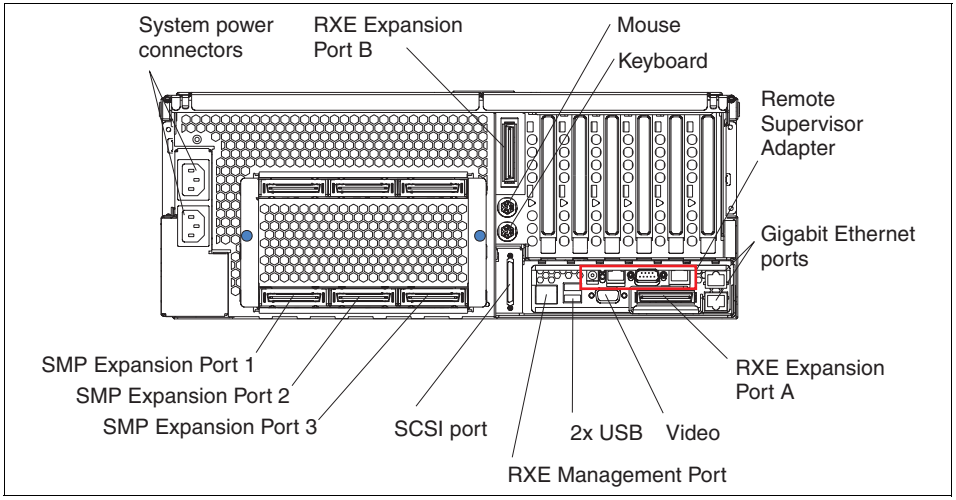

Figure 3-4 x445 rear view

3.3.1 SMP Expansion Modules connectivity

When only one SMP Expansion Module is installed in an x445, an air baffle, as shown in Figure 3-5 on page 73, must be installed instead of the second SMP Expansion Module to maintain proper air flow. If you disposed of the baffle by accident, you can order a replacement using the FRU part number 02R9132.

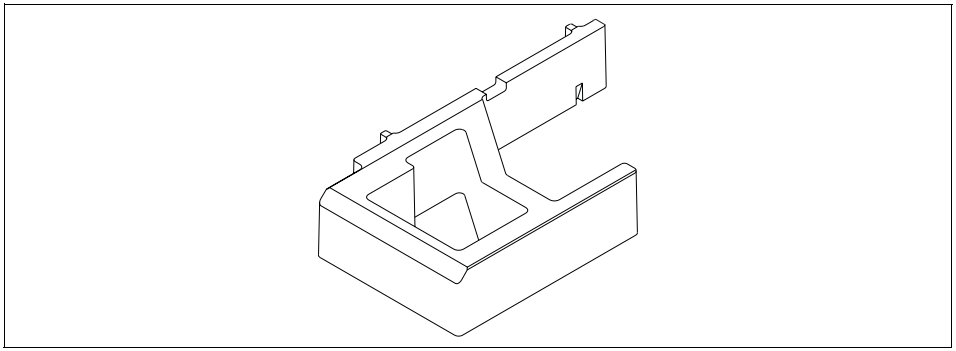

Figure 3-5 SMP air baffle

When a second SMP Expansion Module is installed in the x445, an additional three SMP Expansion Ports are made available on the rear of the chassis, giving a total of six.

With single-chassis x445 configurations, four of these ports are used to connect the two internal SMP Expansion Modules together, which allows for an 8-way configuration. One cable would suffice for basic operation, but we use two cables for higher availability and better performance. You need to connect the modules with two 25cm (10”) scalability cables as shown in Figure 3-6 on page 73:

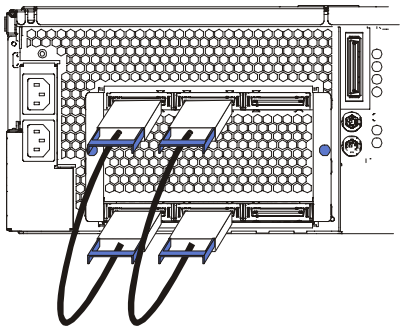

8-way connections:

- Port 1 on the upper module to port 1 on the lower.
- Port 2 on the upper module to port 2 on the lower.
- Port 3 is left unused in the 8-way system.

Figure 3-6 Connecting two SMP Expansion Modules within a single chassis

Cabling for a 16-way configuration

When forming a 16-way configuration, connect the scalability ports as shown in Figure 3-7 on page 74. Basically, you have to connect all the SMP Expansion Modules to each other. As you need to connect each SMP Expansion Module to the remaining three modules, you have to do it in the following manner:

- Use ports 1 and 2 for connections to the two modules in another chassis.

- Use port 3 to connect to the other module in the same chassis.

In addition, you have to connect the Ethernet ports of the Remote Supervisor Adapters together (either via a cross-over cable or via an Ethernet switch).

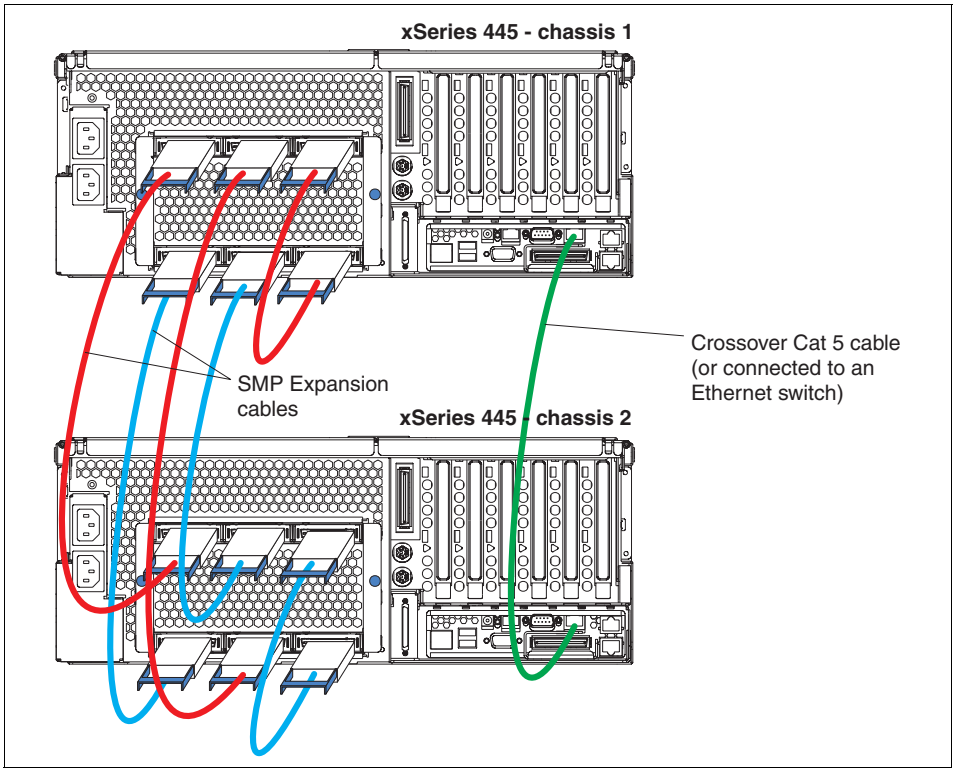

Figure 3-7 Connecting the two x445s together in a 16-way configuration

This cabling requires the installation of four additional 2.5 m copper-colored scalability cables to complete the configuration. These cables are supplied with the x445 Two Chassis 16-way Configuration Kit, part number 02R2013.

Key points relating to the SMP Expansion Module cabling:

- The cables are labelled to make it easier to correctly cable the two systems together.
- The SMP Expansion Module ports do not support use as high-speed interconnects for clustering purposes. Use the integrated Gigabit Ethernet controller instead.

- The connections do offer redundancy. If a connection is lost (for example, if a cable is pulled), the server will disable the failing port and switch traffic to the remaining path.
- In 16-way configurations, the Ethernet port on the Remote Supervisor Adapter in one system is connected to the Ethernet port of the Remote Supervisor Adapter in the other system. This connection is used during system startup and shutdown.

  We recommend either of the following connections:

  – Connecting the two using the provided crossover cable, as shown in Figure 3-7.

  – Connecting the two over an isolated LAN segment using a switch or hub.

    Connecting over a LAN segment will enable you to maintain Ethernet connectivity directly to the Remote Supervisor Adapter for out-of-band management.

Tip: We recommend that you assign static IP addresses to the Remote Supervisor Adapters on both servers.

3.3.2 Remote Supervisor Adapter connectivity

The x445 features an integrated Remote Supervisor Adapter. For detailed information on functionality and configuration 9.5 “Remote Supervisor Adapter” in the redbook *Implementing Systems Management Solutions using IBM Director*, SG24-6188. This document talks about the Remote Supervisor Adapter as a separate adapter; however the functionality and location of ports is consistent with the integrated version of the adapter in the x445.

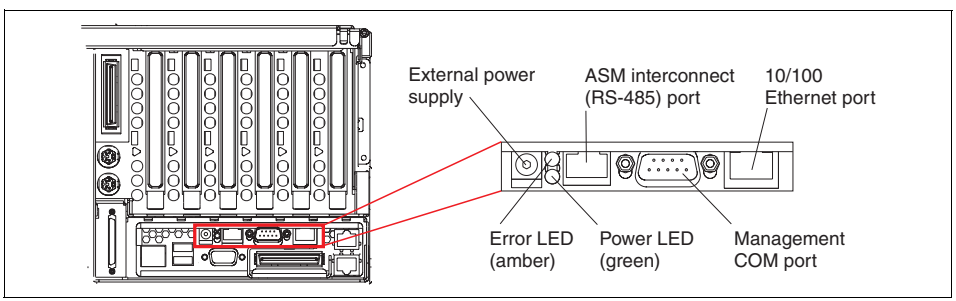

Figure 3-8 Remote Supervisor Adapter connectors

The following Remote Supervisor Adapter connections need to be considered when cabling the x445 (see Figure 3-8):

- External power supply connector. This connector allows the adapter to be connected to its own independent power source. This external power supply is not included with the x445 and will need to be ordered as an option (order a ThinkPad® 56 W AC Adapter with a suitable power cord for your country/region).

  Note: The connection to the Remote Supervisor Adapter of an external power source is not officially supported.

  If this power supply is not used, the adapter will draw power from the server as long as the server is connected to a functioning power source.

- 9-pin Serial port, which supports systems management functions through null modem or modem connections.

- Ethernet port, which provides system management functions over the LAN.

  As described in “Cabling for a 16-way configuration” on page 73, in two-chassis configurations (such as the 16-way), the Remote Supervisor Adapter Ethernet ports of the two servers must be connected together either using a crossover cable or via a 100 Mbps Ethernet switch. The use of a switch is recommended if you also want to perform out-of-band management activities.

- Advanced Systems Management (ASM) RS-485 Interconnect port to facilitate advanced systems management connections to other servers.

  For detailed instructions on cabling ASM interconnect networks, refer to section 9.11 “ASM Interconnect” in the redbook *Implementing Systems Management Solutions using IBM Director*, SG24-6188.

  Note: The x445 does not include the necessary dongle to connect the Remote Supervisor Adapter to an ASM interconnect bus using the RS-485 port on the adapter. Consequently, you will need the Advanced System Management Interconnect Cable Kit (part number 03K9309) for connection to an ASM interconnect network.

In November 2003, IBM announced an option, part number 13N0382, to allow existing x445 models to be upgraded to the Remote Supervisor Adapter II, replacing the existing Remote Supervisor Adapter. This option must be installed by an IBM service technician or authorized business partner.

3.3.3 RXE-100 connectivity

The x445 can be connected to a single RXE-100 to provide an additional six or 12 PCI-X slots to the server. Currently, only one RXE-100 is supported (with either six or 12 slots) per x445 server or per 16-way two-node configuration.

Note: Previously, Red Hat and SuSE Linux only supported the first six slots of the RXE-100. This restriction does not apply anymore and all 12 PCI-X slots are fully accessible to the operating system.

Similarly, VMWare ESX Server 1.5.x could only use the first six slots. It would crash if all 12 slots were installed in the RXE-100. VMWare ESX Server 2.0 does not have this restriction.

The bandwidth of the connections between the x445 and the RXE-100 is 1 Gbps or 2 Gbps, depending on the model of the RXE-100:

- 8684-1RX supports 1 Gbps
- 8684-2RX supports 2 Gbps

The sharing of one RXE-100 between two four-way x445s, two eight-way x445s, or two 16-way x445 configurations is also supported as described in “Sharing the RXE-100 with two x445 configurations” on page 83.

The RXE-100 has six 64-bit PCI-X slots as standard and can accept adapters with speeds ranging from 33 MHz to 133 MHz. With the optional six-slot expansion kit (part number 31P5998) installed, the RXE-100 has 12 slots. Each set of six adapter slots is divided into three buses of two slots each, as shown in Figure 3-9 on page 77.

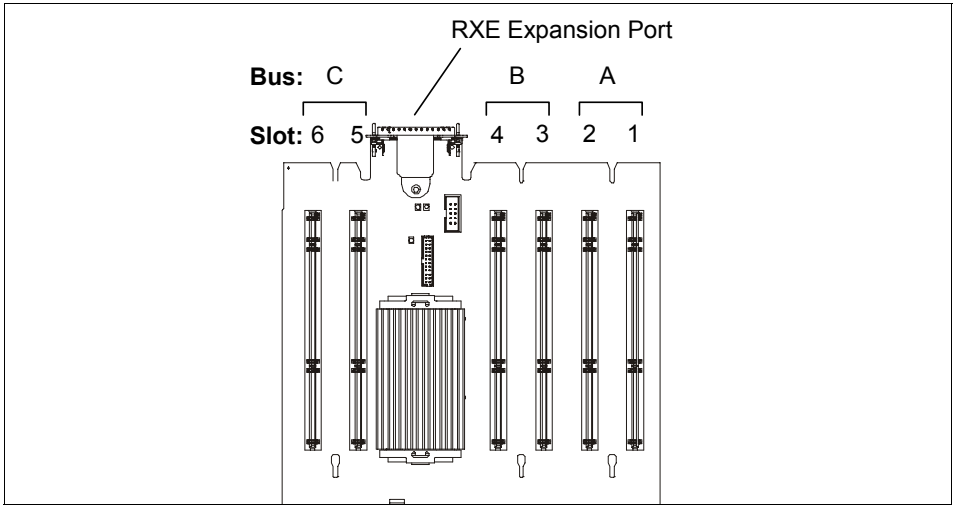

Figure 3-9 RXE-100 PCI-X expansion board (6 slots)

For each of the three buses (A, B, C), one of the following can be installed:

- One 64-bit 3.3 V PCI-X 133 MHz adapter (in the odd-numbered slot), running at 133 MHz
- Two 64-bit 3.3 V PCI-X 133 MHz adapters running at 100 MHz
- Two 64-bit 3.3 V PCI or PCI-X, 33 or 66 MHz adapters

Note: The PCI slots supports 3.3 V adapters only. 5 V adapters are not supported.

Like the x445, these slots can accept adapters rated at speeds ranging from 33 MHz to 133 MHz. When deciding which adapters to put in which slots, consider the following:

- Each adapter has a maximum rated speed and each bus also has a maximum rated speed.
- Installed adapters will operate at the slowest of three speeds:
  - The rated speed of adapter 1 in the bus
  - The rated speed of adapter 2 in the bus
  - The rated speed of the bus
- 32-bit adapters can be installed in any of the slots and will run in 32-bit mode. 32-bit and 64-bit adapters can coexist in 64-bit slots in the same bus. The 32-bit adapters will run in 32-bit mode, and the 64-bit adapters will run in 64-bit mode.
- When installing a 133 MHz PCI-X adapter, it must be installed in the first or odd-numbered slot in the bus (that is in slots 1, 3 or 5).
- Like the x445, a PCI-X and a PCI adapter can be installed in slots on the same bus in the RXE-100. However, these two adapters will both operate in PCI mode.

  In addition, if you have a PCI-X adapter installed, you cannot hot-add a PCI adapter to the same bus. This is because with just the PCI-X adapter installed, the bus is running in PCI-X mode, and you cannot hot-add a PCI adapter into a bus that is in PCI-X mode.
- With Windows NT 4.0 Enterprise Edition, certain token-ring adapters do not work in some slots in the RXE-100. See RETAIN® tip H175383 for more information:

  http://www.pc.ibm.com/qtechinfo/MIGR-42139.html

Cabling considerations

There are two types of cables used to connect the RXE-100 to the x445:

- Remote I/O cable, for data

  This cable connects from the x445 RXE Expansion Port A or B to the RXE-100. Two lengths are available:

  – 3.5 m Remote I/O cable kit (part number 31P6102)
  – 8 m Remote I/O cable kit (part number 31P6103)

  The RXE-100 ships with a 3.5 m Remote I/O cable to connect the unit to the x445. This cable is long enough when both devices are in the same rack. For installation in an adjacent rack, use the optional 8 m Remote I/O cable kit.

- Interconnect management cable, for control

  The RXE-100 also includes a 3.5 m interconnect management cable, which in single-node configurations is used to connect the RXE Management Port on the x445 to the RXE Management A (In) Port on the RXE-100.

  Important: This interconnect management cable has standard RJ-45 connectors, but it is not wired the same as an Ethernet or crossover cable. Ensure the proper cable is used to connect the server to the RXE-100.

  The management cable has two twisted wire pairs: pins 2 and 3 connect to one pair, pins 7 and 8 connect to another. With the Ethernet cable, the wire pairs are connected to the following pins: 1 and 2, 3 and 6, 4 and 5, 7 and 8.

  Two lengths are available:

  – 3.5 m interconnect management cable kit (part number 31P6087)
  – 8 m interconnect management cable kit (part number 31P6088)

  If the RXE-100 has the second set of six PCI slots installed, use the short interconnect management cable (supplied with the PCI slot option kit) to connect Management A (out) Port to Management B (in) Port.

  The 8 m interconnect management cable is suitable for inter-rack configurations.

  For 16-way configurations, the management ports must be connected as shown in Figure 3-12 on page 82. You will need to order an additional management cable.

When connecting the RXE-100 to the x445, the recommended cable connections depend on whether the server has one or two SMP Expansion Modules.

The connections to an x445 with one SMP Expansion Module are shown in Figure 3-10 on page 80. The RXE-100 can have six or 12 slots installed.

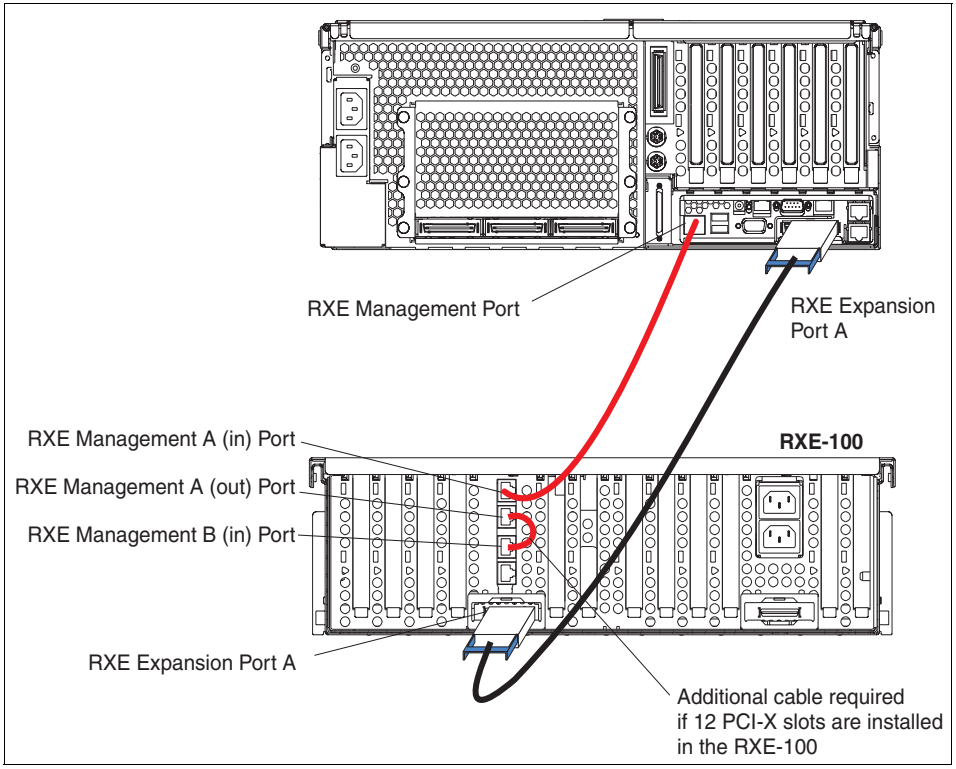

Figure 3-10 Connecting an RXE-100 to an x445 with one SMP Expansion Module

The connections to an x445 with two SMP Expansion Modules are shown in Figure 3-11 on page 81. You must install the second six slots in the RXE-100 and have two cables connected from the x445 to the RXE-100 as shown. The two cables form a redundant pair, which means the system will continue to operate even if one of the two cables fails.

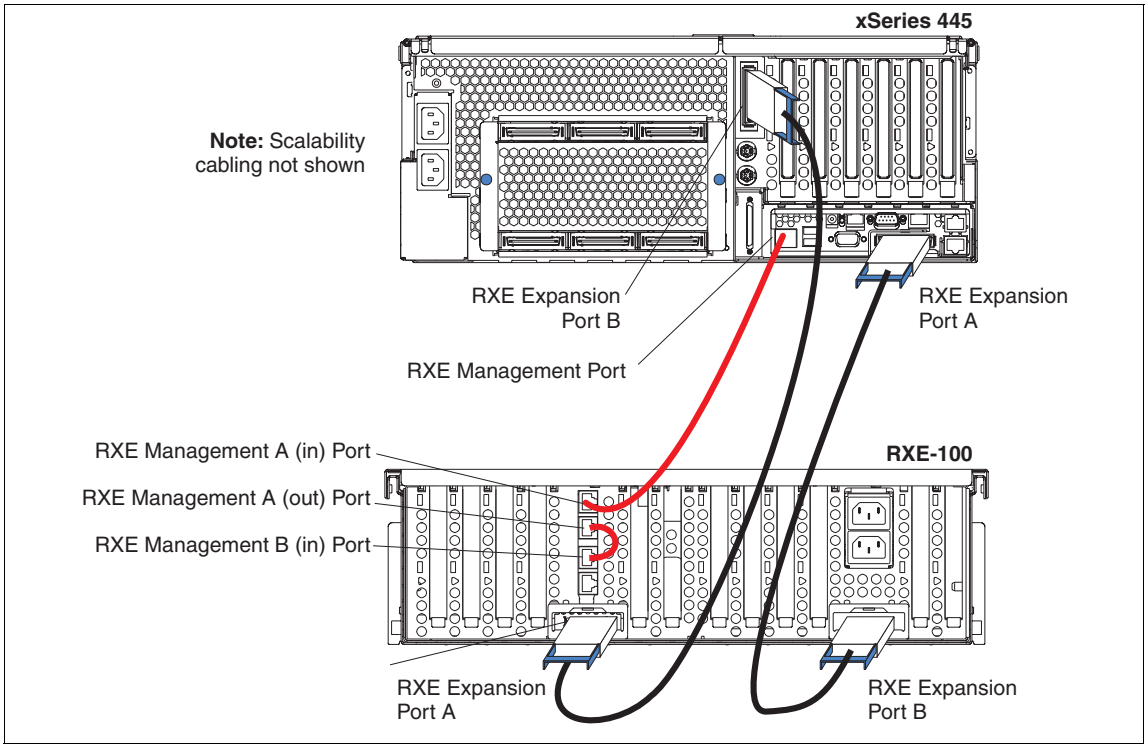

Figure 3-11 Connecting an RXE-100 to an x445 with two SMP Expansion Modules

In the 16-way configuration (that is, two x445 nodes), only one RXE-100 can be connected as shown in Figure 3-12 on page 82. Three Remote I/O cables are used: two to connect the x445s to the RXE-100 and one to connect the two x445s together.

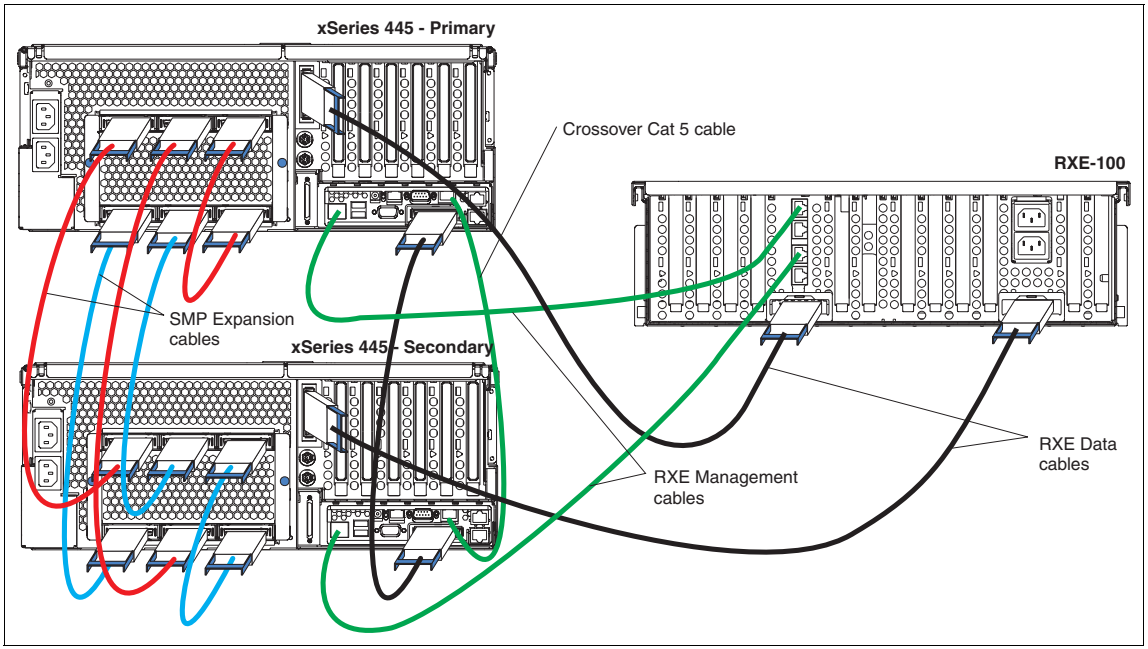

Figure 3-12 Connecting an RXE-100 to a 16-way x445 configuration

Such cabling ensures each SMP Expansion Module has local access to six PCI-X slots, which can increase I/O performance:

- The six internal PCI-X slots are “local” to the bottom SMP Expansion Module in each x445
- Six of the 12 PCI-X slots in the RXE-100 are “local” to each the top SMP Expansion Module in each server.

All 12 slots in the RXE-100 are available to the operating system, with six slots being accessed over each cable. If one cable connection fails, all 12 slots are accessed over the surviving cable connection. It is not currently supported to have each x445 node of a 16-way configuration connected to a separate RXE-100.

One 3.5 m Remote I/O cable ships with the RXE-100. The other two must be ordered separately. Use either the 3.5 m or the 8 m Remote I/O cable.

Tip: Powering the RXE-100 on or off is controlled by the x445, via the interconnect management cable and under the control of the Remote Supervisor Adapter.

Sharing the RXE-100 with two x445 configurations

The x445 will support sharing of the RXE-100 in the following configurations only:

- The RXE-100 must have 12 PCI-X slots installed.
- Supported configurations are four-way, eight-way, and 16-way as shown in Figure 3-13 (note that SMP Expansion cables and RXE management cables are not shown).

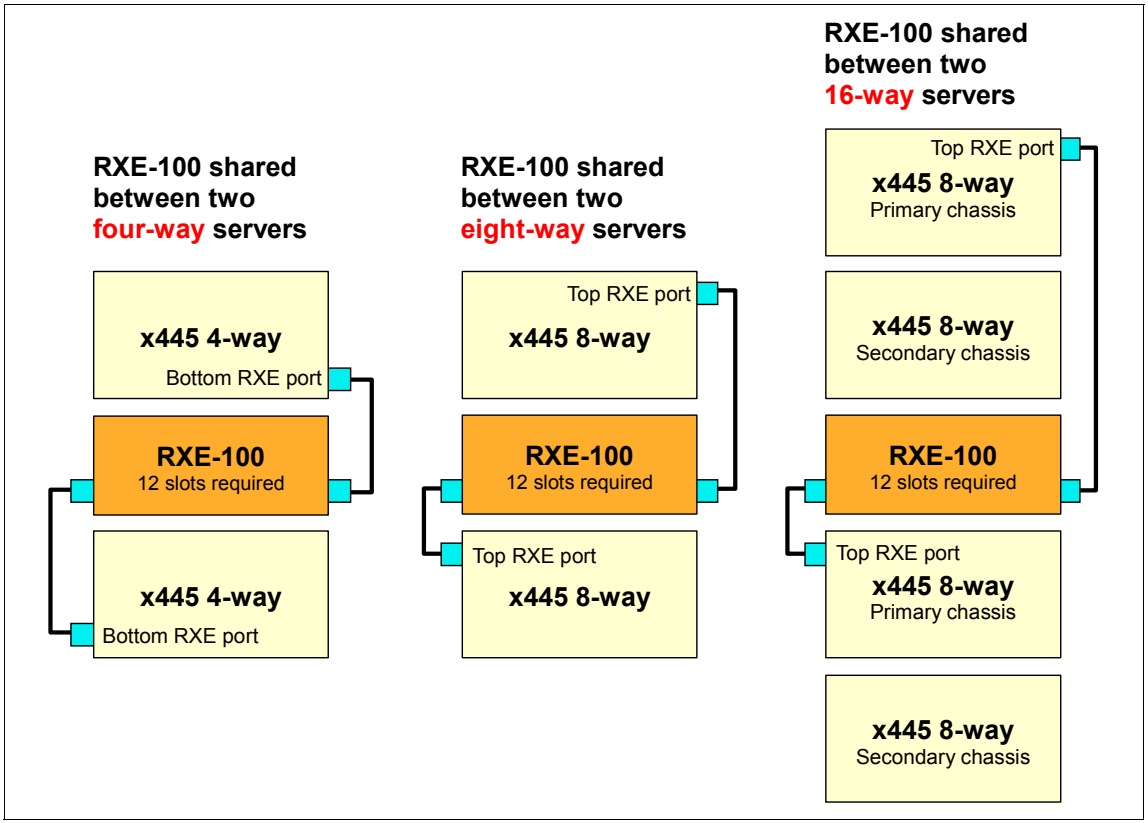

Figure 3-13 Sharing an RXE-100 between two x445 configurations

Restriction: There are no supported configurations for sharing the Remote I/O enclosures between x440 systems, between x440 and x445, between x360 and x445, between x445 and x450, or any other combination.

3.3.4 Serial connectivity

The x445 does not have a standard external serial port, however, the serial port controller is present on the centerplane. If you require the serial port, you can use the external serial connector which is supplied with the x445. Use the PCI-X slot 1 to install the bracket with the serial connector, and connect the ribbon cable to connector J7 on the centerplane. You must enable the serial port in the Setup utility before you can use it.

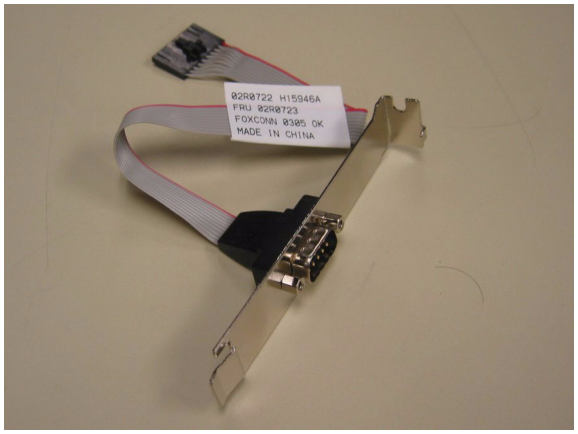

This cable connects to connector J7 on the centerplane to supply an external serial port.

Figure 3-14 Serial cable connector

Another option is the Remote Supervisor Adapter. You can use its serial port as either dedicated to the management tasks, or as shared between the service processor and the operating system. If it is shared, the Remote Supervisor Adapter will have control over the serial port while the server is powered off or booting (running the POST). Once the operating system boots, it takes control over the serial port.

We recommend that you use the dedicated serial port. However, if you need all six internal PCI slots and also the Remote Supervisor Adapter serial connectivity, you still have another option: you can use the USB-to-serial converter, part number 10K3661 (Belkin USB-to-serial converter).

3.4 Storage planning considerations

The x445 has two internal hot-swap 1” (one inch) drive bays (see Figure 1-11 on page 24). Drives up to 15,000 RPM and the converged tray design are supported. These two disk drives would normally be used to install the operating system. The user data would usually reside on an externally-attached high-performance disk subsystem, likely of a SAN variety.

We suggest to install two internal disk drives and configure them as a RAID-1 array, as this will provide higher system availability. You can do this in two ways:

- Using the integrated LSI Logic chip, which supports RAID-1.
- Using a supported ServeRAID adapter (such as the 6M, 4Mx or 4Lx), since these can be used for connection to the hot-swap backplane of the internal drive bays. The ServeRAID-4H is supported when used for external storage enclosure connectivity only, because the adapter is too high for the 4U chassis when the internal SCSI connector is in use.

We recommend you use the integrated LSI Logic SCSI controller, because you can use the ServeRAID channels more efficiently if you connect them to the external expansion enclosures.

Note: The LSI Logic RAID-1 is not compatible with ServeRAID RAID-1. Therefore, if the operating system is installed on drives connected to the integrated LSI chipset, the operating system will need to be reinstalled should the customer decide to move to a ServeRAID adapter. For more information see:

http://www.ibm.com/pc/support/site.wss/MIGR-46174.html

The x445 has two cables pre-installed for use with the internal drive bays, but one is not connected. The shorter cable is initially connected from the onboard SCSI to the drive backplane. When you install a ServeRAID adapter for use with the internal drive bays, you will need to disconnect this cable and connect the longer cable. See “Cabling a ServeRAID adapter” in Chapter 2 of the *IBM @server Series 445 Installation Guide* for details.

3.4.1 ServeRAID with external storage enclosures

The current ServeRAID-4 family of adapters includes the ServeRAID 4H, 4Mx and 4Lx. These 64-bit, Active PCI controllers include advanced features such as Logical Drive Migration, nine RAID levels including RAID 1E, 1E0 and 5E, as well as adapter and cluster failover. The ServeRAID code V6.00 introduced one additional RAID level: RAID-5EE. The latest ServeRAID code level at the time of writing is V6.10, and we recommend you to apply it.

- ServeRAID-6M

  This 133 MHz 64-bit adapter has two Ultra320 SCSI channels with two internal and two external connectors. Two versions of the adapter are available: one with 128 MB of cache and one with 256 MB of cache.

- ServeRAID-4H

  ServeRAID-4H features four Ultra160 SCSI channels, 128 MB of removable battery-backed ECC cache memory, and an IBM PowerPC® 750™ processor onboard. Up to 56 Ultra160 and Ultra2 SCSI devices are supported.

- ServeRAID-4Mx

  This controller features two Ultra160 SCSI channels, 64 MB of battery-backed ECC cache memory, and an Intel i80303 processor. Up to 28 Ultra160 and Ultra2 SCSI devices are supported.

- ServeRAID-4Lx

  ServeRAID-4Lx features one Ultra160 SCSI channel, 32 MB of ECC cache memory, and an Intel i80303 processor. Up to 14 Ultra160 and Ultra2 SCSI devices are supported.

Each ServeRAID adapter supports up to 14 drives (at 160 MB per second throughput) per channel for an aggregate of up to 56 drives (and 640 MBps for the four-channel ServeRAID-4H adapter, for example). Multiple adapters can be installed as needs and available slots dictate.

For the complete list of support RAID adapters, see ServerProven®:

http://www.pc.ibm.com/us/compat/machines/x445.html

For a comparison between the members of the ServeRAID family see:

http://www.redbooks.ibm.com/abstracts/tips0054.html

EXP300

The EXP300 storage expansion unit holds 14 drives and has a maximum 1 TB of disk storage (using 73.4 GB drives) in a 3U package, allowing up to 14 expansion units to be used in a standard 42U rack (meaning that a full rack of EXP300 units can hold 14 TB). The EXP300 provides Predictive Failure Analysis (PFA) on key components, including hot-swap fans, hard drives and redundant power supplies. The EXP300 is optimized for Ultra160 SCSI, with a sustained data transfer rate of 160 MBps.

For more information on IBM SCSI RAID storage solutions go to:

http://ibm.com/pc/ww/eserver/xseries/scsi_raid.html

Important: ServeRAID controllers and EXP300 are supported as the shared storage for the Microsoft Cluster Server in Windows 2000 or Windows Server 2003 environments, but we do not recommend you to use this approach. Instead, we suggest to use the Fibre Channel-based solutions for your shared storage.

3.4.2 The IBM FAStT Storage Servers

The IBM Fibre Array Storage Technology (FAStT) family of Fibre Channel storage solutions is designed for high-availability, high-capacity requirements. FAStT solutions can support transfers over vast distances at rates of up to 200 MBps.

The FAStT Storage Server is a RAID controller device that uses Fibre Channel (FC) interfaces to connect the host systems and the disk drive enclosures. The Storage Server provides high system availability through the use of hot-swappable and redundant components. We briefly discuss the following products:

- The IBM TotalStorage® FAStT200 Storage Server
- The IBM TotalStorage FAStT600 Storage Server
- The IBM TotalStorage FAStT700 Storage Server
- The IBM TotalStorage FAStT900 Storage Server

The FAStT200 Storage Server

The FAStT200 Storage Server is a 3U rack-mountable Fibre Channel RAID controller and disk drive enclosure. It targets the entry and midrange segment of the FC storage market. A typical use of the FAStT200 would be in a two-node cluster environment with up to 30 Fibre Channel disk drives attached to the Storage Server.

Two models are available:

- The FAStT200 Storage Server, with a single RAID controller.
- The FAStT200 High Availability (HA) Storage Server, which contains two RAID controllers and can therefore provide higher availability.

Both models feature hot-swap and redundant power supplies and fans and you can install up to 10 slim-line or half-high FC disk drives. If you need to connect more than 10 disks, you can use the EXP500 FC storage expansion enclosures.

Each EXP500 can accommodate 10 additional disk drives, and up to five EXP500s or up to four EXP700s are supported on the FAStT200 HA. The maximum supported number of disk drives is therefore 60.

The use of hot-swappable and redundant components provides high availability for the FAStT200 Storage Server. A fan or a power supply failure will not cause downtime and such faults can be fixed while the system remains operational. The same is true for a disk drive failure if fault-tolerant RAID levels are used. With two RAID controller units and proper cabling, a RAID controller or path failure will not cause loss of access to data.

Each RAID controller has one host and one drive FC connection. The FAStT200 HA model can use the two host and drive connections to provide redundant connection to the host adapters and to EXP500 enclosures. Each RAID controller unit also contains 128 MB of battery-backup cache.

Important: We do not recommend attaching the EXP500 expansion enclosures to a base-model FAStT200 with a single RAID controller, because a single point-of-failure occurs in such setup. If you want to connect more than just ten internal disk drives, then you should use the FAStT200 HA with two RAID controllers.

The FAStT600 Storage Server

The FAStT600 is an entry level, scalable 2 Gbps Fibre Channel Storage Server. It is a follow-on product to the FAStT200 and it uses similar 3U rack-mountable design, but with some important improvements:

- It uses two RAID controllers as standard. Each controller has 128 MB of cache.
- Four host ports are standard (two per controller)
- It is a 2 Gbps device
- It contains 14 Fibre Channel disk drive bays

Just as the FAStT200, the FAStT600 uses redundant hot-swap power supplies and fans for higher availability.

As the FAStT600 is a 2 Gbps device, you should use the EXP700 expansion unit to increase the storage capacity. You may attach up to two EXP700 expansion units for a maximum of 42 FC disk drives.

The new turbo feature for the FAStT600 offers increased performance, larger cache size, and greater scalability. Available on newly ordered FAStT600 or as an upgrade to an existing FAStT600, the turbo feature expands the ability of the FAStT600. A FAStT600 with the turbo feature includes 2 GB of cache, 1 GB per RAID controller, FAStT Storage Manager V8.4; and authorization for eight storage partitions and attachment of up to seven FAStT EXP700 Expansion Units.

The FAStT700 Storage Server

The FAStT700 Storage Server would typically be implemented in high-end cluster and server consolidation environments and where multiple servers are being consolidated onto a smaller number of x445 systems.

It is a 4U rack-mountable device that supports the 2 Gbps fibre connectivity.

When connected via mini-hubs to the FAStT FC-2 Host Bus Adapter (HBA) and the 2109 F16 Fibre Channel switch, it will give full 2 Gbps fabric.

To avoid single points of failure, it also features dual hot-swappable RAID controllers, dual redundant FC disk loops, write cache mirroring, redundant hot-swappable power supplies, fans, and dual AC line cords.

It supports FlashCopy®, Dynamic Volume Expansion and Remote Mirroring with controller-based support for up to 64 storage partitions. RAID levels 0,1, 3, 5, and 10 are supported and for performance it includes a total of 2 GB battery-backed cache (1 GB per controller).

The FAStT900 Storage Server

The FAStT900 is an enterprise-class Storage Server and a follow-on product to the FAStT700. It supports up to 32TB of storage capacity with 224 146.8 GB FC drives installed in 16 EXP700 disk enclosures.

The FAStT900 excels in performance. The dual high-performance RAID controllers offer the highest throughput of all the FAStT products: up to 772 MBps. For comparison, the FAStT700 supports up to 390 MBps of throughput.

The mechanical design and most of the features are similar to those of the FAStT700.

Additional information on the entire range of FAStT storage solutions can be found at:

`http://www.storage.ibm.com/disk/fastt/index.html`

For more detailed information on the FAStT Storage Servers, we recommend you to have a look at the following redbooks:

- *Fibre Array Storage Technology A FAStT Introduction, SG24-6246-00*
- *IBM TotalStorage FAStT700 and Copy Services, SG24-6808-00*
- *IBM TotalStorage Disk Solutions for xSeries, SG24-6874-00*

FC2-133 Host Bus Adapter considerations

We recommend you use the FC2-133 Host Bus Adapter (part number 24P0960) to connect to a FAStT controller. This will give you connectivity at full 2 Gbps. This adapter is a 133 MHz 64-bit PCI-X adapter and you should install it into the adequate PCI-X slot, so that it can run at maximum speed:

- Slot 5 or 6 in the x445
- Slot 3 or 4 in the x445, providing the other slot on the same bus is empty

- Slot 1, 3 or 5 in the RXE-100 six-pack, providing the other slot in the same bus is empty

FAStT Storage Manager versions

Initially, the FAStT Storage Manager v8.x was not freely available for all the FAStT Storage Servers. Customers who were using v7.x and wanted to upgrade to 8.x, needed to pay for the additional functionality introduced in the associated firmware release. This is no longer the case. Customers can now download the FAStT Storage Manager v8.3 and the associated firmware files from the following link:

http://www.ibm.com/pc/support/site.wss/MIGR-50176.html

3.4.3 The Enterprise Storage Server® (ESS)

The ESS provides integrated caching and RAID support for the attached disk devices. ESS can be configured in a variety of ways to provide scalability in capacity and performance. One ESS can support up to of 56 TB and can utilize 2 Gbps Fibre Channel connectivity.

Redundancy within ESS provides continuous availability. It is packaged in one or more enclosures, each with dual line cords and redundant power. The redundant power system allows ESS to continue normal operation when one of the line cords is deactivated.

The ESS provides an image of a set of logical disk devices to attached servers. The logical devices are configured to emulate disk device types that are compatible with the attached servers. The logical devices access a logical volume that is implemented using multiple disk drives. This allows ESS to connect to all IBM servers, from zSeries to iSeries™, pSeries and xSeries, directly or through a SAN, thus helping the x445 fit into a heterogeneous environment containing a variety of server architectures. The ESS offers several choices of host I/O interface attachment methods, including SCSI and Fibre Channel for the xSeries servers.

For more information on the ESS go to:

http://www.storage.ibm.com/disk/ess/index.html

3.4.4 Disk subsystem performance

Because of the processing capacity of the x445, a poorly designed storage subsystem could become a bottleneck, seriously impacting overall system performance. You should implement a disk subsystem that is able to efficiently

process the potentially massive number of disk I/O requests generated by the processor subsystem.

Assuming that you will be implementing a RAID storage configuration, the rule of thumb is that more physical disks will improve the throughput of your disk subsystem and consequently overall system performance. In almost all applications, adding disks to your RAID configuration will continue to improve performance until another system component becomes a bottleneck. A storage solution with too few physical disks will become a bottleneck for the entire system.

You will need to carefully analyze your storage capacity requirements, your application requirements, and your host requirements before you finalize your storage solution.

A detailed discussion on performance tuning disk subsystems in xSeries servers is available in the redbook *Tuning IBM @server xSeries Servers for Performance*, SG24-5287.

3.4.5 Tape backup

As with your disk subsystem, you need to carefully analyze backup requirements before you select a tape solution. Considerations when selecting a backup solution should include:

- Currently implemented backup solutions

  If you are consolidating a number of servers onto an x445-based solution, for example, you may want to take the opportunity to move from differing and distributed tape technologies (such as DDS and DLT) and consolidate those into a single, high-performance, automated solution. An example is the IBM Ultrium Autoloader.

- Current and projected capacity requirements

  Select a solution that has the ability to scale as capacity requirements increase.

- Performance requirements

  You need to consider the backup window available, as well as the amount of data being backed up when determining what your backup performance requirements will be. It is also important to consider the need for quick access to data committed to tape when selecting a solution.

- Connection requirements

  Will the tape solution be connected to an existing SAN fabric and if so, will this require additional fabric hardware?

- Hardware and software compatibility

  If you implement a new tape solution, you need to ensure that current backup and management software is still suitable. Disaster recovery procedures may also need to be revised.

IBM offers a full range of high-performance, high-capacity and automated tape solutions for the xSeries servers. For detailed information on these products, go to:

http://ibm.com/pc/ww/eserver/xseries/tape.html

The following redbooks discuss IBM tape solutions in greater detail:

- *Netfinity Tape Solutions*, SG24-5218
- *The IBM LTO Ultrium Tape Libraries Guide*, SG24-5946

3.5 Server partitioning and consolidation

The concepts of server partitioning and consolidation are discussed in detail in Chapter 2, “Positioning” on page 27.

Implementing a server consolidation solution using VMware and the x445 introduces a number of important and unique issues that you need to consider during the planning phase of your project.

In particular the x445 configuration needs to be carefully sized to meet the resource requirements of the VMware operating system, the guest operating systems, and the applications being deployed.

A detailed discussion on planning, sizing, and implementing VMware solutions can be found in the redbook *Server Consolidation with the IBM @server xSeries 440 and VMware ESX Server*, SG24-6852. Even though the book discusses the x440 and VMware ESX Server v1.5.x, it contains very useful information that can be applied also to the x445 and VMware v2.0.

3.6 Operating system considerations

In line with the overall message of providing application flexibility to meet the varying needs of our enterprise customers, the x445 is optimized for numerous operating system and application solutions. For the latest operating system support information, go to:

http://www.pc.ibm.com/us/compat/nos/matrix.shtml

As described in 1.4.1, “Intel Xeon Processor MP” on page 12 Hyper-Threading technology allows a single physical processor to appear to the operating system and applications as two logical processors. The logical processors share the core processing engine of the physical processor but are able to execute code streams concurrently.

Operating systems must be “Hyper-Threading aware” before they can “see” the additional processors. When they are, they will “see” twice as many CPUs as there really are (see Figure 5-9 on page 160 for an example).

Simply enabling Hyper-Threading may not guarantee improved overall system performance, however. In order to benefit from enabling Hyper-Threading, the operating system and server applications need to be capable of detecting the additional logical processors and spawning multiple threads, which can exploit the additional processing power.

As well as considering whether the operating system you are installing supports Hyper-Threading, there may be licensing implications to consider before enabling Hyper-Threading technology.

For a more detailed discussion on Hyper-Threading technology, refer to:

http://www.intel.com/ebusiness/hyperthreading/server

Table 3-3 lists the supported operating systems for the x445 and the level of support for Hyper-Threading technology provided by the operating system.

In the column titled Hyper-Threading Support:

- **None** indicates the operating system does not recognize the logical processors that Hyper-Threading enables.
- **Yes** indicates that the operating system recognizes the logical processors and can execute threads on them but is not optimized for Hyper-Threading. From a licensing perspective, if Hyper-Threading is enabled, the operating system will need to be licensed for twice the number of physical processors to take full advantage of the processors’ capabilities.
- **Optimized** indicates that the operating system recognizes the logical processors and the operating system code has been designed to take full advantage of the technology. From a licensing perspective, the logical processors do not count towards the number of processors an operating system is licensed for.
- **Ignores** indicates the operating system recognizes the logical processors when Hyper-Threading is enabled, but the operating system ignores them.

Table 3-3 x445 operating system support

Description	Release	x445 SMP support[1]	Hyper-Threading
Windows NT Enterprise Edition	4.0	Only one SMP Expansion Module is supported (that is, four-way Xeon MP or two-way Xeon DP). Hot-plug PCI not supported	None
Windows 2000 Server	SP3	Supports up to four-way	Yes
Windows 2000 Advanced Server	SP3	Supports up to eight-way	Yes
Windows 2000 Datacenter Server		Supports up to 16-way	Yes
Windows Server 2003, Standard Edition	Initial	Supports up to four-way	Optimized
Windows Server 2003, Enterprise Edition	Initial	Supports up to eight-way	Optimized
Windows Server 2003, Datacenter Edition	Initial	Supports up to 16-way on current x445 models. Future x445 models will support 32-way	Optimized
NetWare	6.0	Supports up to eight-way [2]	Yes
Red Hat Enterprise Linux AS	2.1	Supports up to eight-way [3]	Yes
SuSE Linux Enterprise Server	8.0	Supports up to 8-way on current x445 models. SuSE can operate a 16-way but this is not currently certified. Future x445 models will support 32-way	Yes
VMware ESX Server	2.0	Supports up to 16-way Supports two processors per VM [4]	Ignores [5]
SCO UnixWare	7.1.3	Supports up to eight-way	Yes

Notes to Table 3-3:

1. While operating systems may support eight-way or larger systems, scalability is a function of both the operating system and the application/workload. Applications must be specifically designed to take advantage of larger SMP systems.
2. NetWare notes:
   – NetWare does not support multiple-chassis configurations.

- NetWare 5.1 is not supported, but it should still install. See RETAIN tip H176163 for details on a known shutdown problem:

  http://www.pc.ibm.com/qtechinfo/MIGR-43679.html

- With NetWare 6.0, the server may show extreme CPU utilization values (for example, 13000%). This was fixed with NetWare 6.0 Support Pack 2. See RETAIN tip H176060 at:

  http://www.pc.ibm.com/qtechinfo/MIGR-43532.html

  At the time of writing, the latest NetWare Support Pack is 3.

  NetWare 6.5 has been released recently. However, at the time of writing it was not listed on ServerProven. NetWare 6.5 will support up to 32 processors. For the latest information, see:

  http://www.pc.ibm.com/us/compat/nos/matrix.shtml

3. Ongoing work will improve both Linux and key application scalability. Currently, the general recommendation is to keep system size limited to eight-way and below, and 16 GB and below. Work on scalability beyond eight-way is in progress, and is likely to become available in 2004.

   For the latest information, see:

   SuSE: http://www.suse.com/us/business/products/server/sles/x86.html

   Red Hat: http://www.redhat.com/software/rhel/as/

4. Virtual machines running on VMware ESX Server 2.0 will only be dual processor capable if the VMware ESX Virtual SMP module is also purchased.
5. Hyper-Threading can be enabled. However, VMware ESX Server 2.0 will *not* make use of any logical processors. The guest operating system do not “see” any logical processors.

3.6.1 Windows Datacenter

Windows 2000 Datacenter Server and Windows Server 2003, Datacenter Edition are highly scalable network operating systems designed for mission-critical enterprise-wide applications. High-volume online transaction processing, large-scale data warehousing, and scientific simulations are some of the applications that Datacenter is optimized for.

For more information about the IBM Datacenter program, see Chapter 4, “Windows Server 2003, Datacenter Edition” on page 113.

Datacenter configurations

Two installation options are currently available for Datacenter Server on the x445:

- Factory installation. Table 3-4 shows the available x445 eight-way Datacenter models.

Table 3-4 Datacenter Server certified models

Model	Standard processors	Max SMP	Max SMP two nodes	Std memory and disk
8870-1AX	2x 2.0 GHz Xeon MP, 1 MB L3	8-way	16-way	None
8870-2AX	2x 2.5 GHz Xeon MP, 1 MB L3	8-way	16-way	None
8870-4AX	2x 2.8 GHz Xeon MP, 2 MB L3	8-way	16-way	None

When you order these models, you specify which operating system you want preloaded, as shown in Table 3-5. The kit includes installation guide, system documentation, and recovery CD-ROMs. These preload kits also include five client access licenses (CALs). A CAL is required for each client attached to the server.

Table 3-5 Windows Datacenter preload kits

Preload kits	Order number
Windows 2000 Datacenter 1-8 CPUs	4816-3BU
Windows 2000 Datacenter 1-16 CPUs	4816-3DU
Windows 2003 Datacenter 1-4 CPUs	4816-4AU
Windows 2003 Datacenter 1-8 CPUs	4816-4BU
Windows 2003 Datacenter 1-16 CPUs	4816-4DU

Note: All AX models are to be ordered via an IBM Special Bid process, and all Datacenter solutions require an IBM Solutions Assurance Review, to ensure compliance with the Datacenter certified solution.

The licenses in these kits are cumulative. For a 16-way system for example, you can use four 4-way licenses, two 8-way licenses, or one 16-way license.

- Software-only preload kit.

  The preload kits shown in Table 3-5 allow a customer to install Datacenter Server on an existing x445 configuration that is already certified for use with Datacenter Server.

Subscriptions

Customers should also purchase the annual Software Update Subscription or the Maintenance Update Subscription.

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.

Table 3-6 Software Update Subscriptions

Software Update Subscriptions	Order number
Windows 2000 Datacenter Server 8-way	4816-ABX
Windows 2000 Datacenter Server 16-way	4816-ADX
Windows Server 2003 Datacenter Edition 4-way	4816-CAX
Windows Server 2003 Datacenter Edition 8-way	4816-CBX
Windows Server 2003 Datacenter Edition 16-way	4816-CDX

IBM builds, tests, and provides the complete certified package of these components. IBM provides program updates, as they become available for a period of one year. A program update may contain a new version, release, supplements, or service packs as IBM determines, announced during the subscription term. IBM does not guarantee that updates will be announced during the annual term. The subscription is sold on a per-server basis for xSeries servers.

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).

Table 3-7 Maintenance Update Subscriptions

Maintenance Update Subscriptions	Order number
Windows Server 2003 Datacenter Edition 4-way	4816-DAX
Windows Server 2003 Datacenter Edition 8-way	4816-DBX
Windows Server 2003 Datacenter Edition 16-way	4816-DDX

Datacenter 16-way configurations

The 16-way configurations are comprised of two eight-way Datacenter models which are described in Table 3-4 on page 96.

Two of these can form a single 16-way configuration with the addition of the x445 Two Chassis 16-way Configuration Kit, part number 02R2013. The kit contains the following items:

- Four 2.5 m copper-colored scalability cables
- One Ethernet cross-over cable to interconnect the Remote Supervisor Adapters.

Cluster configurations

Datacenter Edition is supported on the eight-way x445 configurations in one-, two-, three-, and four-cluster node configurations and on a 16-way configuration in one-cluster node and four-cluster node configurations. Check for updates to the Microsoft Hardware Compatibility List (HCL) at:

http://www.microsoft.com/whdc/hcl/default.mspx

3.6.2 Microsoft Windows NT 4.0 Enterprise Edition

Hyper-Threading is not supported by Windows NT 4.0 EE, which has not been enhanced to exploit Intel's ACPI. This operating system is therefore unable to recognize logical processors, which are simply ignored.

There are no plans from Microsoft to patch Windows NT 4.0 Enterprise Edition to support Hyper-Threading technology.

A custom Hardware Abstraction Layer (HAL) for the x445 must be installed during the operating system installation and is available from IBM. The HAL is required to support the Active PCI-X slots in the server and the RXE-100. Download it from:

http://www.pc.ibm.com/qtechinfo/MIGR-42067.html

Only one SMP Expansion Module is supported. With Xeon MP processors, at most four processors are supported. With Xeon DP processors, at most two processors are supported.

3.6.3 Microsoft Windows 2000 Server

The members of the Windows 2000 server family support Hyper-Threading, but they have not been optimized for it. They use a custom Hardware Abstraction Layer (HAL), which should be installed during the operating system installation. It is available from IBM at:

http://www.pc.ibm.com/qtechinfo/MIGR-42325.html

From a licensing point of view, logical processors as provided by Hyper-Threading are counted against the Windows licensing limit. Windows will first count physical processors and, if the license permits more processors, then logical processors will be counted. For example, in a four-way x445, Windows 2000 Server will count four physical processors, then stop, because Windows 2000 Server is limited to four processors. In the same x445, Windows 2000 Advanced Server will count eight processors (four physical and four logical) because the license permits up to eight processors.

Eight-way systems: If the server has eight CPUs installed, then you must ensure Hyper-Threading is disabled in the BIOS before installing Windows 2000 Advanced Server. You can re-enable Hyper-Threading after Service Pack 2 or later is installed. If you do not disable Hyper-Threading during installation, a blue-screen trap will occur.

Windows 2000 Advanced Server supports a maximum of eight CPUs, so there is no performance benefit to be gained from enabling Hyper-Threading on an eight-way Windows 2000 Advanced Server system.

3.6.4 Microsoft Windows Server 2003

The members of the Microsoft Windows Server 2003 family can be compared as shown in Table 3-8.

Table 3-8 Features of the Windows Server 2003 family

	Standard	Enterprise	Datacenter	Web
Processors	1-4	1-8	1-32	1-2
Memory	4 GB	32 GB (32-bit)	64 GB (32-bit)	2 GB
Clustering	N/A	1-8 nodes	1-8 nodes	N/A
Network Load Balancing	Yes	Yes	Yes	Yes
Scalability features		NUMA Hot-add memory	NUMA Hot-add memory	
Windows System Resource Manager (WSRM)	No	Yes	Yes	No
Hyper-Threading	Yes	Yes	Yes	Yes
Active Directory	Yes	Yes	Yes	Partial
IIS 6.0	Yes	Yes	Yes	Yes
Terminal Server	Yes	Yes	Yes	N/A

Microsoft has added “NUMA enhanced” CPU-affinity code and Hyper-Threading optimization to maximize the performance of XpandOnDemand NUMA-based servers such as the x445.

Windows Server 2003 operating systems understand the concept of physical processors versus logical processors. In the case of Windows Server 2003, only physical processors will count against the license limit. For example, Windows Server 2003, Standard Edition running on a two-way system with Hyper-Threading enabled will recognize and use the processing capabilities of both physical and logical processors. However, only two physical processors will be counted for licensing purposes.

These operating systems are also aware of the NUMA architecture of the system and use the resources of the system in a more structured and efficient manner, resulting in better overall performance.

NUMA optimization

NUMA technology is based on a cell/node topology. With the x445, the node is the SMP Expansion Module. When the system is configured as an eight-way, there are two nodes in the system each node with its own memory and processors, which introduces the concept of remote and local memory.

Windows Server 2003 Enterprise and Datacenter Editions attempt to use the local resources first, only using remote resources if the local resources are busy. This is accomplished by using a Static Resource Allocation Table (SRAT), an ACPI 2.0 table that is stored in firmware. The SRAT uses the *proximity domain* concept—local components are grouped into domains and the firmware then creates a proximity domain for each node.

The NUMA description within the SRAT allows Windows Server 2003 to assign all tasks to the local node, if available, potentially increasing performance by 30%.

With respect to process scheduling, processes are assigned a home node in addition to a home processor at initialization. Node assignment is done using a round-robin algorithm. The scheduler attempts to schedule threads in the following sequence:

- First on the home processor
- Second on the home node
- Only schedules off-node if the home node is very busy and other nodes are idle

3.6.5 Novell NetWare

Novell NetWare 6.0 is listed on the x445 ServerProven compatibility list.

NetWare support is currently limited to eight-way systems only. When installing, you will need to use the Novell NetWare 6.0 Support Pack 2 or higher Overlay CD. The latest version is Support Pack 3.

3.6.6 Red Hat and SuSE Linux

Red Hat Linux Advanced Server 2.1 Enterprise and SuSE Linux Enterprise Server 8.0 are architected to exploit the capabilities of the x445 server, including support for Hyper-Threading technology. The most recent version of the Linux 2.4.x kernel scales effectively to four-way and eight-way, although results vary depending on the workload and the scalability of the application.

The xSeries development team, the IBM Linux Technology Center, and other parts of IBM are working with the Linux community, Red Hat, SuSE, and other Linux alliance partners to develop advanced Enterprise X-Architecture features in upcoming releases of Linux. These features will include memory optimization, additional hot-swap/hot-add capabilities, dynamic partitioning, and additional scalability improvements.

3.6.7 VMware ESX 2.0 Server

VMware ESX Server is virtual machine software for consolidating and partitioning servers. It is a cost-effective, highly scalable virtual machine platform with advanced resource management capabilities. VMware ESX Server is used to minimize the total cost of ownership of server infrastructure by maximizing server manageability, flexibility, and efficiency across the enterprise.

Key new features in ESX Server 2.0 are the following:

- Support for two-way SMP virtual machines using an add-on product
- Improved support for NUMA architecture in the x445
- Support for NetWare as a guest operating system
- Support for NIC-teaming for high availability

Support of VMware ESX 2.0 Server is supported on the x445. See ServerProven for the latest information:

http://www.pc.ibm.com/us/compat/nos/matrix.shtml

Part numbers to order are based on the number of CPUs to be installed in the new hardware. Order numbers are shown in Table 3-9. If you already have ESX

Server 2.0 installed and add processors, you can purchase an upgrade license to match the new number of CPUs.

Note: VMware ESX Server 2.0 must be ordered in conjunction with a supported xSeries server. These offerings cannot be ordered separately from an xSeries server.

Table 3-9 Order numbers for ESX Server 2.0

CPUs installed in the server	Order number	Upgrade
1-2 processor license	4817-22X	Not applicable
3-4 processor license	4817-24X	4817-29X
5-8 processor license	4817-28X	4817-2AX
9-16 processor license	4817-26X	4817-2BX
Media and documentation	4817-21X	Not applicable

New to ESX Server 2.0 is support for two-way SMP virtual machines. This support is made with the add-on component VMware ESX Virtual SMP. You need to order this component in addition to the base ESX Server license.

Table 3-10 Order numbers for VMware ESX Virtual SMP

CPUs installed in the server	Order number
1-2 processor license	4817-2CX
3-4 processor license	4817-2DX
5-8 processor license	4817-2EX
9-16 processor license	4817-2FX

In addition to the above feature additions, ESX Server 2.0 has the following enhancements:

- Provides SAN failover

  ESX Server 2.0 maximizes storage area network (SAN) availability by providing multipathing or automatic path failover in case of HBA, switch, or storage controller failure.

- Improved VMware ESX Server File System

  ESX Server 2.0 includes a new VMFS Version 2 (VMFS-2). VMFS-2 has the ability to span multiple partitions and the ability for multiple ESX Servers (and the virtual machines on these servers) to access VMFS-2 concurrently.

- Improved VMware Management Interface

  ESX Server 2.0 introduces a new VMware Management Interface that contains improved navigation, better performance, and an easier-to-use interface.

- Support for new guest operating systems:
  - Microsoft Windows XP Professional
  - Microsoft Windows Server 2003
  - Red Hat Linux 9.0
  - SuSE Linux 8.x operating systems

- Virtual LSI Logic SCSI adapter

  ESX Server 2.0 supports the LSI Logic SCSI virtual adapter. By default, ESX Server assigns a virtual BusLogic SCSI adapter to Linux, Microsoft Windows NT 4.0, Windows 2000 or Windows XP Professional guest operating systems. Similarly, ESX Server assigns a virtual LSI Logic SCSI adapter to Windows Server 2003 guest operating systems.

  You can change these default settings by editing the virtual machine's configuration file.

- NUMA enhancements include:
  - Support for manually specifying affinities for best possible CPU and memory location.
  - You can specify a minimum CPU percentage, a maximum CPU percentage, or both for each virtual machine. The minimum CPU percentage represents an absolute fixed lower limit, while the maximum CPU percentage represents an absolute fixed upper limit.

    For example, if you create a virtual machine and a minimum CPU percentage of 30% and a maximum CPU percentage of 80%, then this virtual machine will always have at least 30% and never more than 80% of a physical CPU's resources.
  - ESX Server 2.0 can also automatically set affinities for optimized CPU and Memory location based on best-guess heuristics.
  - Smart location of shared pages to increase performance through page migration and remapping.

For more technical information, IBM employees may use the following URL:

http://ats.greenock.uk.ibm.com/vmware_new.htm

Notes:

- ESX Server 2.0 must be ordered in conjunction with a supported xSeries server. It cannot be ordered separately, unless you contact VMware directly.

- Each VMware ESX Server 2.0 includes a one-year Software Update Subscription Service from VMware. During this period, product and maintenance updates are made available to registered owners. After the first year, this is a fee-based offering.
- In addition to the Software Update Subscription Service, other relevant fee-based services include:
  - Support services are provided on a fee basis through SupportLine offerings. VMware ESX Server 2.0 is a supported product under the Microsoft and Linux SupportLine offerings. For additional information, visit:

    http://www.ibm.com/eserver/xseries/vmware
  - 8x5 or 24x7 telephone coverage.
- The x445 has two onboard Gigabit Ethernet controllers. However, these cannot be split between the console OS and the virtual machines. The Ethernet controllers can be either assigned to the console OS or, more likely, to virtual machines, but not both. You will need an additional network adapter for connectivity to the console OS.

ESX Server works by letting you transform physical computers into a pool of logical computing resources. Physical servers are partitioned into secure virtual servers. Operating systems and applications are isolated in these multiple virtual servers that reside on a single piece of hardware. These resources can then be distributed to any operating system or application as needed.

ESX Server provides dynamic logical partitioning. It runs directly on the hardware to partition and isolate server resources, using advanced resource management controls to let you remotely manage, automate, and standardize these server resources.

Dynamic logical partitioning involves:

- Partitioning server resources

  The ESX Server acts as the host operating system, provides dynamic logical partitions to hold other operating systems, and virtualizes most system resources, including processors, memory, network capacity, and disk controllers.
- Isolating server resources

  With ESX Server, each hosted operating system thinks it owns the entire computer, yet it sees only the resources that the administrator (through ESX Server) assigns to it. As shown in Figure 3-15, ESX Server resides between the hardware and the various operating systems and applications. Partitions can be administered remotely, even down to the BIOS level, just as individual servers are.

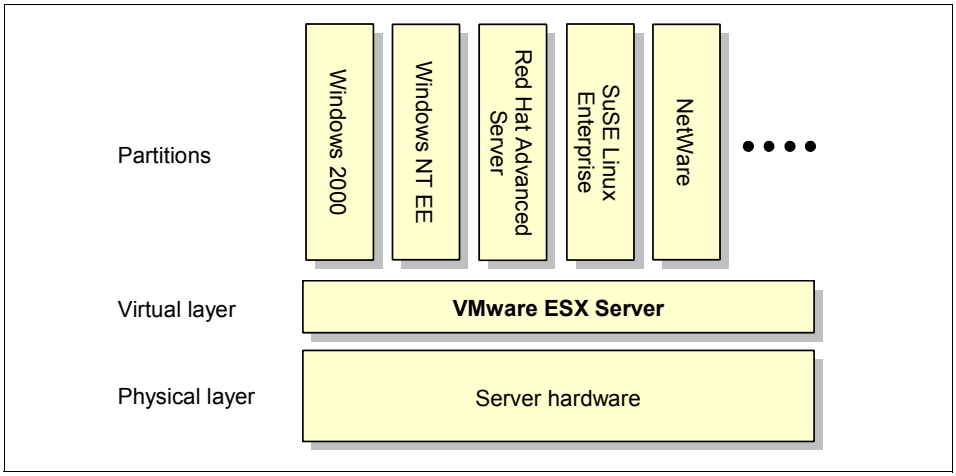

Figure 3-15 ESX Server resides between the server hardware and server resources

- Managing server resources

  The ESX Server's advanced resource management controls allow you to guarantee service levels. CPU capacity can be allotted on a time-share basis. Memory can be assigned dynamically based on partition workloads and defined minimums. If the allocated amount is insufficient in one partition, ESX Server can temporarily borrow memory from one partition and lend it to another, and then restore it to the original partition when needed. Network sharing is determined by token allocation or consumption based on the average or maximum bandwidth requirements for a partition.

Full support is provided for the following operating systems as guests under ESX Server:

- Windows Server 2003 (Enterprise, Standard, and Web Editions)
- Windows XP Professional (Service Pack 1)
- Windows 2000 Server (Service Pack 3 or 4)
- Windows 2000 Advanced Server (Service Pack 3 or 4)
- Windows NT 4.0, Service Pack 6a
- Red Hat Linux 7.3, 8.0 and 9.0.0
- Red Hat Linux Advanced Server 2.1
- SuSE Linux 8.2
- SuSE Linux Enterprise Server (SLES) 8

ESX 2.0 Server detects the logical processors provided by the Hyper-Threading technology. But it does not use them at all - it simply ignores them.

For support information including supported servers, see:

http://www.pc.ibm.com/ww/eserver/xseries/vmware.html

For more information about VMware, refer to the redbook *Server Consolidation with the IBM @server xSeries 440 and VMware ESX Server*, SG24-6852.

Support for applications running on ESX Server

Ensure that the applications you plan to run on the x445 running VMware ESX Server are supported by the application vendor.

- Microsoft

  The statement below is taken from the Microsoft support Web site:

  http://support.microsoft.com/default.aspx?scid=kb;en-us;Q273508

  "Microsoft does not support issues that occur in Microsoft operating systems or programs that occur in the virtual machine until it is determined that the same issue can be reproduced outside of the virtual machine environment."

- IBM software

ESX Server is currently not a supported platform for WebSphere, Lotus, and Tivoli® products. To obtain support for a problem, it must first be reproduced without ESX Server. Work is underway to support ESX Server as a supported platform. Check the support site for the specific products you are planning to install for the latest support information.

3.7 Application considerations

As well as operating systems, there are enterprise applications currently available that are licensed on a per-processor basis. You should be aware that enabling Hyper-Threading or adding physical processors on systems running these applications may have licensing implications for your organization. This will need to be considered in the planning phase of your deployment and, if required, additional licenses purchased prior to enabling Hyper-Threading.

Microsoft has stated that current server products licensed on a per-processor basis will require one license per physical processor. For example, on a two-way system with Hyper-Threading enabled and running Microsoft SQL Server, a two-processor license is required, even though the application may process threads on four logical processors.

Performance benefits versus licensing costs may be a consideration before enabling Hyper-Threading or adding processors and may require testing to confirm that there will be a substantial benefit to application performance. In most applications there will be a performance gain as processors are added; however, this gain does not generally remain linear with the continued addition of processors. The performance improvements seen will depend largely on application scalability, which is discussed in more detail in the next section.

3.7.1 Scalability and performance considerations

Adding processors improves server performance because software instruction execution can be shared among the additional processors. However, this requires software to detect the additional CPUs and generates additional work in the form of threads or processes, which execute on the additional processors. This does not happen automatically. The operating system provides a platform that enables the capability of multiprocessing, but it is up to the application to generate the additional threads and processes to execute on all processors. This is referred to as *application scalability*.

Having faster machines in the server hardware space means more parallelism (more processors, larger memory, larger disk arrays, additional PCI buses, and so on). The obvious case of software that does not scale is DOS. Run DOS on a server with eight CPUs, 16 GB of memory, equipped with 250 15K RPM disks in RAID arrays, and you get similar performance as if you have one CPU, one disk, and 640 KB of memory. Obviously, the server isn't slow. The problem is that the software (in this case DOS) does not scale.

Software scalability is a complex subject, one most people don't consider until it is too late. Often people purchase new high-performance SMP servers expecting huge performance gains with old applications, only to learn the bottleneck is in the server application. In this case there is little they can do to efficiently utilize the new SMP server until the application is modified.

A scalable application makes use of greater amounts of memory, generates scalable I/O requests as the number of disks in a disk array increases, and will utilize multiple LAN adapters when a single LAN adapter limits bandwidth. In addition, a scalable application has to detect the number of installed processors and spawn additional threads as the number of processors increases to keep all processors busy.

Hyper-Threading increases the number of logical processors and demands that the software spawn additional threads to run at maximum efficiency. However, most applications don't yet do this. This is why Hyper-Threading performs well with two-way and four-way SMP systems, because many applications already generate sufficient threads to keep four physical/logical CPUs busy. But at

eight-way and 16-way, the applications have to do even more than they do today to efficiently utilize Hyper-Threading. All of these things must be engineered into the server application and operating system. In general, the only applications that scale past four-way are database applications.

3.7.2 SMP and server types

SMP has a direct relationship to the type of application server being used. If the server is used as a file server, adding a processor to it does not significantly improve performance, whereas it can result in a very high performance gain for an application server.

As you can see from Figure 3-16, a file/print server benefits only marginally from the addition of a second processor and can actually degrade performance when the third and fourth processors are added. However, with a database or application server, the addition of one to three processors makes a marked improvement in processing power.

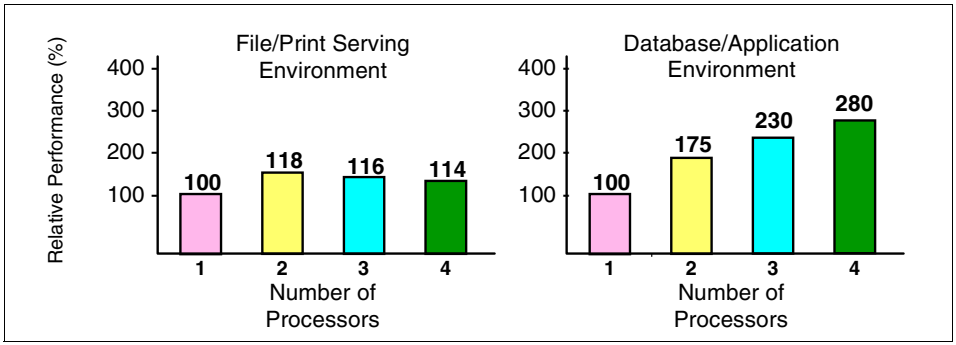

Figure 3-16 Effect of adding processors under file/print and application environments

SMP will not provide a linear improvement in processing power as additional processors are added. You might achieve a 70-90% performance increase from the second processor, but each additional processor after the second will provide a diminishing return on investment as other system bottlenecks come into play.

For more information regarding multiprocessor systems and performance tuning xSeries servers, refer to the redbook *Tuning IBM @server xSeries Servers for Performance*, SG24-5287.

3.8 Rack installation considerations

The maximum number of x445 servers that can be effectively used in a 42U rack is determined by which IBM or vendor power distribution units (PDUs) the customer installs in the rack, how fully the x445 servers are configured, and how the line cords are plugged into the PDUs. Customers must therefore plan for rack power distribution.

Because the x445 chassis is 27.5” deep, it can be used in all NetBAY racks except the NetBAY3/3E (which can hold only a 3U-or-smaller chassis). However, due to increased airflow requirements, racks with perforated doors are required, or the front glass door (if installed) must be removed to allow proper airflow.

It is recommended that the x445 be installed in a 19-inch rack cabinet designed for 28-inch deep devices, such as the NetBAY42 Enterprise Rack (ER), NetBAY42 Standard Rack (SR), or NetBAY25 SR. Tower-like or mini-rack xSeries 445 installations can be created with the NetBAY11 Standard Rack Cabinet. The Enterprise Rack is deeper and therefore provides more room for cable management.

The 8 m Remote I/O cable and interconnect management cable allow the RXE-100 unit to be installed in racks adjacent to the one containing the x445 chassis.

If you use a non-IBM rack, the cabinet must meet the EIA-310-D standards with a depth of at least 28 inches. Also, adequate space (approximately two inches for the front bezel and one inch for air flow) must be maintained from the slide assembly to the front door of the rack cabinet to allow sufficient space for the door to close and provide adequate air flow.

Make sure all the cables attached to the x445 are long enough to permit the server to be slid out of the rack. This would include the normal cables such as power, network, and fiber cables, but also includes the Remote I/O cable for connecting to the RXE-100 and the SMP expansion cables for connecting to another x445.

Since the x445 is rack optimized, the IBM xSeries rack configurator should be used to ensure correct placement. The configurator can be downloaded from:

- For EMEA: http://www.pc.ibm.com/europe/configurators/
- For USA: http://www.pc.ibm.com/us/eserver/xseries/library/configtools.html
- For other countries or regions:

  a. Go to http://www.ibm.com
  b. Click **Select a country**
  c. Select your country

d. Click **Products and Services** at the top of the page
e. Under Servers, click **Intel-based**
f. Click **Tools** from the left navigation area
g. Scroll down to find the Rack Configurator section

3.9 Power considerations

The x445 ships with two redundant, hot-swappable power supplies that produce 1050 W each at 220 V, or 550 W each at 110 V. When the x445 is populated with more than two processors, the power supplies may not be redundant if they are connected to a 110 V power source. Therefore, IBM recommends that the x440 be connected to a 220 V power source to ensure power supply redundancy for all possible configurations.

Tip: If power is not redundant, the Nonredundant LED will be lit in the light path diagnostic panel (see Figure 1-12 on page 25).

Two system power-cord connectors are available on the back of the x445, one for each of the power supplies. Connect each of these power connectors to separate power circuits to ensure availability if one circuit should fail.

The x445 ships with two 2.8 m/9 ft. IEC 320-C13 to IEC 320-C14 power cables for intra-rack power distribution. Models shipped in the US also include two 2.8 m/9 ft. IEC 320-C13 to NEMA 6-15P power cords for attachment to high-voltage power sources.

3.10 Solution Assurance Review

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. For example, simpler solutions may need only a self review. However, the customer environment risk combined with the complexity of the solution may require than an expert level SAR take place, facilitated by a Quality Assurance practitioner and supported by a team of technical experts.

The three levels of Solution Assurance are:

- **Self**. In a self review, the solution designer checks his/her own work using a general or product-specific checklist, as appropriate.
- **Peer**. A peer review is conducted by a peer or colleague of the solution design team using available checklists, general and/or product-specific. The solution designer explains the solution to the peer reviewers, who assess the technical

viability of the solution and record their action items and comments on the checklist.

- **Expert**. During an expert review, the solution design team explains the solution to one or more subject matter experts. A Solution Assurance Quality Assurer facilitates the expert review and records the results in the Solution Assurance database used by that geography.

If an xSeries x445 is to be included as part of a solution, completion of an *expert level* Solution Assurance Review is required before you present the associated proposal to the customer.

Note: The EXAct Business Partners themselves may be required to perform the Solution Assurance Review.

The Solution Assurance Web sites are:

- 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 further information about what is required, refer to the following documents in the Solution Assurance Web sites (listed above) for more information in your specific geography:

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

4

Windows Server 2003, Datacenter Edition

This chapter describes how Microsoft Datacenter 2003 and the x445 are an ideal match. 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:

4.1 Description

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

Table 3-8 on page 99 lists a summary of the differences between the family members.

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.

4.1.1 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/windowsserver2003/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 Window 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 4.4, “Clustering” on page 127. 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 *Implemeting 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/TCPIP03.asp

4.1.2 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 link directly with this URL (from the online PDF of this redbook, the link is active):

http://www.microsoft.com/windows/catalog/server/default.aspx?subID=22&xslt=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

4.1.3 Comparing Datacenter to Enterprise and Standard Editions

Table 4-1 on page 123 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 4-1 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

Features	Datacenter Edition	Enterprise Edition	Standard Edition	Web Edition
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
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				

Features	Datacenter Edition	Enterprise Edition	Standard Edition	Web Edition
Windows Media Services	Yes	Yes	Yes	No

4.2 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.

4.3 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.

4.4 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.

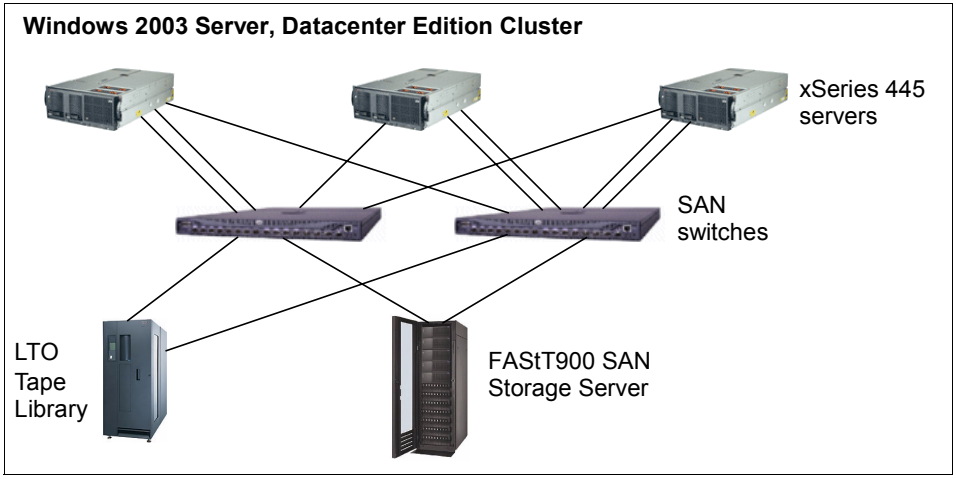

Figure 4-1 Typical Windows Server 2003, Datacenter Edition cluster

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 FAStT900 SAN Storage Server (or other certified SAN storage servers) and to the LTO tape library.

The FAStT900 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.

4.5 Datacenter Edition and server consolidation

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 4-2 on page 130 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.

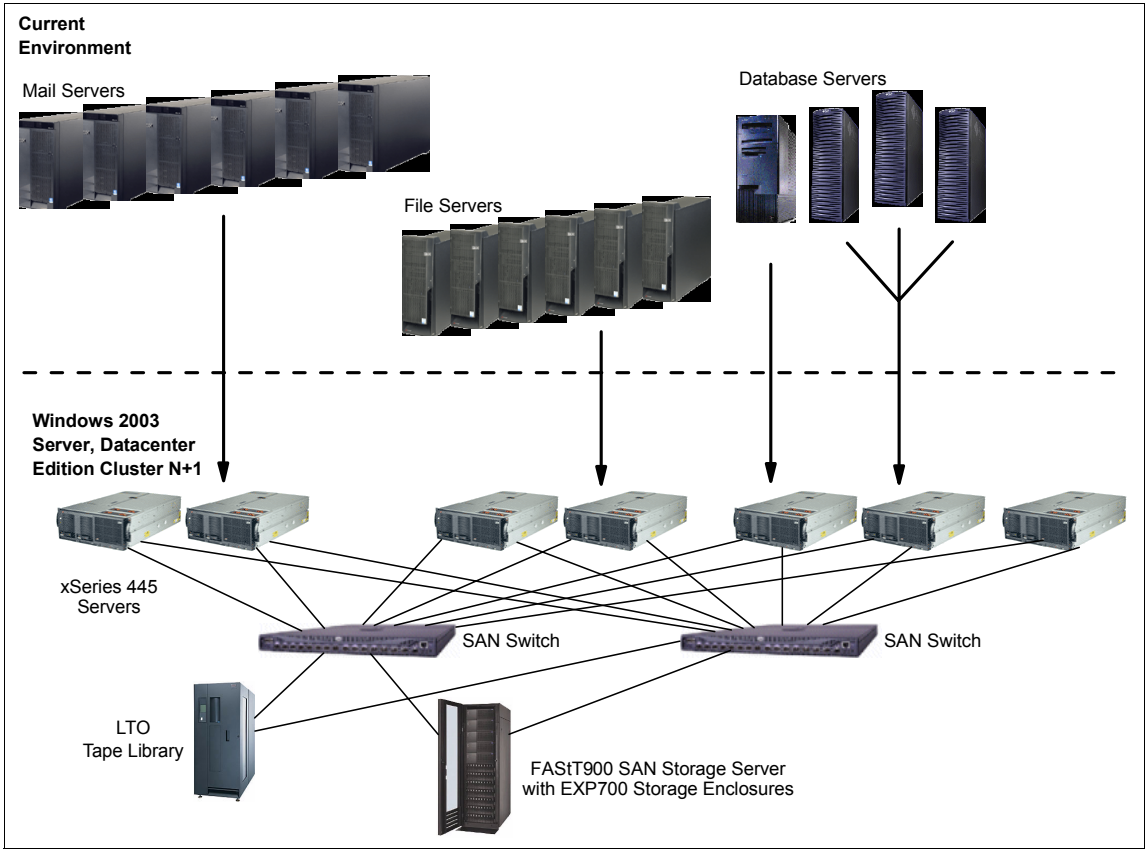

Figure 4-2 Before and after view of server consolidation and Datacenter Edition clustering implemented

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.

4.6 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.

4.6.1 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 4-2 Test comparison between 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	► 14 day no-fail test ► DC Driver Program
Kernel mode software (applications and drivers)	14-day no-fail test	► 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	► 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.

4.6.2 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/

4.6.3 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.

4.6.4 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 "Subscriptions" on page 97) 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

4.7 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.

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

4.8 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 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 added 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.

4.9 NUMA support in Windows Server 2003

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.

Affinitize means link or tie. The operating system ensures that processes are run on specific CPUs.

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.

4.9.1 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.mspx

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.

4.10 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 Lotus Notes®, on server D01DBM20, database x_dir\u20a\xserie1.nsf.

4.11 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/MIGR-4CMSC7.html

- 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 , "HARQ and the JSQ" on page 145, 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.

  See "Subscriptions" on page 97 for details on what to order.

- 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).

  See "Subscriptions" on page 97 for details on what to order.

- 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.

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.mspx

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, therefore 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.

5

Installation

In this chapter we describe procedures specific to the installation of Windows Server 2003, Windows 2000, Linux, Novell NetWare, and VMware operating systems on the x445 platform. The topics covered are:

5.1 System BIOS settings

This section describes system BIOS settings that you may need to configure prior to installing an operating system.

5.1.1 Updating BIOS and firmware

We recommend you check the BIOS and firmware levels on the items listed below and update to the most current revision, as part of your installation procedure for the x445:

- System BIOS
- Remote Supervisor Adapter firmware
- Onboard diagnostics
- Additional devices if installed, such as ServeRAID adapters and FAStT Fibre Channel host adapters

The latest BIOS and firmware code can be found at the driver matrix page for the x445:

http://www.ibm.com/pc/support/site.wss/MIGR-52102.html

Follow the installation instructions provided with each package.

5.1.2 Memory mirroring, hot-add, and hot-swap memory settings

Memory mirroring, part of IBM's Active Memory technology, provides an additional level of fault tolerance to the memory subsystem.

- 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 hot-swap memory feature allows you to replace failed DIMMs without turning off the server. Hot-swap memory is operating system independent.

For detailed information and guidelines on Active memory, please see “Active memory considerations” on page 62.

To enable these features on the x445, perform the following steps:

1. Press F1 when prompted during system startup to enter the System Configuration Utility.
2. From the main menu select **Advanced Setup** → **Memory Settings**. Figure 5-1 appears.

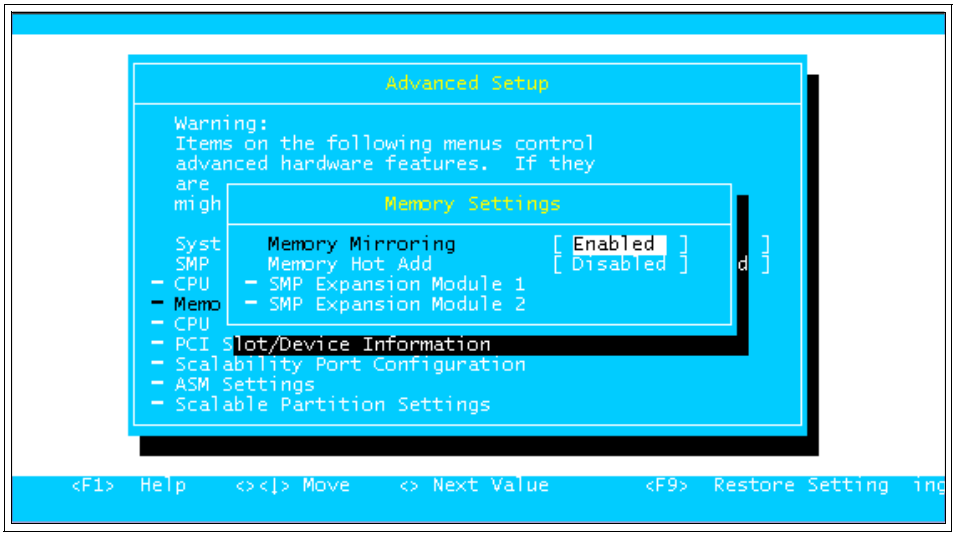

Figure 5-1 Memory settings

3. Select Memory Mirroring or Memory Hot Add, and press the right arrow key to change the value to **Enabled**.
   - To enable memory mirroring, set Memory Mirroring to **Enabled** and Memory Hot Add to **Disabled**.
   - To enable hot-swap memory, set Memory Mirroring to **Enabled** and Memory Hot Add to **Disabled**.
   - To enable hot-add memory, set Memory Mirroring to **Disabled** and Memory Hot Add to **Enabled**.

   Note: To use the hot-add memory feature memory mirroring must be disabled. To use the hot-swap memory feature memory mirroring must be enabled.

4. Exit the System Configuration Utility and save any changes.
5. If you enabled memory mirroring, you will see the following messages during POST:

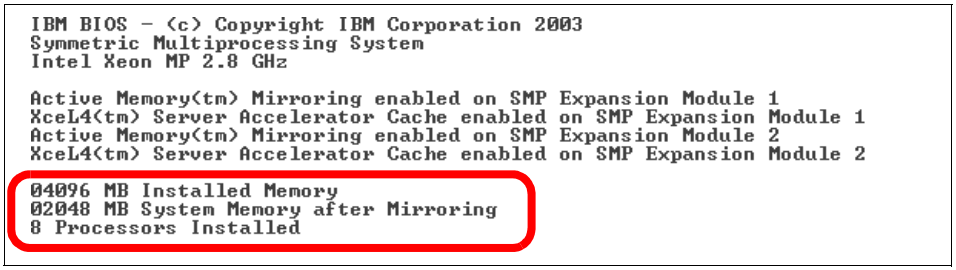

Figure 5-2 POST screen after enable Memory Mirroring

5.1.3 Enabling the serial port

If you want to install the serial port, as described in 3.3.4, “Serial connectivity” on page 84, you will also need to enable it in BIOS. This is achieved by changing the Serial port to Enabled in the Devices and I/O Ports menu item.

5.1.4 Hyper-Threading settings

Hyper-Threading technology allows a single processor to execute two separate instruction threads concurrently, effectively operating as two separate logical processors. The installed operating system sees these logical processors as two separate physical processors. This is demonstrated in Figure 5-9 on page 160.

There are a number of important factors to consider before enabling Hyper-Threading. These considerations are discussed in detail in 3.6, “Operating system considerations” on page 92. It is important to fully understand the configuration rules, performance issues and potential licensing implications related to Hyper-Threading before you proceed.

Hyper-Threading technology is disabled by default on x445, to enable it, do the following:

1. Press F1 during system startup to enter the System Configuration Utility.
2. From the main menu, select **Advanced Setup** → **CPU Options**. Figure 5-3 appears.

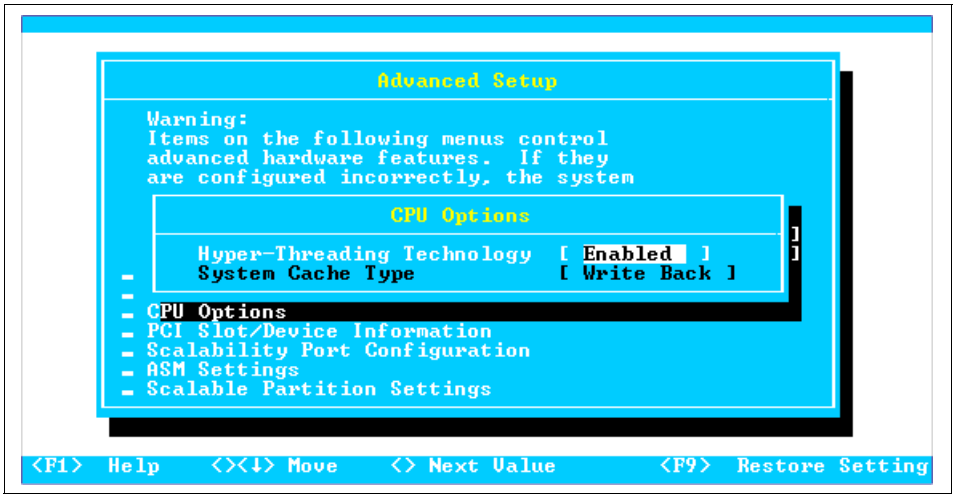

Figure 5-3 Hyper-Threading setting

3. Press the right arrow key to change the value to **Enabled**.
4. Save changes and exit System Configuration Utility.

5.2 Create RAID-1 using integrated LSI SCSI controller

Complete the following steps to start the LSI Logic Configuration Utility program and create a mirrored pair of SCSI hard disk drives:

1. Turn on the server.
2. When the following prompt appears, press Ctrl-C

   ```
   Press CTRL-C to start LSI Logic Configuration Utility
   ```

3. If you have set an administrator password, you are prompted to type the password.
4. Use the arrow keys to select the first channel from the list of SCSI channels; then, press Enter.
5. Select **Mirroring Properties**, and press Enter.
6. Use the arrow keys to highlight the first drive in the pair; then, press the Minus key to change the mirroring value to **Primary**.
7. Use the arrow keys to highlight the second drive in the pair; then, press the Minus key to change the mirroring value to **Secondary**.

8. When you have finished mirroring settings, press Esc to exit from the mirroring properties menu; select **Save changes then exit this menu** and press Enter to save the settings that you have changed.
9. Press Esc and select **Exit the Configuration Utility** to exit LSI Logic Configuration Utility program.

5.3 Device drivers

Specific device drivers are available for the x445 that may not be included with the base operating system.

Drivers that you should obtain separately from the operating system include:

- Hardware Abstract Layer (HAL) Support for Microsoft Windows NT 4.0, Windows 2000 and Windows Server 2003.
- Broadcom NetXtreme Ethernet controller
- Active PCI device
- Remote Supervisor Adapter management processor
- ATI Rage XL video controller
- LSI Ultra320 SCSI controller
- IBM ServeRAID

The required drivers as well as the latest versions of BIOS, service processor firmware, and diagnostics firmware are available from:

`http://www.ibm.com/pc/support/site.wss/MIGR-52102.html`

If you are implementing Microsoft Cluster Service (MSCS), check the Microsoft Cluster Service Hardware Compatibility List (HCL) to confirm that hardware, drivers and firmware have been Microsoft Certified. The HCL can be found at:

`http://www.microsoft.com/whdc/hcl/default.mspx`

5.4 Operating system installation

This topic will discuss the installation of the following operating systems on the x445 hardware platform:

- 5.4.1, “Microsoft Windows Server 2003” on page 153. This section also includes Windows 2000 installation information.
- 5.4.2, “Red Hat Linux Advanced Server 2.1” on page 164

- 5.4.3, “SuSE Linux Enterprise Server 8.0” on page 169
- 5.4.4, “SCO UnixWare Release 7.1.3” on page 169
- 5.4.5, “Novell Netware 6.0” on page 169
- 5.4.6, “VMware ESX Server 2.0” on page 170

Tip: The latest x445 operating system installation instructions are available on the Web:

- Customers and business partners: http://www.pc.ibm.com/support
- IBM employees: http://nosinfo.raleigh.ibm.com

In the following discussions, we assume that disk subsystems, such as RAID arrays, for both operating system installations and data storage have been configured.

5.4.1 Microsoft Windows Server 2003

This section describes how to install Windows Server 2003. The instructions also largely apply to Windows 2000 as well.

Tip: The latest installation instructions can be found at:

http://www.ibm.com/pc/support/site.wss/MIGR-51935.html

The following activities should be performed before you begin the operating system installation.

- Update BIOS and firmware

  Before commencing the installation of Windows, ensure that BIOS and firmware levels have been updated to the latest levels. Refer to 5.1.1, “Updating BIOS and firmware” on page 148.

- Using RAID with a ServerRAID adapter or the onboard LSI controller

  If you plan to create a RAID array, you will need to configure that before you begin. For the onboard LSI controller, follow the instructions in 5.2, “Create RAID-1 using integrated LSI SCSI controller” on page 151. For ServeRAID, boot the ServeRAID CD to configure the adapter.

  In addition, if you plan to boot from a disk connected to a ServeRAID controller, you will need the device driver on diskette. 5.3, “Device drivers” on page 152.

  For Windows 2000 installations, you will also need a driver diskette for the onboard LSI controller.

- x445 custom hardware abstraction layer software (HAL)

  Download the HAL and copy it to diskette. See 5.3, “Device drivers” on page 152 for instructions.

  You will need to install the custom Hardware Abstraction Layer (HAL) on your x445 servers regardless of the quantity of processors or SMP Expansion Modules installed. This HAL exploits the xSeries architecture: It allows for server flexibility and for the server to operate in various configurations. HAL also allows for an increased number maximum total interrupt inputs from all APICs.

- IBM Remote Expansion Enclosure 100 (RXE-100)

  We recommend you connect the RXE-100 to the x445 before installing the operating system.

  You can also connect the RXE-100 after installation and the operating system will detect and configure the new PCI-X slots. See “Attaching an IBM Remote Expansion Enclosure (RXE-100)” on page 158 for instructions.

- Hyper-Threading Technology

  Hyper-Threading is disabled by default on the x445. If you plan to enable it, you should do so before the starting the installation process. To enable Hyper-Threading, refer to 5.1.4, “Hyper-Threading settings” on page 150 for instructions.

  Note: Hyper-Threading only offers performance improvements when the server has two or four processors installed. There is little or no value in enabling Hyper-Threading on servers with more processors.

Installing the operating system

Tip: The latest installation instructions can be found at:

`http://www.ibm.com/pc/support/site.wss/MIGR-51935.html`

1. Turn on the server. Insert the operating system CD into the CD-ROM drive, then, restart server.
2. When prompted, press any key to boot from the CD.
3. When you are prompted with the following message, do one or both of the following.

   ```
   Setup is inspecting your server's hardware configuration
   ```

   – Press F5 to specify that you want to install a custom HAL

– Press F6 if you plan to install the operating system on drives connected to a ServeRAID adapter. For Windows 2000, you will need to press F6 to specify the driver for the LSI controller as well.

4. As a result of pressing F5, you will see the screen shown in Figure 5-4. Select **Other** from list, and press Enter.

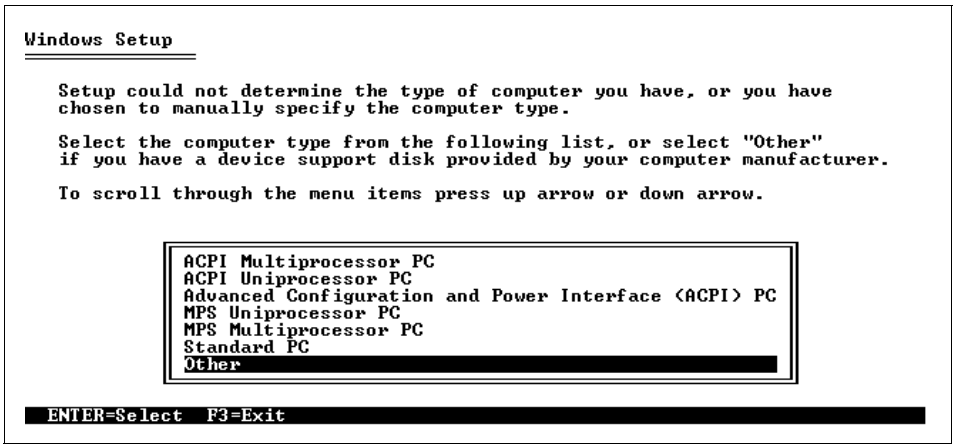

Figure 5-4 Selecting computer type

5. Insert the diskette that contains the HAL and press Enter.
6. Select the **IBM eServer xSeries 440/445 (Windows 2003 HAL)** option as shown in Figure 5-5, and press Enter to install.

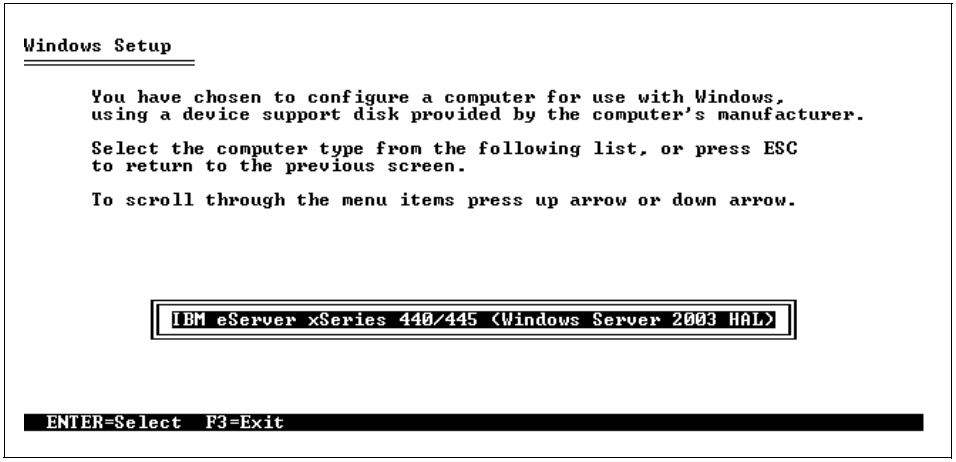

Figure 5-5 x445 HAL for Windows Server 2003

7. As a result of pressing F6, The Setup window is displayed where you can specify a disk controller. Select **S** to specify additional devices.

   If you are booting from a ServeRAID-attached disk, load the ServeRAID driver.

   For Windows 2000, you will also need to load the driver for the onboard LSI RAID controller.

8. Continue the operating system installation as normal.

Installing and updating device drivers

After Windows is installed, Device Manager will report a number of unknown devices, as shown in Figure 5-6. In this section, we describe how to install and update drivers for these unknown devices.

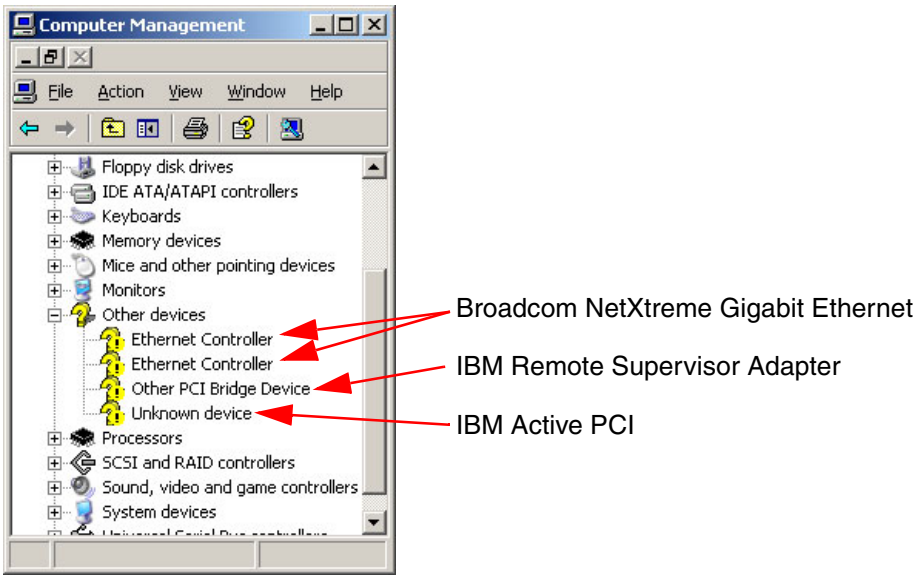

Figure 5-6 Device Manager before updating device drivers (Windows Server 2003)

For each of the missing drivers, download the drivers and run each EXE you download to unpack them. See 5.3, “Device drivers” on page 152.

Follow these steps to install the drivers:

- Active PCI driver

  a. In Device Manager, right click on **Unknown device** (see Figure 5-6) and click **Reinstall Driver**.

  b. Select **Install from a list or specific location** and click Next.

c. Specify the directory where you unpacked the driver to.

d. Follow the instructions to install the driver.

- Remote Supervisor Adapter driver

  a. Run setup.exe from the WINDOWS2000 directory where you unpacked the driver to.

  b. Follow the instructions to install the driver.

- Broadcom Gigabit Ethernet driver

  a. In Device Manager, right click on either **Ethernet Controller** entry (see Figure 5-6) and click **Reinstall Driver**.

  b. Select **Install from a list or specific location** and click Next.

  c. Specify the directory where you unpacked the driver to.

  d. Follow the instructions to install the driver.

  e. If necessary, repeat for the other device listed in Device Manager

Figure 5-7 shows the Device Manager after these three drivers have been installed.

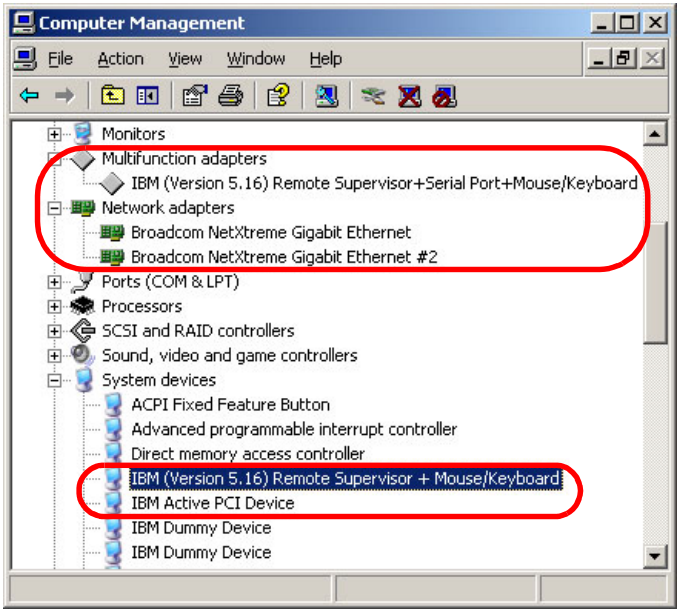

Figure 5-7 New devices detected after installing the drivers (Windows Server 2003)

Attaching an IBM Remote Expansion Enclosure (RXE-100)

If you chose to connect the RXE-100 to the server after installing the operating system, complete the following steps:

1. Turn off the server.
2. Refer to 3.3.3, “RXE-100 connectivity” on page 76 for instructions of connect RXE-100 to x445 server.
3. If the RXE-100 power cord is plugged in, unplug the power cord, wait 10 seconds, and plug it back in.
4. Connect the RS485 management cable from the RXE-100 to the server.
5. Turn on the server. If the RXE-100 is not powered on, repeat steps 3 and 4.
6. After Windows Server 2003 has booted, the system will automatically install the device driver for RXE-100 and you will see more PCI bus and PCI standard host CPU bridge devices under System devices in Device Manager as shown in Figure 5-8.

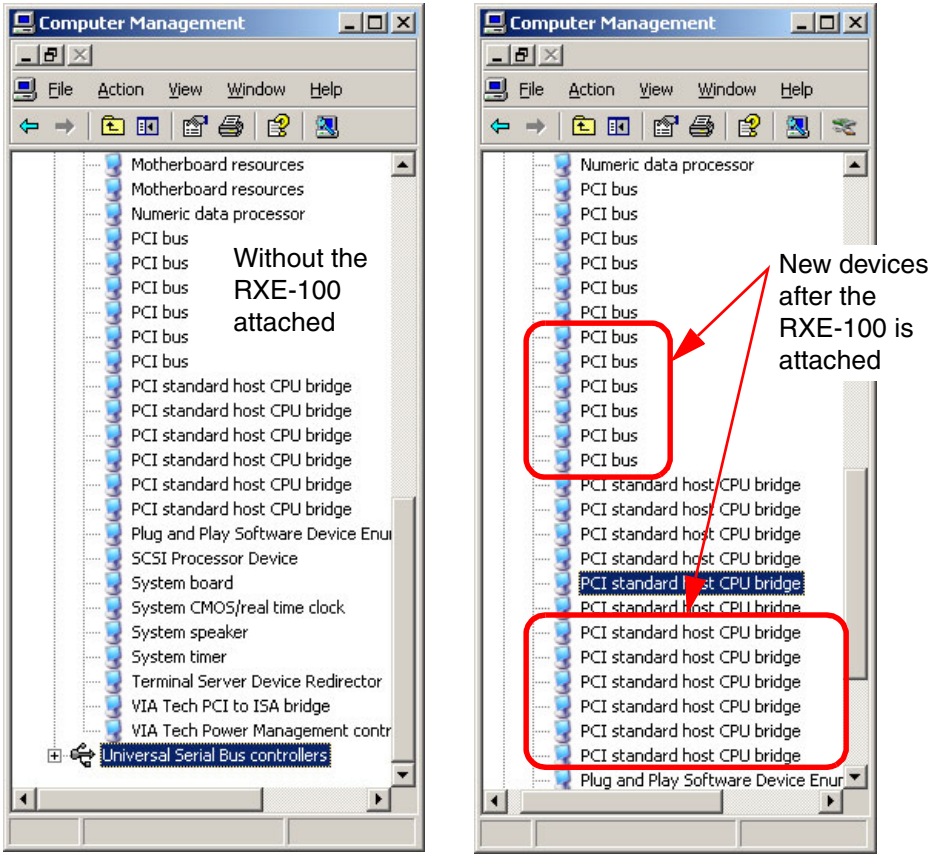

Figure 5-8 Device Manager with and without the RXE-100 attached

Hyper-Threading

Refer to 3.6, “Operating system considerations” on page 92 for a discussion of the performance and planning issues related to Hyper-Threading technology.

Once you enable Hyper-Threading, the Performance tab in Task Manager will show twice the number of CPUs installed as there are physical processors. Figure 5-9 shows the windows before and after enabling Hyper-Threading on a 8-way x445.

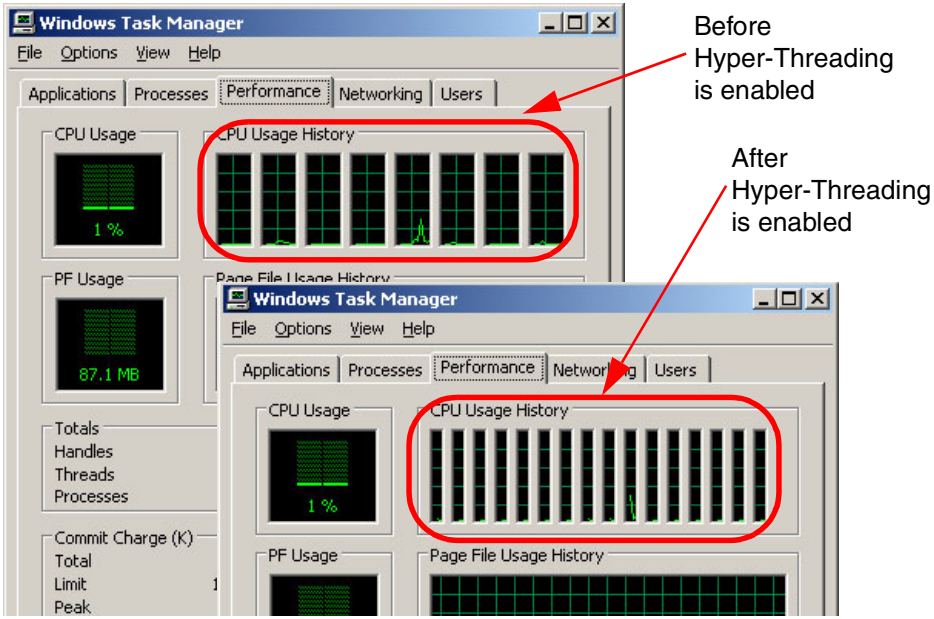

Figure 5-9 Before and after enabling Hyper-Threading on a 8-way x445 server

Server with more than 4GB of memory

By default, Windows Server 2003, Enterprise Edition and Datacenter Edition will only use up to 4 GB of memory. To increase the amount of memory the server can use, PAE (Physical Address Extension) must be added into the BOOT.INI file.

Figure 5-10 shows the memory size in System Properties window before enabling PAE option in the x445 with 16 GB memory installed.

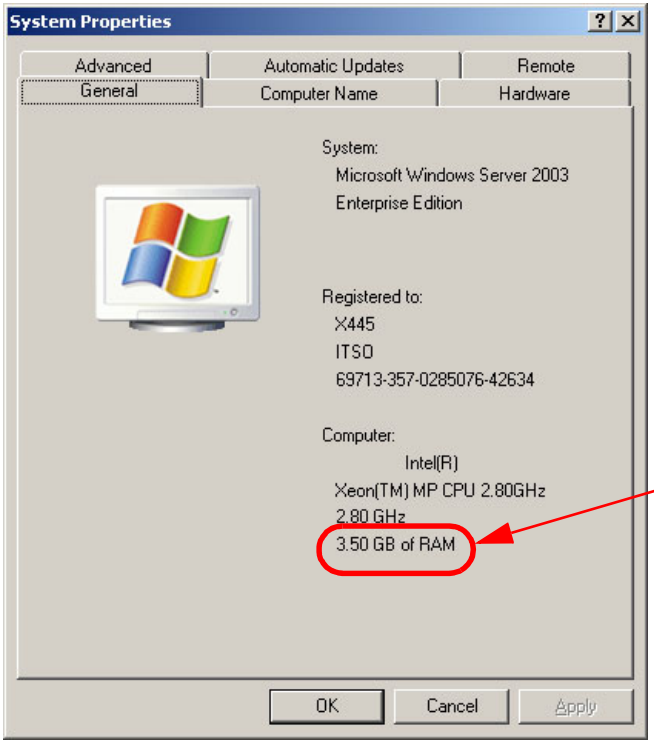

Figure 5-10 Memory size in System Properties before enable PAE option

The BOOT.INI file is a hidden operating system file. To view and modify the BOOT.INI file, complete the following procedures:

1. Open Explorer and click **Tools** → **Folder Options**.
2. Click the **View** tab.
3. Under Hidden files and folders select **Show hidden files and folders**.
4. Deselect **Hide protected operating system files,** click **OK**.
5. The BOOT.INI file should now be visible on the C: drive.

To add PAE to the BOOT.INI file, complete the following steps:

1. Double-click the BOOT.INI file.
2. Type the switch `/PAE` at the end of `Windows Server 2003, Enterprise` as shown in Example 5-1.
3. Save the BOOT.INI file.
4. Reboot the server.

Example 5-1 BOOT.INI

```
[boot loader]
timeout=30
default=multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
[operating systems]
multi(0)disk(0)rdisk(0)partition(1)\WINDOWS="Windows Server 2003, Enterprise" /PAE /fastdetect
```

Figure 5-11 shows the memory size of system after enable PAE option in BOOT.INI file.

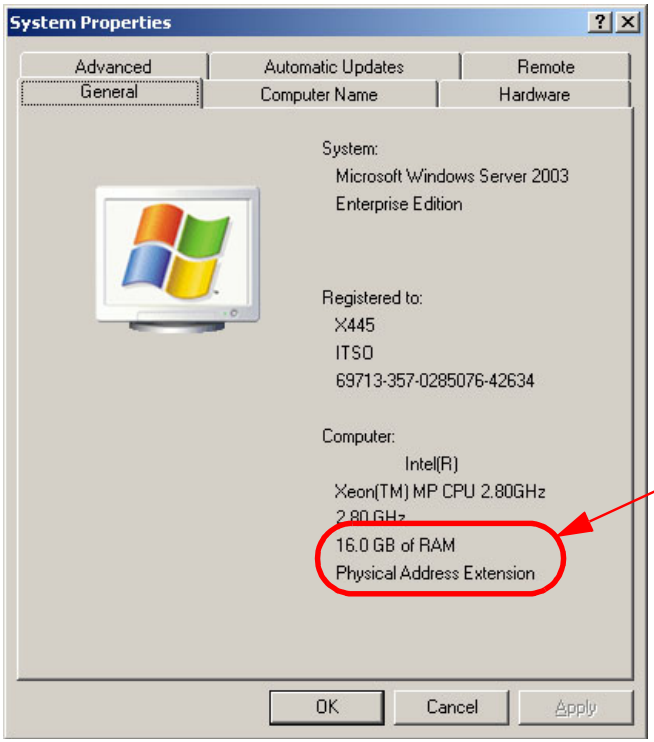

Figure 5-11 Memory size in System Properties after enable PAE option

Hot-adding memory in Windows Server 2003

This section describes how to hot-add DIMMs to x445 server running Windows Server 2003, Enterprise Edition and Datacenter Edition.

Before you begin, read 3.1.2, “Memory” on page 60 about the Active Memory considerations and memory configuration rules. In particular, Memory Mirroring must be disabled, Hot-Add enabled, and one of the memory ports must be empty.

You should also read the safety information in the x445 *Options Installation Guide* with respect to inserting memory while the server is still running.

Complete the following steps to hot-add DIMMs to x445 server SMP Expansion Module:

1. Open the top cover of x445 server.
2. Open the DIMM access door of the empty memory port on the SMP Expansion Module.
3. Wait until the Memory Port Power LED is turned off, then continue with following steps. *Do not* add any DIMMs if the Memory Port Power LED is on.
4. Add new DIMMs. Please follow the memory configuration rules that described in 3.1.2, “Memory” on page 60 to install the new DIMMs. DIMMs must be installed in matched pairs.
5. Close the DIMM access door. The Memory Port Power LED will turn on again.
6. Close the top cover of x445.

To verify success, you can open the System Properties window (right click on **My Computer**) to check the change of memory size as shown in Figure 5-12.

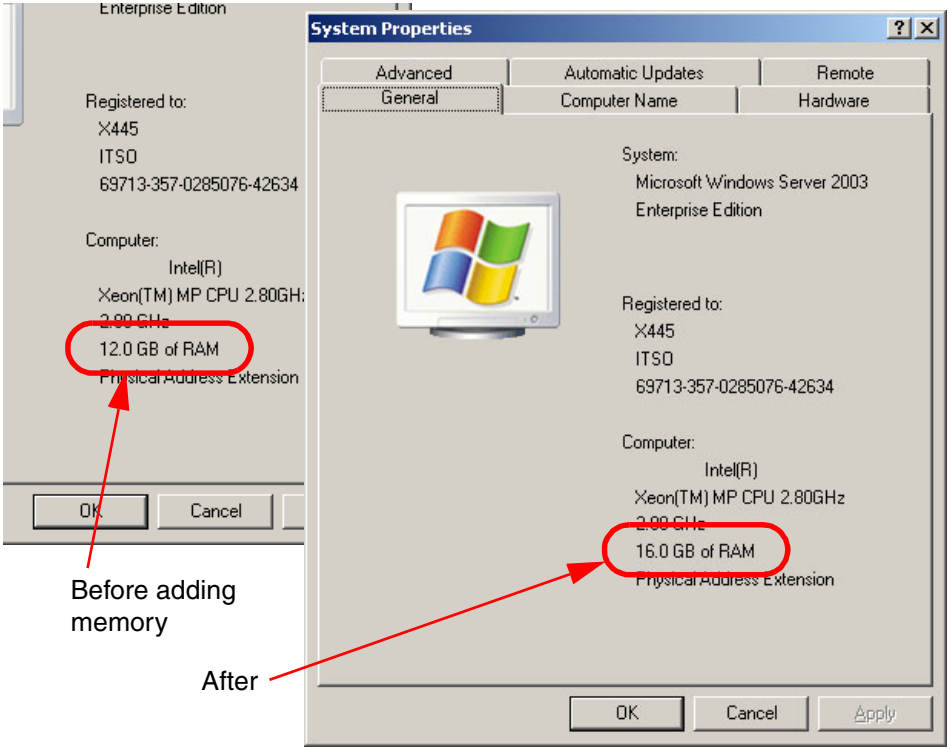

Figure 5-12 Hot-adding memory on Windows Server 2003 with the server power on

5.4.2 Red Hat Linux Advanced Server 2.1

In this section we discuss the installation of Red Hat Enterprise Linux Advanced Server 2.1 on the x445 server.

Tip: For the latest Red Hat installation instructions, go to:

`http://www.pc.ibm.com/support`

Select the x445 (machine type 8870), then select document type **OS Installation**.

Before you begin

To install Red Hat Enterprise Linux AS 2.1 on x445 server, you will need following:

- A registered copy of Red Hat Enterprise Linux AS 2.1.
- One blank diskette to create a boot disk.

- LSI SCSI driver for Red Hat Enterprise Linux AS 2.1 on diskette.
- If you are using the a ServeRAID Adapter, you will need the IBM ServeRAID Support CD Version 5.10 or later.

Refer to 5.3, “Device drivers” on page 152 for information about where to download the device drivers.

Before installing Red Hat Linux, update firmware and BIOS levels as discussed in 5.1.1, “Updating BIOS and firmware” on page 148.

Installing the operating system

Complete the following steps to install Red Hat Enterprise Linux AS Version 2.1 on x445. See “Installing Red Hat Enterprise Linux AS Version 2.1 Update 2” on page 168 if you are using Update 2 installation CDs.

1. Power on the server and insert Red Hat Enterprise Linux AS (v.2.1) CD 1.
2. In the Welcome to Red Hat Enterprise Linux AS Version 2.1 window, at boot prompt enter:

   ```
   linux dd
   ```

3. At the Devices window, insert the LSI SCSI driver diskette, select **Yes** and then, select **OK**.
4. Following the prompt on screen to select language, keyboard, mouse and select **Advanced Server** to install in the Installation Type window.
5. In the Package Group Selection window, select **Software Development** and Select **individual packages** and click **Next**, Select **Development** → **System** → **kernel-source package**, then click **Next**.
6. When you have completed the installation and server restarts, select **2.4.9-e.3enterprise** kernel in the GRUB menu to boot, and do the following steps to install the summit kernel (2.4.9-e.12 or later).

   You can get the 2.4.9-e.24 summit kernel from Red Hat Enterprise Linux AS Version 2.1 Update 2 CD 1 or the latest version of summit kernel from the Red Hat Network site (http://rhn.redhat.com).

   Tip: The summit kernel shipped on the V2.1 CD is not suitable for the x445 and will cause a kernel panic if used. Instead, use the summit kernel from the Update 2 CD, or download it from the Red Hat site listed above.

   To install the summit kernel from the Update 2 CD, do the following:

   a. Log in as root

b. Insert the Red Hat Enterprise Linux AS Version 2.1 Update 2 CD 1 and mount the CD (if it is not auto-mounted) with the following command:

```
mount /mnt/cdrom
```

c. Install the summit kernel RPM with the following command:

```
rpm –ivh /mnt/cdrom/RedHat/RPMS/kernel-summit-2.4.9-e.24.i686.rpm
```

d. Restart the server and select the **2.4.9-e.24summit** kernel in the GRUB menu to boot.

e. During the summit kernel booting, if the Welcome to Kudzu window appears, press any key to continue; then the system will report find a tg3 device on the “Hardware Added” window. Select **Ignore** and press Enter.

Note: If you experience a system hang when booting the enterprise kernel to install the summit kernel, reset the x445 server and press F1 to enter the Configuration/Setup Utility. In the **Advanced Options** and **CPU Options** menu, enable Hyper-Threading. Save changes, exit the Utility and restart the server.

Installing Broadcom Ethernet controller device driver

Complete the following steps to install the Broadcom device driver:

1. Ensure the summit kernel is installed, as described in “Installing kernel source for summit kernel” on page 167.
2. If the tg3 module is loaded, using following commands to remove the tg3 module:

```
ifdown eth0
ifdown eth1
rmmod tg3
```

3. Download the Broadcom device driver version using the TAR file. Refer to 5.3, “Device drivers” on page 152.
4. Create a directory and extract the files by type following command (we use Broadcom device driver Version 6.2.11 as sample):

```
tar xvzf bcm5700-6.2.11.tar.gz
```

5. Build the driver `bcm5700.o` as a loadable module for the running kernel:

```
cd src
make
```

6. Load the driver:

```
insmod bcm5700.o
```

7. Install the driver and man page:

   ```
   make install
   ```

8. Edit the `/etc/modules.conf` file. If there is an alias entry in your `/etc/modules.conf` file referencing the `tg3` driver, make sure that you replace `tg3` with `bcm5700`. Otherwise, add the entry below if necessary.

   ```
   alias eth0 bcm5700
   alias eth1 bcm5700
   ```

9. Configure the network protocol and address for Broadcom Ethernet controller.

Installing kernel source for summit kernel

If you are compiling updated drivers on the summit kernel for the x445 server, you will need to install the kernel source and apply the summit kernel patch before you compile the driver.

If you updated the summit kernel from Red Hat Network (RHN), use the following instructions with the appropriate kernel version. For example, the 2.4.9-e.24 smp kernel would be changed to 2.4.9-e.25 for the 2.4.9-e.25 summit kernel.

1. Install the kernel-source RPM from either the Red Hat Enterprise Linux AS Version 2.1 Update 2 CD 2 or the Red Hat Network (http://rhn.redhat.com) by running the command:

   ```
   rpm -ivh kernel-source-2.4.9-e.24.i386.rpm
   ```

2. Install the kernel SRPM from either the Red Hat Enterprise Linux AS Version 2.1 Update 2 CD 4 or the Red Hat Network by running the command:

   ```
   rpm -ivh kernel-2.4.9-e.24.src.rpm
   ```

3. Run the following commands:

   ```
   cd  /usr/src/linux-2.4.9-e.24
   patch -p1 < /usr/src/redhat/SOURCES/linux-2.4.9e12_summit-2003-03-14.patch
   patch -p1 < /usr/src/redhat/SOURCES/linux-2.4.9e12_686-notsc_A0.patch
   patch -p1 < /usr/src/redhat/SOURCES/linux-2.4.9e12_cyclone-timer_B2.patch
   patch -p1 < /usr/src/redhat/SOURCES/linux-2.4.9e12_lost-tick_A0.patch
   patch -p1 < /usr/src/redhat/SOURCES/linux-2.4.9e12_get_irq_list.patch
   patch -p1 < /usr/src/redhat/SOURCES/linux-2.4.9-summit-notsc-soft.patch
   cp -f configs/kernel-2.4.9-i686-summit.config .config
   ```

4. Edit the `Makefile` to change the `EXTRAVERSION` field from `custom` to `summit`:

   ```
   vi Makefile
   ```

   Now find the line that reads `EXTRAVERSION = -e.12custom` and change it to `EXTRAVERSION = -e.12summit`

5. Run **`make clean`**.

6. Run **`make menuconfig`**.

7. Select **Exit**, When prompted at the "`Do you wish to save your new kernel configuration?`" screen, select **Yes**.
8. Run **`make dep`**.
9. Compile the driver. See the documentation that comes with the driver for more information.

Installing Red Hat Enterprise Linux AS Version 2.1 Update 2

Complete the following steps to install Red Hat Enterprise Linux AS Version 2.1 Update 2:

1. Power on the x445 server and insert Red Hat Enterprise Linux AS (v2.1) Update 2 CD 1 in to CD-ROM drive.
2. In the Welcome to Red Hat Enterprise Linux 2.1 AS window, press Enter.
3. Following the normal installation process to select Language, keyboard, mouse and setup disk.
4. As Red Hat Enterprise Linux AS (v2.1) Update 2 installation CD has the Broadcom Tigon3 driver for Broadcom Ethernet controller, so the Network Configuration window will appear, configure the two network adapter and click **Next**. We will install Broadcom bcm5700 driver replace the Tigon3 driver for onboard Ethernet controller after we complete Red Hat Linux installation.
5. In the Package Group Selection window, select **Software Development** and Select **individual packages** and click **Next**, Select **Development** → **System** → **kernel-source package**, then, click **Next**.
6. When you have completed the installation and server restarts, select **2.4.9-e.24summit** kernel in the GRUB menu to boot, then continue with “Installing Broadcom Ethernet controller device driver” on page 166.

Hyper-Threading support under Linux

Hyper-Threading support is part of the summit kernel. Once the kernel has been installed and compiled, and you enable Hype-Threading support in the BIOS, Red Hat Enterprise Linux Advanced Server will recognize all your physical and virtual CPUs.You can view the processor information by running command:

```
cat /proc/cpuinfo
```

RXE-100 support under Linux

Red Hat Enterprise Linux AS 2.1 supports the full 12 slots of the RXE-100. Once the summit kernel is installed and configured, you have access to your RXE-100 and adapters added to the external enclosure.

5.4.3 SuSE Linux Enterprise Server 8.0

Obtain the latest installation instructions from:

http://www.ibm.com/pc/support/site.wss/MIGR-52209.html

Note: There is a known problem with SLES 8 Service Pack 2a on the x445, which causes a system hang. Do not install SP2a. The problem should be resolved with Service Pack 3.

Refer to 5.3, “Device drivers” on page 152 for information about where to download the device drivers.

Before installing SuSE Linux, update firmware and BIOS levels as discussed in 5.1.1, “Updating BIOS and firmware” on page 148.

5.4.4 SCO UnixWare Release 7.1.3

The latest installation instructions can be found at:

http://www.ibm.com/pc/support/site.wss/MIGR-52460.html

5.4.5 Novell Netware 6.0

The latest installation instructions can be found at:

http://www.ibm.com/pc/support/site.wss/MIGR-52367.html

You will need:

- The latest NetWare Overlay CD: http://support.novell.com/tools/csp/
- The latest LSI Logic SCSI driver: http://www.ibm.com/pc/support/site.wss/MIGR-44870.html
- The latest Broadcom NetXtreme Gigabit Ethernet driver: http://www.ibm.com/pc/support/site.wss/MIGR-52019.html

If you use any additional PCI or PCI-X adapters (for example, ServeRAID or Fibre Channel HBA), download the latest drivers for them as well. See 5.3, “Device drivers” on page 152.

In addition to these tasks, we recommend you to update the following components to the latest level:

- x445 BIOS, server diagnostics, and the RSA firmware
- BIOS and/or firmware of additional PCI or PCI-X adapters (such as ServeRAID)

See 5.1.1, “Updating BIOS and firmware” on page 148.

You are now ready to install the operating system. Start the server and insert the NetWare 6 SP 3 Overlay CD into the CD-ROM drive. Follow the instructions on the screen, provide information and make selections as appropriate. During the installation, use the LSI SCSI and the Broadcom Ethernet device drivers, plus the drivers for optional any other devices you are using.

Notes: You cannot use a USB keyboard or mouse during the installation, because the USB driver will not be activated at that point. Once NetWare is successfully installed and the USB driver started, you can use the USB devices.

At least one network adapter must be installed in the server before you can start the NetWare installation. Additional network adapters can be configured after the installation completes.

Hyper-Threading

The Hyper-Threading support has two prerequisites:

- You must enable it in BIOS. See 5.1.4, “Hyper-Threading settings” on page 150.
- You have to install the ACPI modules (ACPIDRV.PSM). Support Pack 3 will also be needed for this. Use the nwconfig utility install the ACPI support.

5.4.6 VMware ESX Server 2.0

In this section we will go through the process of installing and configuring the VMware ESX Server software on x445 server.

For more detail information about how to install ESX Server 2.0, how to create virtual machines, how to manage your virtual machines from a remote workstation and how to work with the advanced features of VMware ESX Server, refer to the product publications:

- *VMware ESX Server Installation Guide*
- *VMware ESX Server Administration Guide*

You can download those books from VMware Web site:

`http://www.vmware.com/support/esx2/doc/`

Before you begin

To install ESX Server 2.0 on x445, you need the following:

- A registered copy of VMware ESX server Version 2.0, which includes the VMware Service Console, VMware ESX Server software, and VMware Remote Console software.
- Two or more Ethernet controllers. Supported controllers include:
  - Broadcom NetXtreme 570x Gigabit controllers
  - Intel PRO/100 adapters
  - Intel PRO/1000 adapters
  - 3Com 9xx based adapters

  Note: Since there will be minimal network activity to the VMware management console, it may make sense to install a 10/100 adapter for use by the console, and allocate the onboard Broadcom Gigabit Ethernet controller to the virtual machines.

  If you want to do this, you must disable the onboard adapter prior to the installation of ESX Server.

  For information on supported hardware of ESX Server 2.0, download the *VMware ESX Server Hardware Compatibility Guide* from the VMware Web site at:

  `http://www.vmware.com/support/esx2`

  For a full list of adapters support in the x445, see:

  `http://www.pc.ibm.com/us/compat/lan/matrix.html`
- ESX Server should be installed on local hard disks. VMware does not support installing ESX Server on a hard disk that is on a SAN.
- Update BIOS and firmware levels as discussed in 5.1.1, “Updating BIOS and firmware” on page 148.

Using onboard Gigabit Ethernet controller for VMs

ESX Server will use the first network adapter it finds to be the network connection for the console operating system.

If you want to use the onboard Broadcom Gigabit Ethernet controllers in the x445 dedicated to the VMs and have an additional card (for example, a 10/100 Ethernet adapter) for the console operating system and management, you will

need to disable the onboard Ethernet, install ESX Server, and then later re-enable it.

1. Boot the x445 and disable the onboard Ethernet by pressing F1 during the startup as shown in Figure 5-13.

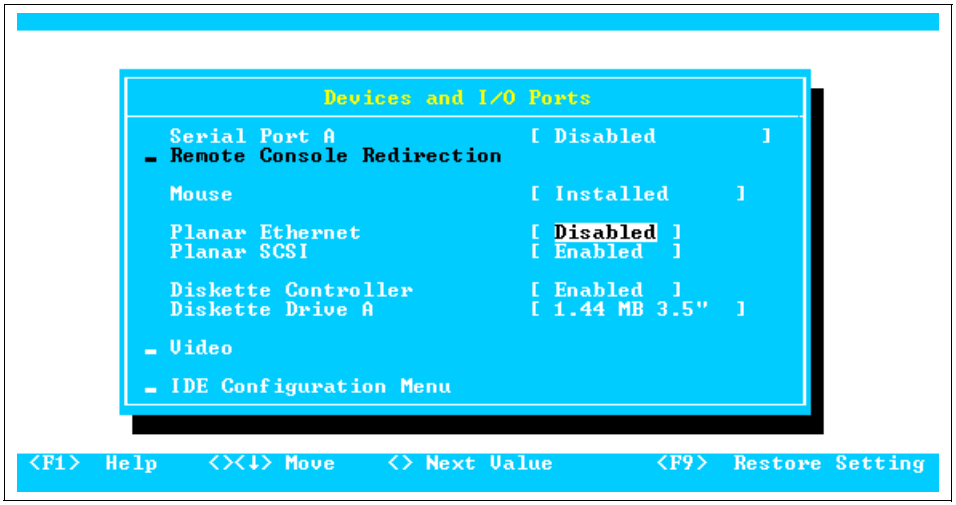

Figure 5-13 Disable Onboard Ethernet controller

2. Install an additional 10/100 Ethernet adapter.
3. Install ESX Server and configure the 10/100 adapter for use by the console.
4. Continue with the installation until you save the configuration of your VMkernel and, immediately before creating the VMs, reboot the system, enter the System Setup, and re-enable the onboard Ethernet controller.
5. As the system reboots, make sure that the current kernel is still loaded.
6. Log on to the console system and if you can do this it means you still are using the 10/100 Ethernet adapter you originally configured for the role.
7. Edit the current VMkernel configuration through the graphical interface and you can now move the onboard Ethernet card to the VMs.

Installing ESX Server 2.0

In this section, we describe how to install VMware ESX Server 2.0 on an x445. Before you begin, ensure that you have all the relevant information about your network:

- The host name for the server, including the full domain name, if applicable
- An IP address for the ESX Server and IP addresses for all the virtual machines you will be creating

- The subnet mask
- The IP address of the gateway
- The IP address of the name server

Note: We recommend you use a static IP for your VMware server.

Complete the following steps to install VMware ESX Server 2.0.

1. Make sure the network cable is plugged into the network adapter you plan to use with the console operating system (we recommend an additional 10/100 Ethernet adapter) before starting the installation.
2. Power on the x445 server and insert the VMware ESX server Version 2.0 CD into the CD-ROM drive. The boot process will begin.
3. At the Install Supplemental Drivers window, select **Continue** to proceed with the installation.
4. At the Welcome to VMware ESX Server window, select **Install**.
5. At the Disk Partitioning Setup window, select **Manual**, and press Enter.
6. At the Partitioning window, create the following suggested partitions:
   a. Boot partition with mount point **/boot**, file system type **ext3**, and **50 MB** in size.
   b. Swap partition at least twice the memory assigned to the service console. Recommended: **512 MB**.
   c. Root partition with mount point **/**, file system type **ext3**, and **1800 MB** in size.
7. At the Boot Loader Configuration window, select **Master Boot Record**, and press Enter.
8. At the Network Configuration window, configure your first network card and select the **OK** to continue.
9. At the Hostname Configuration screen, enter the host name of the server.
10. At the Time Zone Selection window, select the correct time zone and select the **OK** to continue.
11. At the Root Password window, type and confirm your root password, then select the **OK** to continue.
12. At the Add User window, add a user that will be able to login to the VMware Management interface, then select the **OK**.

    These accounts will not have administrative privileges for ESX Server, but will be able to create and administer VMs created with that user ID. You can add

as many accounts as required, and additional accounts may be created after the VMware installation is completed. In our example, we have created a Linux user who will have access to only Linux VMs (user ID of “linux”) and a Windows user (user ID of “windows”) who will only have access to Windows VMs.

13. At the User Account Setup window, add any additional users, then select the **OK** to start the actual install.
14. At the ESX Server Installation Complete screen, select **Reboot**, and press Enter. Remove the CDROM once ejected.
15. At the boot menu select **linux-up** to start the server.

When the system has rebooted, you will see the screen shown in Figure 5-14. ESX Server is now ready for configuration.

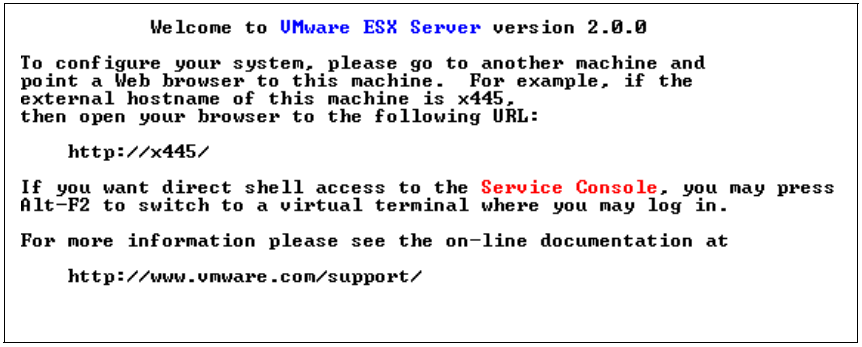

Figure 5-14 ESX Server installation completed

Accepting the VMware ESX Server 2.0 Certificate with IE

Complete the following steps to accept the VMware ESX Server 2.0 certificate with Internet Explorer.

1. Launch Internet Explorer and enter the IP address you assigned to the service console to load the VMware Management Interface.
2. At the Security Alert window, click on **View Certificate**.
3. At the Certificate window, click on **Install Certificate**.
4. At the Certificate Import Wizard window, select **Place all certificates in the following store**, then click **Browse**.
   a. Select the **Show physical stores** check box, then expand the **Trusted Root Certification Authorities** folder.
   b. Select **Local Computer** folder and click **OK**.
   c. Click **OK** at the import was successful message.

d. Click **OK** again to close the Certificate window, then click **Yes** to accept the certificate.

Configuring ESX Server

Complete the following steps to configure VMware ESX Server 2.0.

1. Launch Internet Explorer and enter the IP address you assigned to the service console to load the VMware Management Interface.
2. At the browser security certificate screen, follow the steps in , “Accepting the VMware ESX Server 2.0 Certificate with IE” on page 174 to accept the certificate.
3. Next, login to the VMware Management Interface using the root account you created during installation. A new window will open and the System Configuration Wizard will be displayed.
4. At the Welcome screen, click **Next** to start the wizard.
5. At the End User License Agreement screen as shown in Figure 5-15, check **I Accepted the terms in the license agreement** and enter a valid VMware ESX Server 2.0 serial number and if appropriate a valid VMware Virtual SMP for ESX Server serial number. Click **Next** to continue.

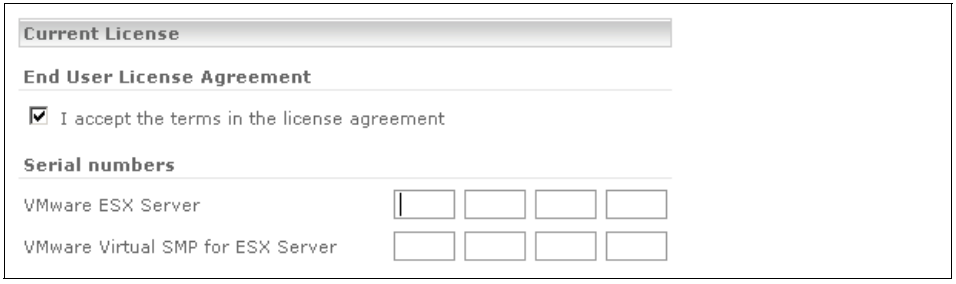

Figure 5-15 ESX Server license

6. At the Startup Profile screen as shown in Figure 5-16, configure the service console’s hardware as follows:
   a. Reserved memory — This is the amount of memory reserved for the service console.
      - 192 MB - 8 virtual machines
      - 272 MB - 16 virtual machines
      - 512 MB - 32 virtual machines
   b. SCSI storage controller — This is where you allocate SCSI controllers to be used by the service console or the virtual machines. Both the service console and the virtual machines need access to a device. When you are configuring the storage, you may find the user interface is somewhat

confusing. If you only have one hard drive, deselect **Service Console**, Select **Virtual Machines**, and then select **Shared with Service Console**.

c. Ethernet controllers — This is where you add Ethernet controllers to be used by the service console. Select dedicated to **Service Console** here.

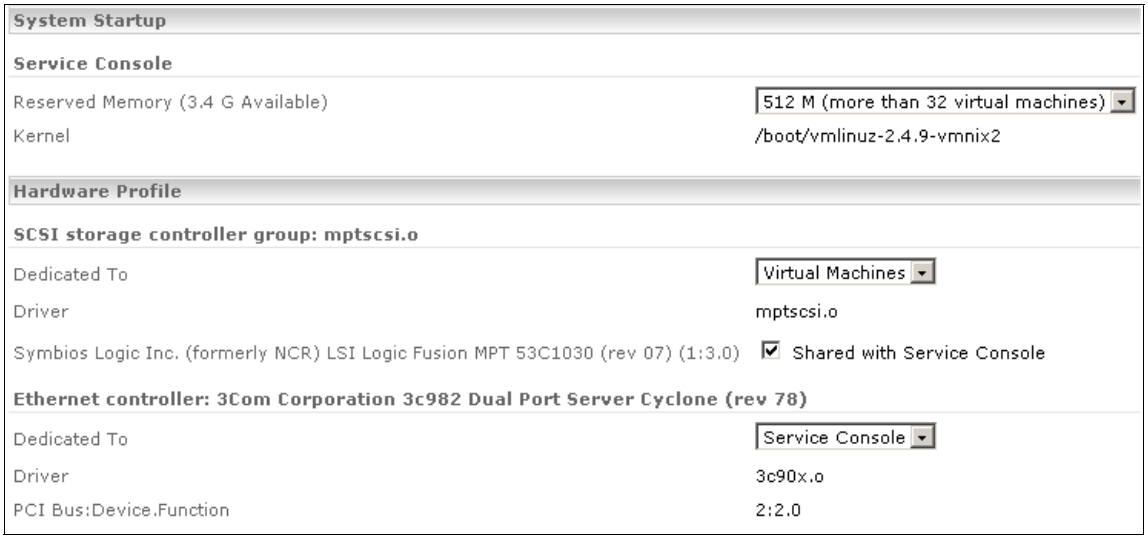

Figure 5-16 ESX Server 2.0 Startup Profile configuration

7. Click **Next** to continue, your system will be rebooted.
8. At the server's boot menu, select the **vmnix** kernel and press the Enter to start the server.
9. Once the server reboots completely, click **Next** on the VMware Management Interface's Reboot screen to continue (this may take up to 5 minutes before you get a response, do not click **Next** again).
10. At this point you may be prompted to login again, if so, login as root.
11. At the Storage Configuration screen, create a VMFS volume by clicking **Create Partition** next to freespace.
    a. When asked if you want to create a core dump partition, click **Yes**.
    b. A dialog box will be displayed asking you to give your VMFS volume a unique label, enter a name and click **OK**.
    c. Click **Next** once back at the Storage Configuration screen.

12. At the Swap Configuration screen, setup the swap space that ESX server uses as part of its memory management by clicking **Create** under the configured Swap Files menu.
    a. At the Edit Swap File Configuration screen, enter the desired File Size (set it as your system memory size is recommended), then click **Next**.
    b. At the Swap Configuration window, click **Activate** under Configured Swap Files.
13. At the Network Connections screen, allocate network adapters to be used by virtual machines on the x445 server. Click **Next**.

> **Note:** If you have disabled the onboard Gigabit Ethernet controller, you will have no NICs to configure at this point. Skip to step 13.

14. At the Security Settings screen, select the desired level of security and click **Next**.
15. At this point your configuration is complete, click **OK** and the VMware Management Interface Login page appears.

> **Note:** If you disabled the onboard Broadcom Gigabit Ethernet controller, you will need to enable it via System Setup in BIOS, and then allocate it before proceeding:
>
> 1. Reboot the server, and press F1 to enter BIOS Setup.
> 2. Enable the Planar Ethernet, Exit and save. Reboot the server.
> 3. Once the system has booted, Log in to the Management Interface as root from a remote console.
> 4. Click **Options** from Management Interface, then click **Startup Profile**. The two Broadcom Ethernet controllers will appear on the window.
> 5. Select them dedicated to **Virtual Machines**, click **OK**.
> 6. Reboot server by clicking **OK**.

16. Click **Network Connections** from **Options** table as shown in Figure 5-17, you can configure and tune the network adapters that your virtual machines will use at here.

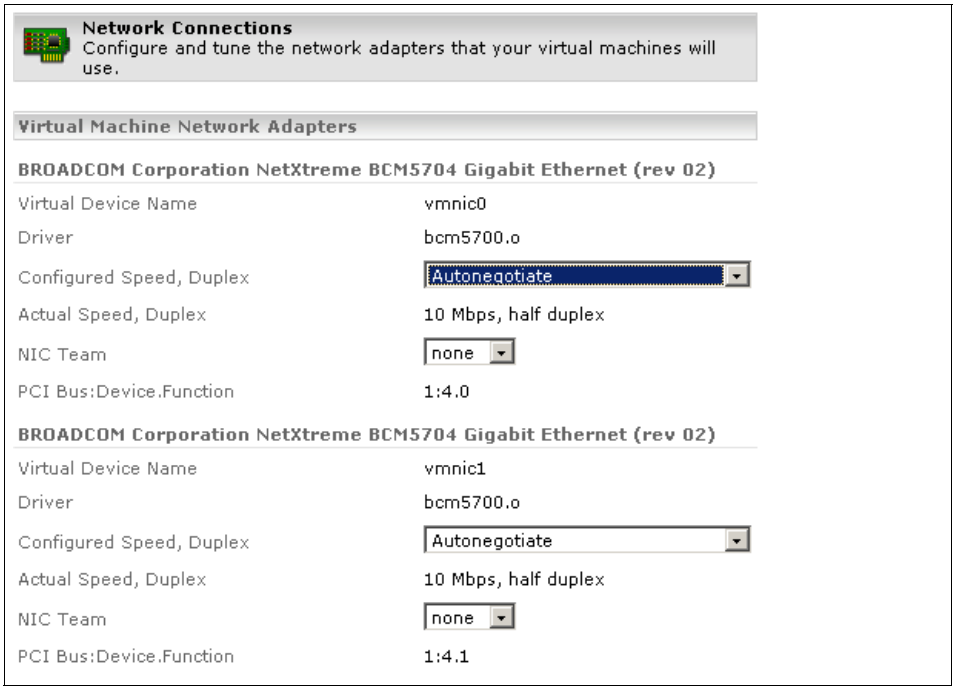

Figure 5-17 Network setting

You are now ready to create your virtual machines.

Creating a virtual machine

For more detail information about how to install ESX Server 2.0, how to create virtual machines, how to manage your virtual machines from a remote workstation and how to work with the advanced features of VMware ESX Server, please refer to the following product publications

- *VMware ESX Server Installation Guide*
- *VMware ESX Server Administration Guide*

You can download those books from VMware Web site:

`http://www.vmware.com/support/esx2/doc/`

Installing VMware Virtual SMP package

If you have purchased VMware Virtual SMP for ESX Server, then you can create Symmetric Multiprocessing (SMP) virtual machines, with single or dual virtual processors. You install VMware Virtual SMP for ESX Server when you enter its serial number during your ESX Server 2.0 configuration.

You also can update ESX Server 2.0 to support Virtual SMP function after installation and configuration is completed simply by entering the VMware Virtual SMP serial number in the Licensing and Serial Numbers window on the remote management interface.

If you enter a new serial number for a license, you will be prompted to reboot the server for the new license to take effect.

5.5 16-way x445 setup instructions

This section provides the basic information that is needed to setup an x445 16-way “merged” configuration. It is called merged because the two eight-way systems are merged at boot time to form a single 16-way image.

Tip: You can also use Scalable Systems Manager to configure the 16-way image. This is described in 6.2, “Scalable Systems Manager” on page 210.

If you configure the 16-way system using BIOS as described here, then when you run Scalable Systems Manager, it will import the BIOS configuration one time. You can then manage the partitioning using SSM.

Once the 16-way configuration is set up, you will be given the option at boot time to either continue in a merged 16-way configuration, or to boot the two systems separately at two eight-way servers.

Some steps might require the use of other x445 documentation that comes with your server.

Note: The following hardware restrictions apply when installing a 16-way configuration:

- Only 220V (200-240V) power connections will be supported.
- The diskette driver, IDE drivers, video, and USB will not be available on the secondary server once the system is configured as a 16-way. Keyboard and mouse also will not work on the secondary server, only on the primary server.
- The USB functions are only enabled after the BIOS merges the two eight-way systems. Consequently, when a USB keyboard is used, it will only work *after* the merge. PS/2® keyboards are active prior to the merge and two are therefore required if you need to boot the servers in an unmerged state.

5.5.1 Configuring the static partition

1. Connect power cords, mouse, keyboard, and monitor to each of the two x445 servers.
2. Using the two short SMP Expansion cables that come with your servers, complete the following steps to cable each server separately:
   a. Connect one end of the first SMP Expansion cable to port 3 of the top SMP Expansion Module on the primary server; then, connect the opposite end of the cable to port 3 of the bottom SMP Expansion Module on the primary server.
   b. Connect one end of the second SMP Expansion cable to port 3 of the top SMP Expansion Module on the secondary server; then, connect the opposite end of the cable to port 3 of the bottom SMP Expansion Module on the secondary server.
3. Update BIOS and RSA firmware for each of the two x445 servers. Both servers must have the same levels for:
   – System BIOS
   – System diagnostics
   – Remote Supervisor Adapter firmware

   No special 16-way system BIOS is necessary.
4. Decide which one of the two eight-way systems is to be the primary server and which one will be the secondary server. In a 16-way merged configuration, the operating system will boot from drives in the primary server. In an standalone configuration, the operating systems on each eight-way will boot from drives in each system.
5. Start the Configuration/Setup Utility program on the primary server.
   a. Turn on the primary server and watch the monitor screen.
   b. Press F1 when prompted during system startup to enter the System Configuration Utility.
   c. Select **Load default Settings** and press Enter for change current settings to default values.
   d. Enable Hyper-Threading Technology if desired. Enable Memory Mirroring or Hot Add Memory if desired.
   e. Save and exit setup.
6. Configure network settings for RSA.

Note: You can also setup an IP address for Remote Supervisor Adapter's network interface in x445 system BIOS. However, the network interface of RSA may be disabled — you cannot re-enable it via BIOS.

To set the IP address and verify the Remote Supervisor Adapter is enabled, we recommend you boot from the Remote Supervisor Adapter firmware update diskette, as described below.

a. Insert the RSA firmware update diskette into the primary server's diskette drive and boot the server.
b. Select **Configuration Settings** in the main menu of IBM Remote Supervisor Utility program, then select **Ethernet Settings**.
c. Enable the network interface of RSA, if a DHCP server is not used, then you must assign a IP address. A static address is recommended, although fixed names can be cross referenced to the new address.
d. Exit the utility and turn off the server.

7. Repeat steps 5 and 6 on the secondary server.
8. Connect the Remote Supervisor Adapters's network interface of each server to the network or use the Ethernet crossover cable to connect them to each other. A separate hub and network for RSA is recommended so that you can also connect to the Remote Supervisor Adapter for management purposes.
9. If you have not already done so, connect the RXE expansion and RXE management cables between the server and the remote expansion enclosure. Refer to "Cabling an RXE-100" on page 185 for instructions on how to connect the RXE-100 to 16-way system, then, return here.

Note: Unlike eight-way or less configurations, we recommend you connect the RXE-100 before operating system installation for 16-way configurations.

The remaining steps describe how to create the 16-way partition.

Tip: Instead of configuring the 16-way system via BIOS, you can also use IBM Director and Scalable Systems Manager as described in 6.2, "Scalable Systems Manager" on page 210. This has the added advantage of the use of discovery and diagnostic tools.

10. Turn on the primary server and start the Configuration/Setup Utility by press F1 when prompted. The secondary server must be powered off, but with AC

power still supplied (that is, the green power LED on the front panel is blinking)

11. Select **Advanced Setup** from the main menu.
12. Select **Scalable Partition Settings**.
13. Select **Create Scalable Partition Settings**. Figure 5-18 appears.
14. Select **2 x445 chassis**, as the Scalable System Configuration.
15. Enter a Scalable Partition ID as 0001.
16. For the Primary Scalable Node select **SMP ID 0**. The SMP ID can be 0 or 2, but it must be different from the SMP ID of the secondary server.
17. Type the IP address of the Remote Supervisor Adapter in the secondary server in the Secondary ASM Host Name field.

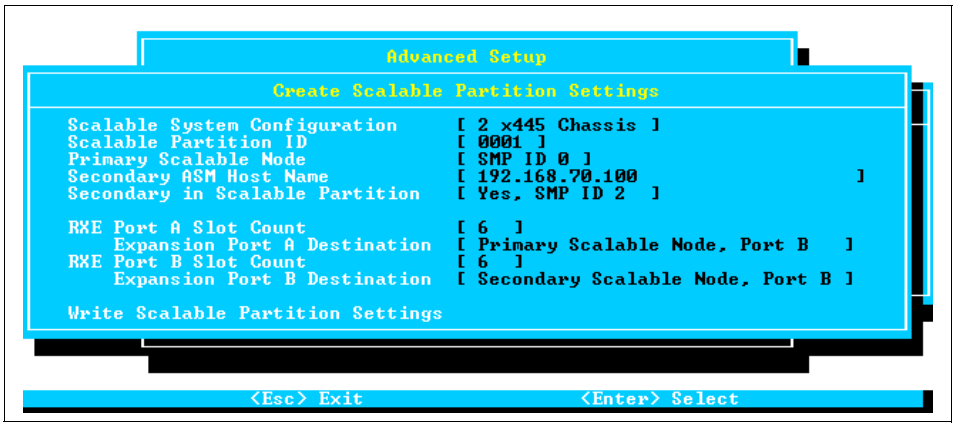

Figure 5-18 Scalable partition settings

18. From the Secondary in Scalable Partition menu, select **Yes, SMP ID 2**.
19. Select the number of PCI-X slots available to the RXE Expansion Port in the remote expansion enclosure. The following choices are available:
    - 0 - No PCI-X slots are assigned to the RXE Expansion Port.
    - 6 - Only six of the 12 slots are assigned to the RXE Expansion Port.
    - 12 - All 12 PCI-X slots are assigned to the RXE Expansion Port.
20. If 12 PCI-X slots are available to port A, select **Primary Scalable Node, Port B** for the RXE Port A destination.

    If six PCI-X slots are available to port A and six PCI-X slots are available to port B, then select **Primary Scalable Node, Port B** for the RXE Port A destination and **Secondary Scalable Node, Port B** for the RXE Port B destination. For performance reason, the latter setting is recommended as shown in Figure 5-18.

21. Select **Write Scalable Partition Settings** and follow the prompts.
22. Exit the Configuration/Setup Utility Program.
23. Turn off the primary server, then continue with the next section.

5.5.2 Cabling the servers

After create scalable partition, we should connect two x445 servers together to form the 16-way system as shown in Figure 5-19.

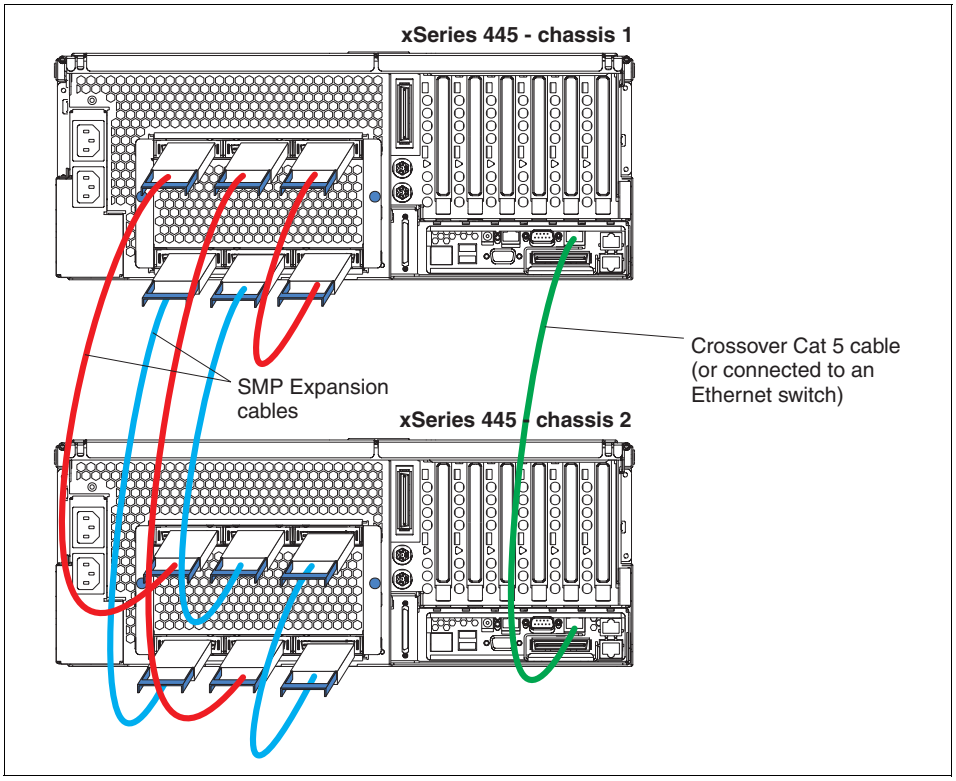

Figure 5-19 Connecting two x445 together in a 16-way configuration

Turn off both servers and connect the cables as follows:

Table 5-1 16-way cable connections

Cable	Cables connect on the primary...	to the secondary...
Long	Bottom module, port 1	Bottom module, port 1
Long	Top module, port 1	Top module, port 1

Cable	Cables connect on the primary...	to the secondary...
Long	Bottom module, port 2	Top module, port 2
Long	Top module, port 2	Bottom module, port 2
Short	**Primary** server: Top module, port 3 to Bottom module, port 3	
Short	**Secondary** server: Top module, port 3 to Bottom module, port 3	

5.5.3 Booting the 16-way system

Now that the 16-way configuration is cabled, the server can be booted:

1. Turn on the primary server, if the 16-way x445 configuration settings and cabling is correct, the secondary server and RXE-100 will power on automatically.

   Note: The procedure for turning your 16-way system on and off is no different than the procedure on a single x445 server. However, pressing the power-control button on either the primary or secondary server in the 16-way system will turn on or turn off both servers and RXE-100.

2. The following information will display on primary server:

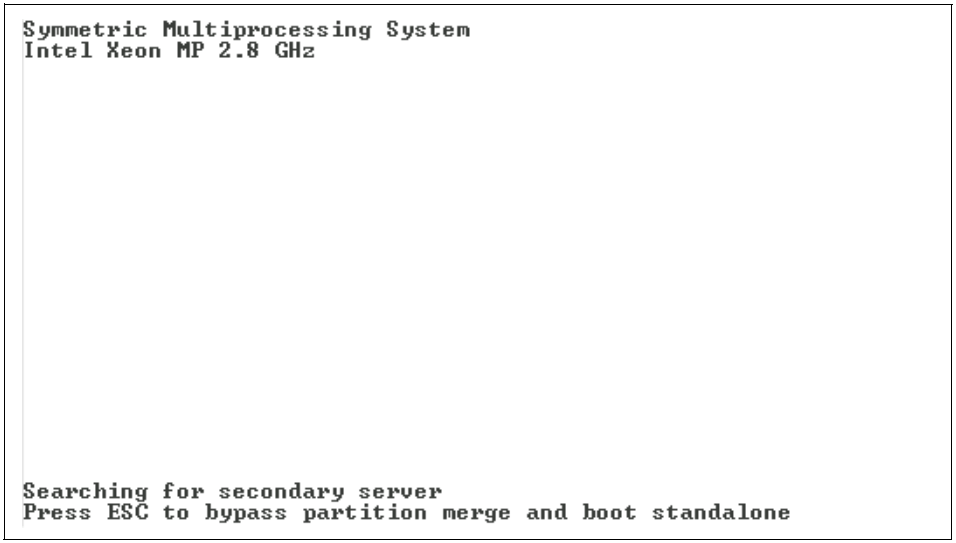

Figure 5-20 Primary server start up screen

3. If you want boot to standalone server, press Esc at this time.

4. The secondary server will display `Waiting for primary server` on screen.
5. If two servers are successfully merged as a 16-way system, a message will display on primary server as shown in Figure 5-21.

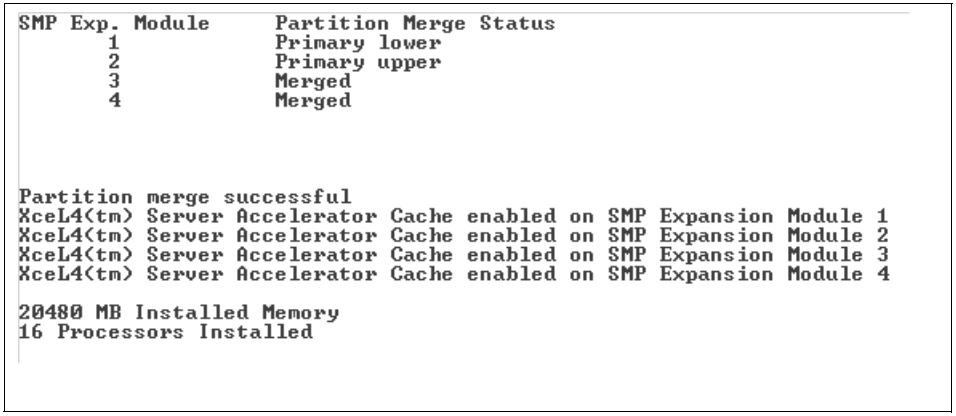

Figure 5-21 Two x445 servers merged as a 16-way system

6. Once the two server are merged together, the video on the secondary server will display

   `Merge complete -- see primary server display`

7. Install your operating system per the instructions in 5.4, “Operating system installation” on page 152. The operating system is installed on disks in the primary server.

5.5.4 Cabling an RXE-100

In the 16-way configuration, only one RXE-100 can be connected as shown in Figure 5-22.

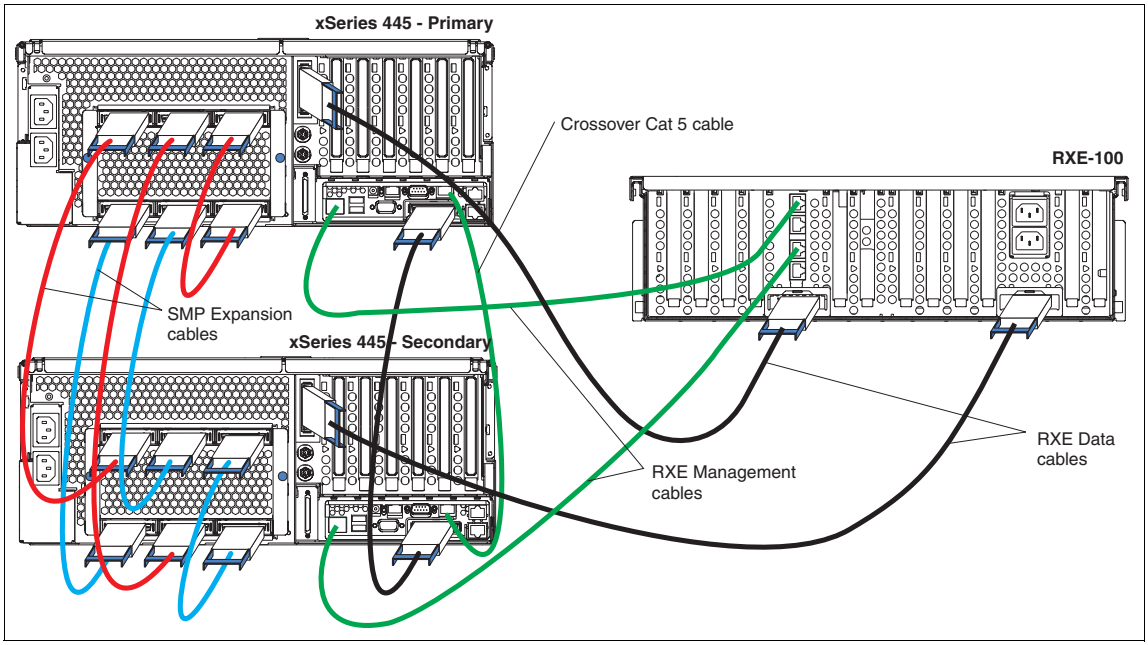

Figure 5-22 Connecting an RXE-100 to a 16-way x445 system

Complete the following steps to connect the IBM RXE-100 Remote Expansion Enclosure to a 16-way system.

1. Turn off both servers.
2. Using an RXE Expansion cable, connect RXE Expansion Port A on the primary server to RXE Expansion Port A on the secondary server.
3. Using an RXE Expansion cable, connect RXE Expansion Port B on the primary server to RXE Expansion Port A on the RXE-100.
4. Using an RXE Expansion cable, connect RXE Expansion Port B on the secondary server to the RXE Expansion Port B on the RXE-100.
5. Using an RS485 RXE Management cable, connect the RXE Management Port on the primary server to RXE Management A (in) Port on the RXE-100.
6. Using an RS485 RXE Management cable, connect the RXE Management Port on the secondary server to RXE Management B (in) Port on the RXE-100.

5.6 Troubleshooting

This section provides additional information that is needed to troubleshoot servers during the installation.

The following are some hints and tips published on the IBM Web site related to x445 server installation. For an up-to-date list of tips, go to the following URL and display the Hints & Tips for the x445 (machine type 8870):

http://www.pc.ibm.com/support

- x445 - Microsoft Windows 2003 installation hangs at final shutdown

  http://www.ibm.com/pc/support/site.wss/MIGR-52480.html

- View Blue Screen option not supported remotely under Microsoft Windows Server 2003

  http://www.ibm.com/pc/support/site.wss/MIGR-51942.html

- Microsoft Windows 2003 only supports NUMA configurations with PAE enabled

  http://www.ibm.com/pc/support/site.wss/MIGR-52036.html

- x445 - Cannot remote install Microsoft Windows Server 2003 until Microsoft provide a hot-fix

  http://www.ibm.com/pc/support/site.wss/MIGR-52038.html

- Blue screen with ServeRAID 6.xx and Director 4.1 and Microsoft Windows Server 2003

  http://www.ibm.com/pc/support/site.wss/MIGR-52387.html

- x445 - Hot-adding memory under Microsoft Windows 2003 results in slow performance

  http://www.ibm.com/pc/support/site.wss/MIGR-52034.html

- x445 - System Management tests unavailable under diagnostics Version 1.00

  http://www.ibm.com/pc/support/site.wss/MIGR-52030.html

- x445 - Novell NetWare restarts when NCMCON is loaded

  http://www.ibm.com/pc/support/site.wss/MIGR-52043.html

- x445 - Broadcom 5704 Gigabit Ethernet device must be updated to Version 6.67 or later

  http://www.ibm.com/pc/support/site.wss/MIGR-52039.html

- Microsoft Windows Server 2003 numerates ServeRAID disks in reversed order

  http://www.ibm.com/pc/support/site.wss/MIGR-52389.html

- Tape operations result in Windows Server 2003 Blue Screen

  http://www.ibm.com/pc/support/site.wss/MIGR-52368.html
- x445 - Processor VRM LEDs always on when processor VRM not present

  http://www.ibm.com/pc/support/site.wss/MIGR-52037.html
- Fans increase speed after hot swapping the power supply

  http://www.ibm.com/pc/support/site.wss/MIGR-52031.html
- x445 - Remote video corrupted via Remote Superviser Adapter with 8-bit color modes

  http://www.ibm.com/pc/support/site.wss/MIGR-52040.html

For additional information on the installation of the x445, refer to the following product publications posted on the Web:

- *IBM @server xSeries 445 Installation Guide*
- *IBM @server xSeries 445 Option Installation Guide*
- *IBM @server xSeries 445 User's Guide*
- *IBM @server xSeries 445 Hardware Maintenance Manual*
- *IBM @server xSeries 445 Rack Installation Guide*
- *IBM Remote Expansion Enclosure Installation Guide*

These product publications can be downloaded in PDF format from:

http://www.pc.ibm.com/support

6

Management

IBM Director is the systems management software provided with the IBM @server xSeries servers. The latest version at the time of writing is 4.12, which is effectively the same as the base Version 4.1 updated with the Service Pack 1. This chapter covers the three plug-ins to IBM Director that are of particular relevance to the x445:

- 6.1, “Active PCI Manager” on page 190

  Active PCI Manager guides you when installing new PCI adapters and can analyze an existing PCI configuration and suggest changes to improve performance. Additional function introduced in V4.1 is the management of network controllers in a teaming configuration. Active PCI Manager is one of the components in the Server Plus Pack and comes at an extra charge.

- 6.2, “Scalable Systems Manager” on page 210

  Scalable Systems Manager is used to partition an x445 scalable system (8 or 16-way) so that each partition behaves as an independent server. You can create one partition on an 8-way system, or one or two partitions on a 16-way system.

- 6.3, “Application Workload Manager” on page 227

  The Application Workload Manager enables you to control the CPU and memory resource allocation to individual applications. This can be especially useful in server consolidation scenarios.

6.1 Active PCI Manager

IBM Director 4.1 introduced a new look to the Active PCI Manager in the IBM Director Console. As shown in Figure 6-1 on page 191, it now has two components:

- FTMI (Fault Tolerant Management Interface)

  This is a new component that allows you to work with the network controller teams, which have been created with various vendors' network teaming utilities. Examples of these utilities are:

  - 3Com DynamicAccess software for server configuration
  - Intel PROSet
  - Broadcom Advanced Control Suite

  The actions available are placing an adapter in a team offline or online, ejecting an adapter, and viewing information about the adapters and teams.

- Slot Manager

  This task is what the Active PCI Manager used to be in V3.1. Because the FTMI component was added in Version 4.1, and the Active PCI Manager became the parent name, the name of this task changed to Slot Manager.

  Slot Manager is used to plan the placement of PCI and PCI-X adapters to the x445 and RXE-100 expansion enclosures, to best use the bus architecture of the systems. It can also analyze existing adapter arrangements to determine if the configuration is optimal, and if not, suggest alternative configurations. If required, it then graphically assists with changing the adapter placement.

This section concentrates on Slot Manager.

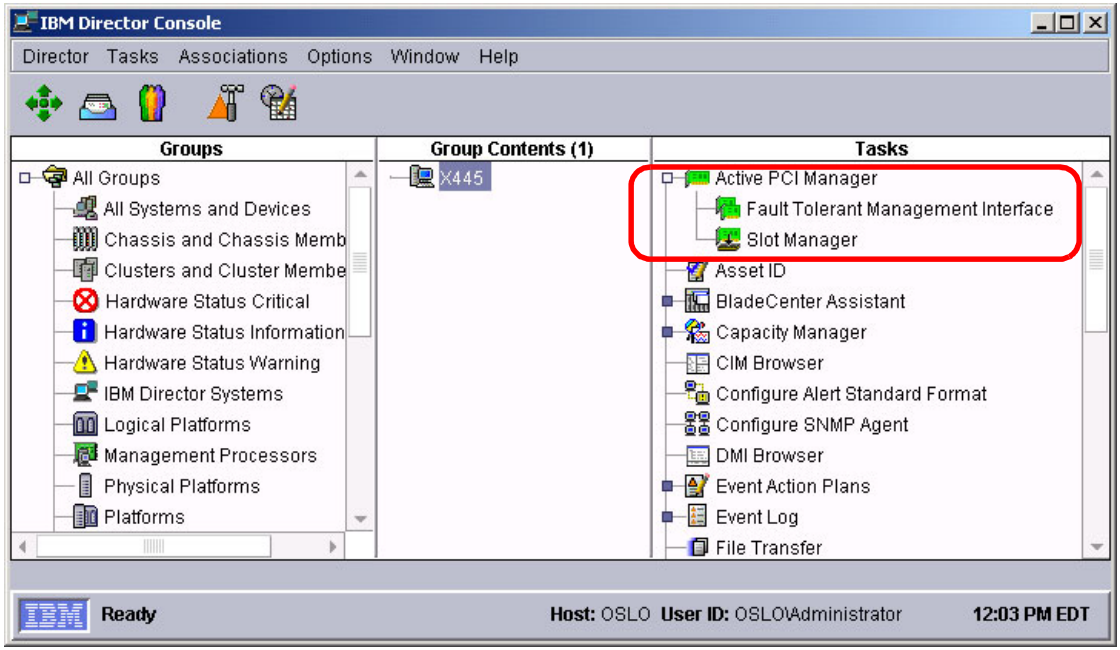

Figure 6-1 Active PCI Manager components

Like other IBM Director plug-ins, there are three components:

- Agent component, installed on the x445
- Server component, installed on the IBM Director Server
- Console component, installed on each IBM Director console

Note: All Active PCI Manager functions are performed from the IBM Director console.

The agent has the following prerequisites:

- Windows 2000 Server or Advanced Server with Service Pack 2 or higher, or Windows Server 2003
- IBM Active PCI Software Version 5.0.2.0 or later. This software contains the following components:
  - IBM Active PCI filter driver
  - IBM Active PCI alert driver
  - IBM Active PCI alert service

  The latest version is 5.1.1.1 and you can download it from:

  `http://www.pc.ibm.com/qtechinfo/MIGR-4J2QEQ.html`

Once installed, the Active PCI Manager appears as a parent component for the FTMI and Slot Manager, as shown in Figure 6-1 on page 191.

6.1.1 Using the Slot Manager

The Slot Manager user interface is shown in Figure 6-2. The details of your system configuration are available in three ways by clicking the appropriate tab:

- Slot view, Figure 6-2
- Tree view, Figure 6-4 on page 196
- Table view, Figure 6-5 on page 197

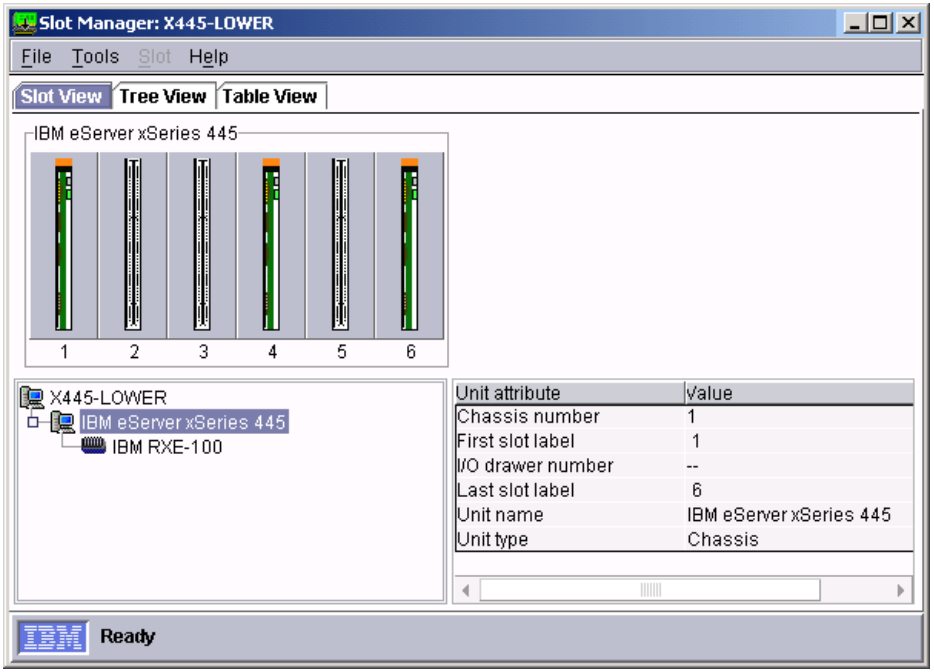

Figure 6-2 Slot Manager - slot view

In the slot view, we see a graphical representation of the PCI-X slots in the x445 or in the RXE-100. The x445 in Figure 6-2 contains three adapters in slots 1, 4 and 6. The orange bar above the adapter indicates that the adapter present in that slot is hot-swappable.

The tree pane in this view (lower left) shows there is an RXE-100 attached to the server. By selecting the RXE-100, the graphical view changes to show the six or 12 slots in that enclosure.

The details pane in this view (lower right) includes information about the currently selected item in the tree pane or the currently selected slot. In Figure 6-2, the server is selected and the details pane shows information about the system and its slots, including:

- Chassis Number: The x445 is assigned chassis number 1. In multi-chassis x445 configurations, each chassis will be identified by a unique number.
- First Slot Label, Last Slot Label: The first and last slot labels in the server are identified. When the RXE-100 is selected, the slot labels are A1 through A6 and B1 through B6. The labels correspond to the names of the slots as printed on the chassis.
- I/O Drawer: This value is only relevant when an RXE-100 is selected. The number is unique and the first drawer will be number 0.
- Unit Name, Unit Type: The unit name is IBM eServer xSeries 445 and the unit type is Chassis. When you select the RXE-100, the unit name is IBM RXE-100 and the unit type is I/O Drawer.

Selecting a slot results in the details of that slot appearing in the bottom right pane, as shown in Figure 6-3.

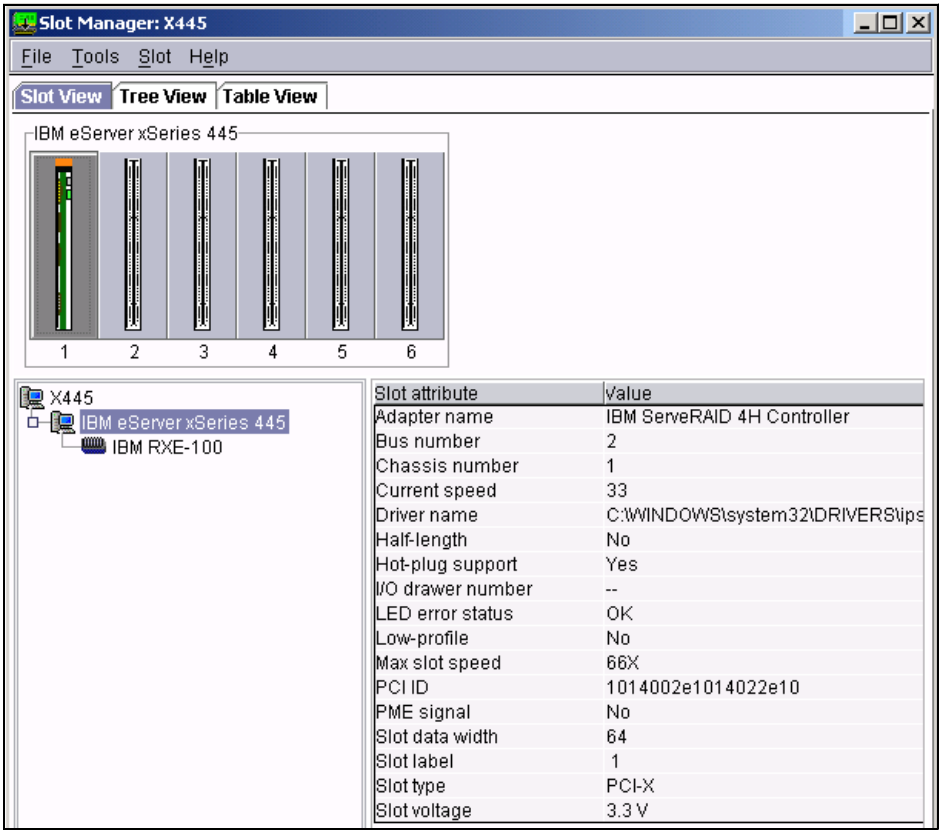

Figure 6-3 Slot Manager - slot 1 details

For information about the attributes, consult the online help. However, some explanation is warranted here:

- Current speed, Max slot speed: As described in 1.7, “PCI subsystem” on page 22, the adapters installed in this slot and other slots in the same bus dictate the speed that the slot will run at. Consequently, the current speed of the slot may not be the maximum possible speed. The “X” in the Current Speed value indicates the slot is running in PCI-X mode (as opposed to running in PCI mode).

  Some RXE-100 slots may indicate “Not supported” for the slot speed (see Figure 6-5 on page 197 for an example). These are even-numbered slots when the odd-numbered slot in the pair is running at 133 MHz. This is because if both slots in a pair are used, the maximum speed of both slots is 100 MHz.

- PME signal: PME (power management events) is a hardware signal that some PCI slots can supply. Most PCI cards do not implement PME and the x445 and RXE-100 slots do not support PME.
- Bus number: a system-assigned number for each bus. Useful for identifying slots that are on the same bus, but the bus number does not correlate to the slot numbers.
- Low-profile: The slot *only* supports cards that have a half-height end plate. See the PCI Special Interest Group specifications of low-profile PCI adapters at http://www.pcisig.com/data/specifications/lowp_ecn.pdf.
- LED error status: Table 6-1 shows the possible LED status messages and their meanings:

Table 6-1 LED status messages

Message	Meaning	Recommended Action
OK (no error)	None of the other conditions exist.	None
Hot eject successful	Adapter removal completed without error.	None
Bus speed mismatch	A second adapter was hot-added but its rated speed does not match the current speed of the bus. The new adapter is held inoperative.	Either: ► Move one of the adapters to an empty bus, or a bus with an adapter of the same speed. ► Reboot the server and the bus will slow down so that both adapters can operate.
Power fault on card in slot	Adapter has short-circuit or other problem preventing normal operation.	Remove the adapter and have it repaired or replaced.
Surprise removal occurred	Operating system activity to slot was not stopped before power to adapter was removed.	This should be avoided. Use the Unplug or Eject Hardware wizard to disable the adapter before removing it.
Slot disabled at current speed	A second 133 MHz adapter was hot-added to a two-slot bus that already had a 133 MHz adapter installed. The new adapter is held inoperative.	Restart the server. The bus will run at 100 MHz and both adapters will be operative.
Too many adapters on bus	Two PCI-X 133 MHz adapters in a slot pair.	Move one of the adapters to an empty bus.
Bus connection error	System has detected a hardware fault.	Remove the adapter. If this does not correct the condition, have system serviced.

The tree view, shown in Figure 6-4, lets you look at all of the slots in the system at once.

The attributes of the selected slot and the adapter installed in it are shown on the right. For information about the attributes, see the discussion following Figure 6-3 on page 194 and the online help.

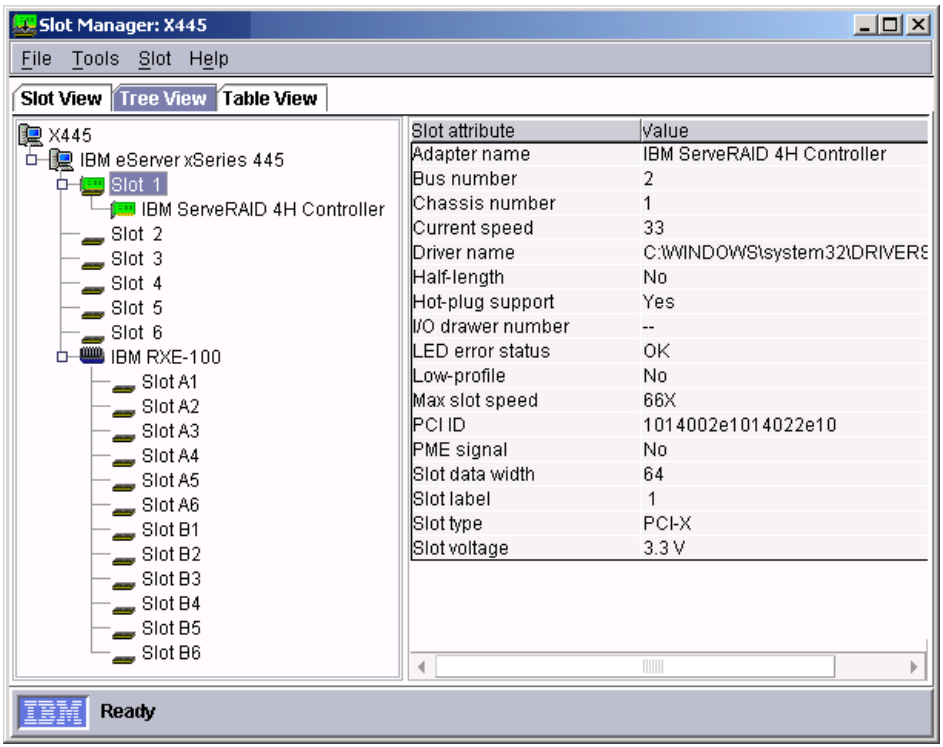

Figure 6-4 Slot Manager - tree view

The table view, Figure 6-5, provides a summary of all of the slot and adapter information in one table. The columns in this view are similar to the attributes in the tree view and slot view.

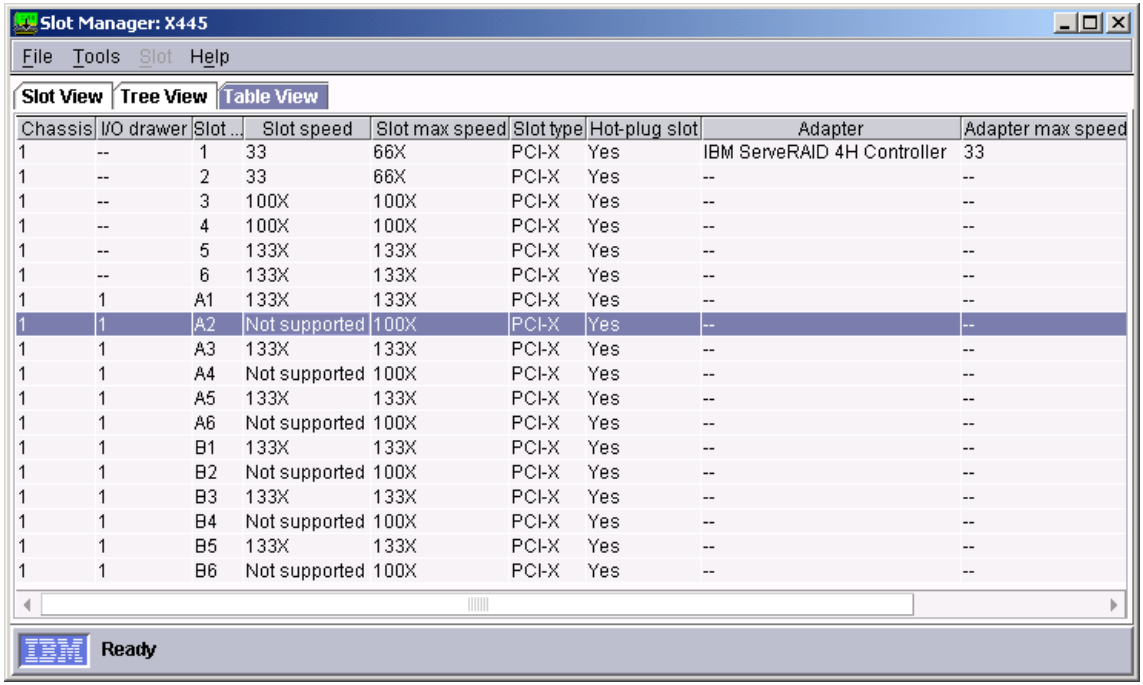

Figure 6-5 Slot Manager - table view

6.1.2 Adding adapters to the system

The process of adding an adapter to an x445 using the Slot Manager is simple. As an example, let's start with a server with no adapters, and then add a ServeRAID-4H.

Note: The x445 and RXE-100 only support 3.3 V PCI and PCI-X adapter. 5 V adapters are not supported.

The processing of adding an adapter depends on whether the adapter is on the list of known adapters or you have to specify the adapter's characteristics.

Adding a known adapter

The process to add a known adapter, such as the IBM ServeRAID-4H, is as follows:

1. Select **Tools** → **Add Card Wizard**. The window shown in Figure 6-6 on page 198 appears.

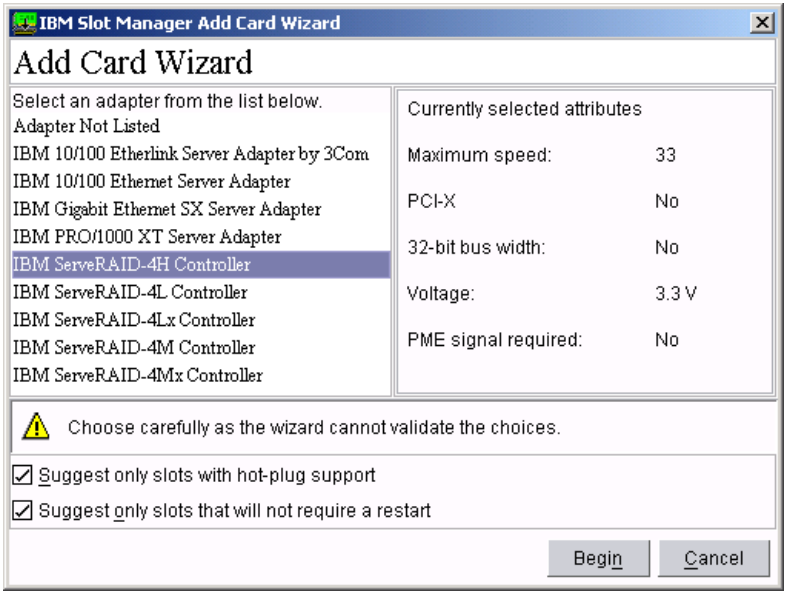

Figure 6-6 Slot Manager - Add Card Wizard

The bottom two check boxes are selected, indicating the following default settings:

- Slot Manager will suggest only slots with hot-plug support (all of the slots in the x445 and RXE-100 support hot-plug).
- Slots that will not require a server restart.

  The Add Card Wizard, when finding the best slot in which to place the new adapter, may find a slot that needs to reset the bus in order for it to work.

  An example would be a 133 MHz PCI-X adapter being added to a bus that already has a 133 MHz PCI-X adapter installed. Since the hardware only supports two adapters in the one bus running at most at 100 MHz, the bus must be reset to run at the lower speed. Restarting the bus requires that the operating system be rebooted.

  With this checkbox unchecked, the wizard may suggest the adapter be added to a 133 MHz bus if it is the best slot, and will ask the user to restart the system. With the box checked, the wizard will not suggest the slot on a 133 MHz bus even though it is the best slot, but will try to find a slot that does not require a server restart.

  By limiting the wizard this way, it may fail to find a slot. If it does find a slot, the adapter will not require the system to be restarted, but the adapter may not be running at the best speed or in the best mode.

2. To add the IBM ServeRAID-4H, select it from the list. The characteristics of the adapter are listed on the right side, as shown in Figure 6-6.
3. To insert the adapter, click **Begin**. Slot Manager analyzes the available slots to find the one that provides the best performance for this adapter. While this is running, the following message is displayed:

   ```
   Wait while the system is being analyzed. This might take several
   minutes.
   ```

4. Upon completion of the analysis, the Slot Manager recommends the best slot for the adapter (if one is available with the right criteria), as shown in Figure 6-7. In our case, this is slot 1in the x445.

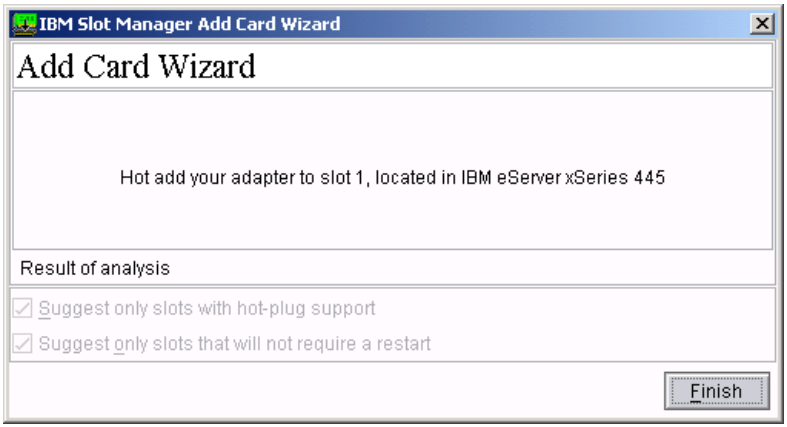

Figure 6-7 Slot Manager - slot recommendation

5. Click **Finish** to close the wizard.
6. From the slot view, select the slot that the Add Card Wizard recommended, then right-click the slot and click **Blink Slot**, as shown in Figure 6-8.

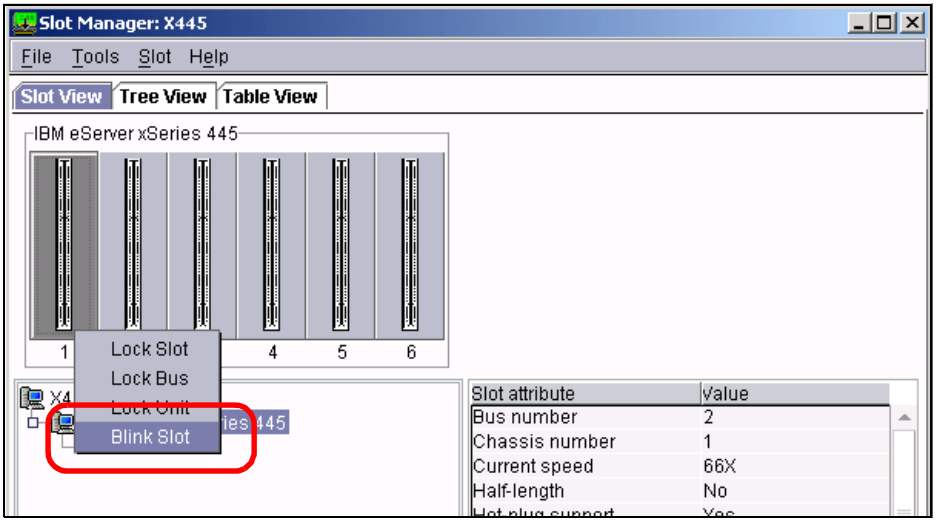

Figure 6-8 Slot Manager - blink slot

The LED on the slot itself will blink as does the graphic of the slot in the slot view. This provides easier identification of the target slot when inserting the adapter.

Lock options: There are lock options for a slot (see Figure 6-8). These will lock a slot, all of the slots on the bus that a slot is a member of, or the unit (the physical chassis or I/O expansion enclosure that a slot is a member of).

Locking the slots excludes them from being processed by:

- Add Card Wizard
- Performance optimizer, described in 6.1.3, “Analyzing an existing configuration” on page 204.

Note: The table view does not provide access to the lock function.

7. To hot add the adapter into the slot, open the cover on the x445 or RXE-100. Push the orange tab above the adapter slot in the direction of the arrow, toward the rear of the x445 or RXE-100, and swing the black adapter retaining arm so that it is vertical. This will allow you to remove the blank filler plate and add the adapter.
8. Once the adapter is installed, return the black adapter retaining arm to the closed position. The orange tab will click back into place, locking the adapter retaining arm down.

Important: Be sure to close all adapter retaining tabs before closing the cover on the x445 or RXE-100, or the cover will not close properly and forcing it will likely break the adapter retaining arm.

9. After the adapter has been added, the slot view window is dynamically updated.

Adding an unlisted adapter

To add an adapter whose characteristics are not known to the wizard, such as the QLogic QLA23xx PCI Fibre Channel Adapter, do the following:

1. Select **Tools** → **Add Card Wizard**. Leave the default selection, **Adapter Not Listed,** highlighted, and click **Next**. Figure 6-9 on page 201 is displayed.

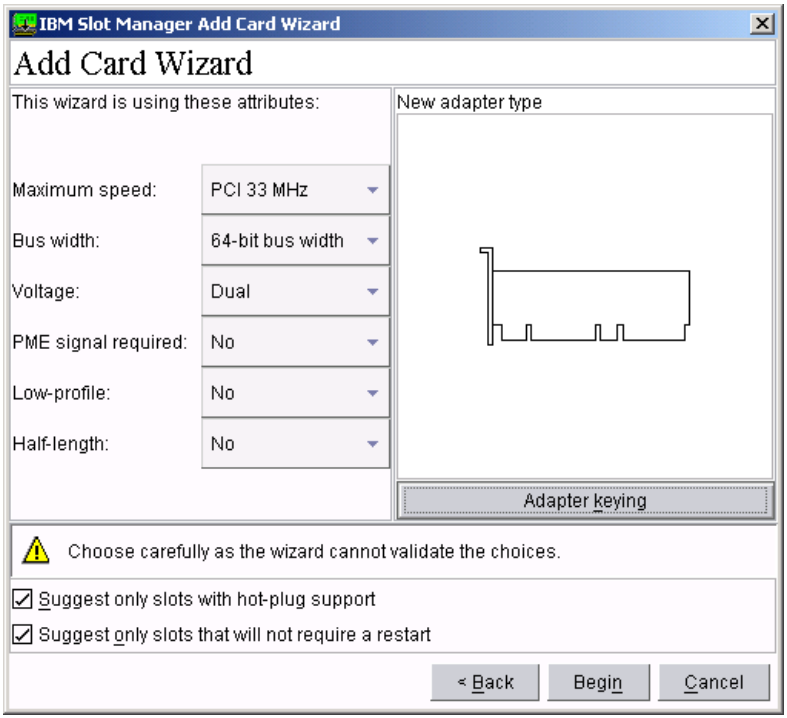

Figure 6-9 Adapter identification window

There are two ways of specifying the characteristics of an unknown adapter:

– Manually select the attributes from the pull-down menus in the left pane in Figure 6-9.

– Click the **Adapter Keying** button to cycle through drawings of adapters in the right pane in Figure 6-9, and select the appropriate adapter from the shape and location of the PCI connector.

  The shape of the PCI connector indicates the voltage support (3.3 V or 5 V or both) and the bus width (64-bit or 32-bit), as shown in Figure 6-10.

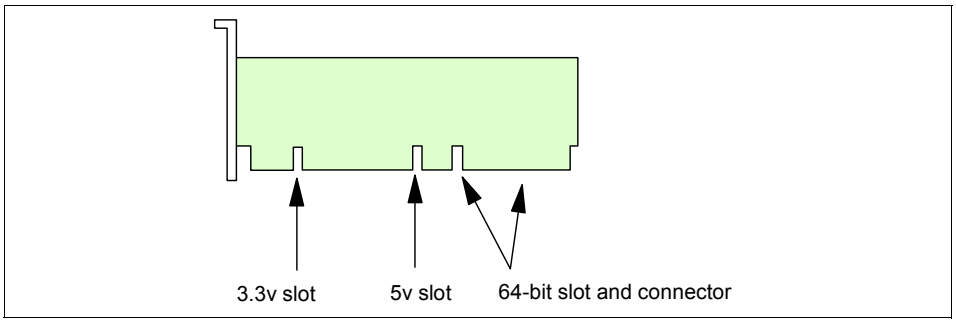

Figure 6-10 PCI connector identification

Table 6-2 shows you the possible choices when you click **Adapter Keying**. As you click the button, both the graphic and the bus width and voltage pull-down menu entries change.

Note: Clicking the **Adapter Keying** button only changes the bus width and voltage pull-down menus. The other fields (maximum speed, PME, low profile, and half length) must be set manually.

Table 6-2 Adapter attributes and associated graphics

Attributes	Graphic
Bus width: 64-bit Voltage: dual (3.3 V and 5 V)	
Bus width: 64-bit Voltage: 3.3 V	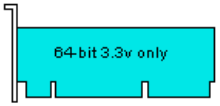
Bus width: 32-bit Voltage: Dual (3.3 V and 5 V)	

Attributes	Graphic
Bus width: 64-bit Voltage: 5 V **Note**: 5 V adapters are not supported in the x445 and RXE-100	
Bus width: 32-bit Voltage: 5 V **Note**: 5 V adapters are not supported in the x445 and RXE-100	

2. Review the documentation that is included with the adapter to determine its characteristics. In our case, they are:
   - PCI-X 66 MHz
   - 64 bit
   - Dual voltage
   - No PME signal
   - Full height (that is, not low profile)
   - Half length
3. Enter the remaining values into the pull-down menus. We enter our parameters and the following message is displayed:

   ```
   Hot add your adapter to slot 3, located in the IBM eServer xSeries 445
   ```
4. Blink the slot as described in step 6 on page 199.
5. Hot add the adapter as described in steps 7 and 8 on page 200.
6. The Slot Manager updates its display to reflect the addition of the adapter, and we see the operating system recognize the addition of the adapter, as shown in Figure 6-11.

Figure 6-11 Windows finds new hardware

7. If required, supply the device driver for the adapter as usual to complete the installation.

The Slot Manager is not required to hot add an adapter, but the use of the blink slot tool will reduce the likelihood of error. The use of the Add Card Wizard can insure that the adapter slot selection will result in an optimal configuration, if used correctly.

6.1.3 Analyzing an existing configuration

The Slot Manager can be used to analyze an existing configuration and assist with the optimization of that configuration, if required. We have intentionally re-arranged the adapters described in the previous sections and added some new ones to illustrate this capability. Figure 6-12 shows the configuration we are starting with.

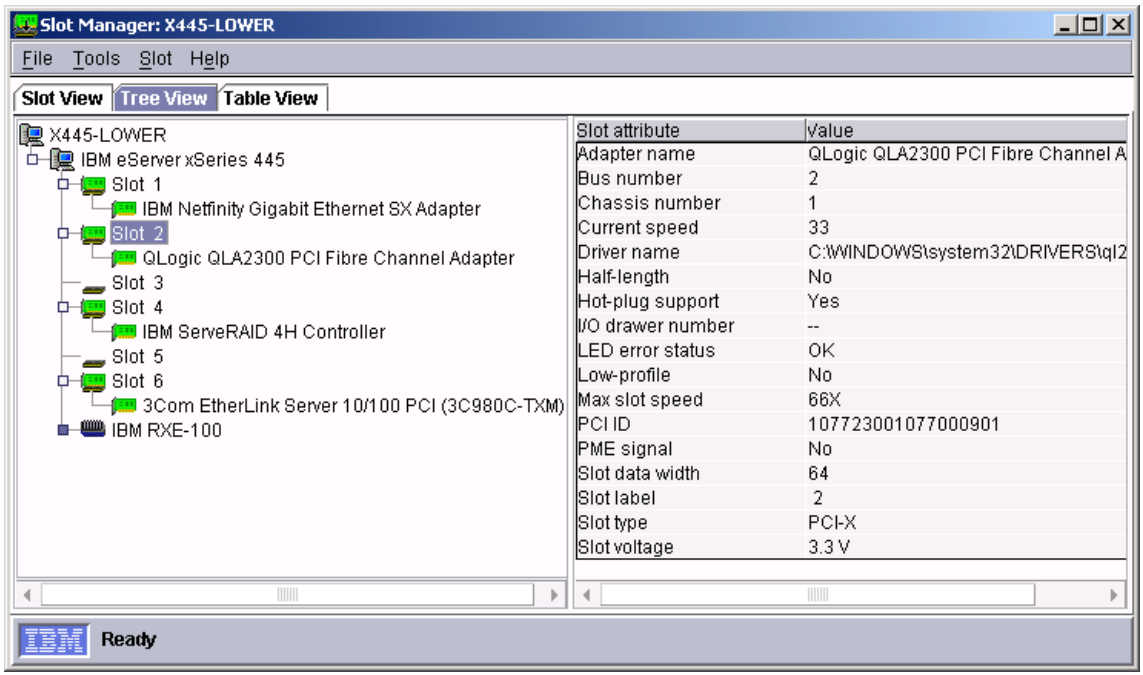

Figure 6-12 Non-optimal adapter arrangement

To analyze the adapter configuration, click **Tools** → **Analyze.** Once the analysis is complete, the Optimization Steps window appears, as shown in Figure 6-13 on page 206. It includes three sections:

- Observations — problems identified
- Suggested Adapter Arrangement — adapter layout after recommended changes are made
- Recommended Actions — recommended changes

In our case a performance problem was discovered.

Note: You will note that the optimizer reports that both slots 2 and 4 contain boot devices. These are *potential* boot devices, not necessarily actual boot devices.

IBM does not support booting from a Fibre Channel device. However, the Slot Manager recognizes the QLogic adapter as a potential boot device. It does not make recommendations that include moving bootable devices, since this may alter the order in which bootable devices are enumerated during startup, resulting in a system that will not boot from its intended boot device (or at all, if no other bootable devices exist).

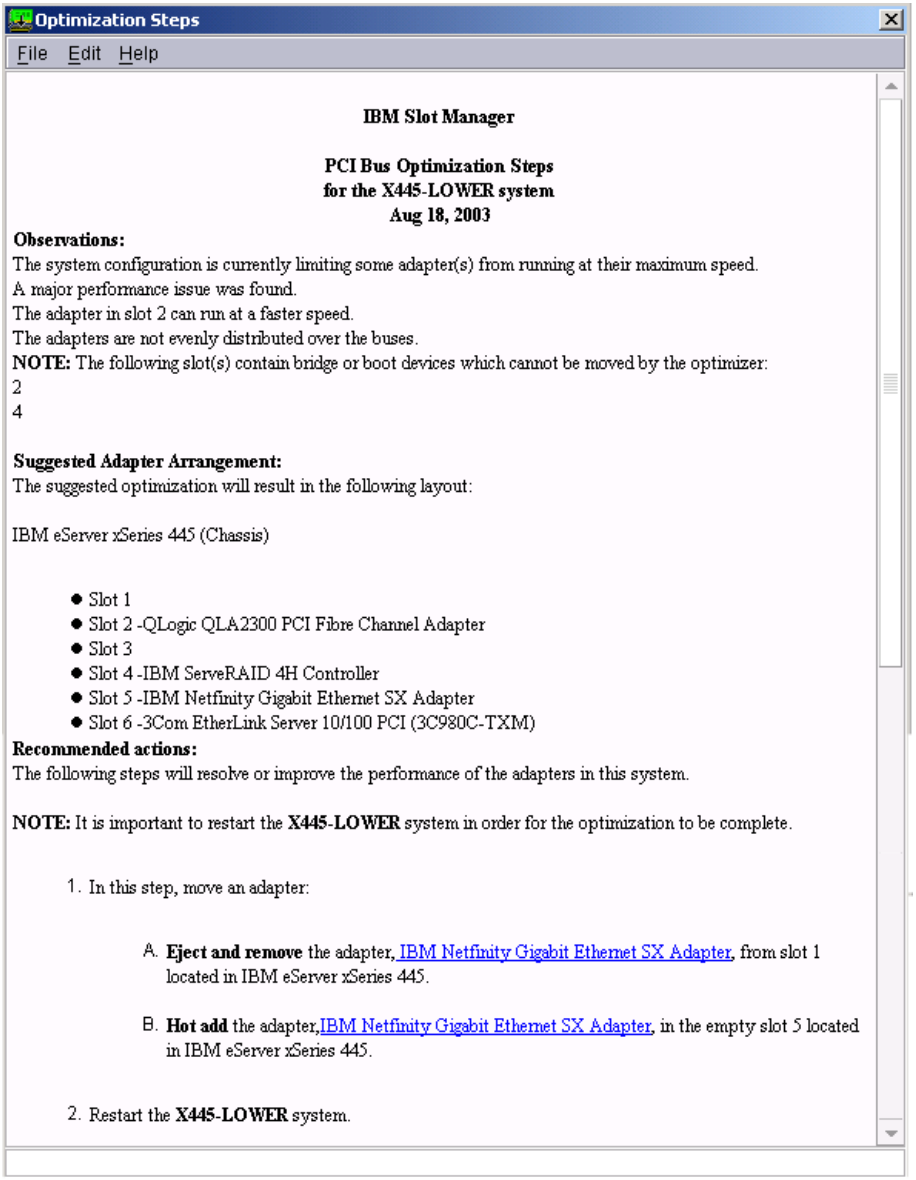

Figure 6-13 Optimization steps

In the Recommended Actions, follow the instructions to move adapters. In our example, the recommendation is to move an adapter from slot 1 to slot 5:

1. Click the link on the line **Eject and Remove** the adapter. Figure 6-14 on page 207 appears. An upward-moving arrow indicates the adapter to remove.

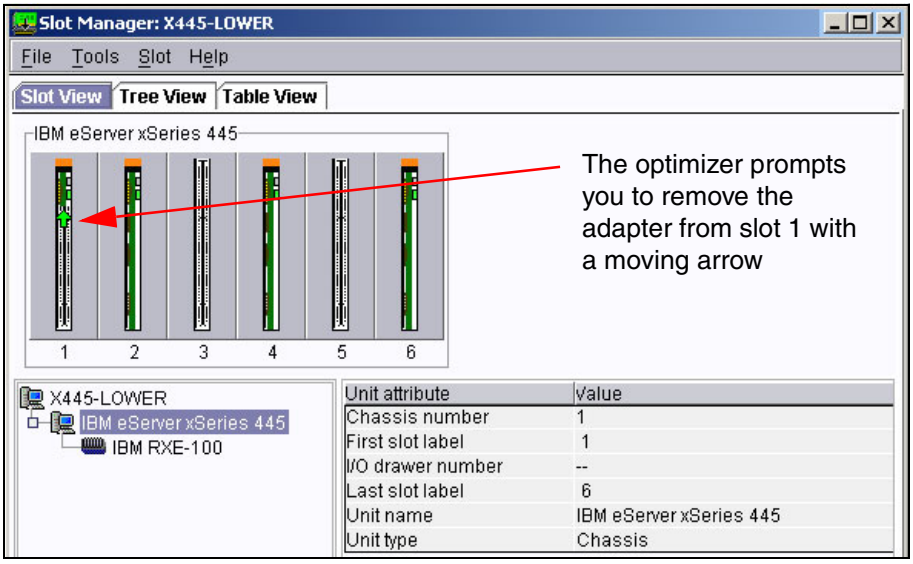

Figure 6-14 Indication of the slot for removal

Tip: The two links in Figure 6-13 merely show you which slot to insert the adapter into or remove the adapter from.

2. Stop the adapter in the operating system. Double-click the  icon in the system tray. Figure 6-15 on page 208 appears.

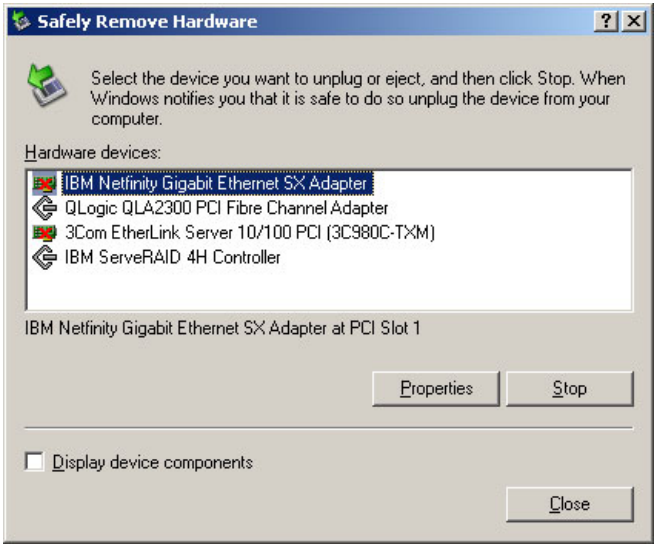

Figure 6-15 Stop adapter use by the operating system

3. Select the adapter you want to stop and click **Stop**.
4. Confirm that you want to stop the device by clicking **OK**.
5. Once the operating system confirms the adapter has been stopped, use the Slot Blink action (see step 6 on page 199) to indicate which adapter to remove.
6. Remove the adapter.
7. The Slot Manager then shows the adapter has been removed, as shown in Figure 6-16 on page 209, by adding an to the bottom of the graphic and changing the LED error status to “Hot eject successful”.

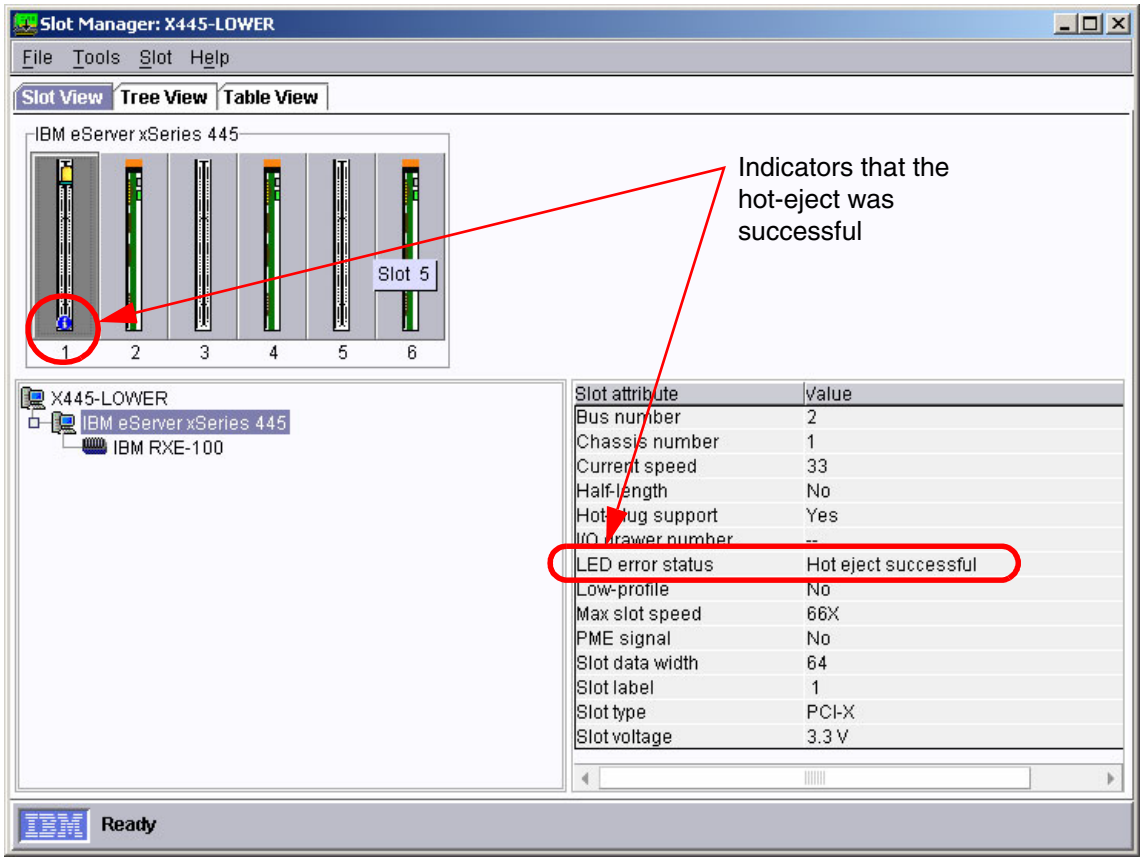

Figure 6-16 Hot eject successful

8. Close the adapter retainer on slot 1. The slot is then reported as empty.
9. Back at the Optimization Steps window, shown in Figure 6-13 on page 206, click the link to hot add the adapter.
10. The slot view appears with an arrow showing you which slot to insert the adapter into, as shown in Figure 6-17.

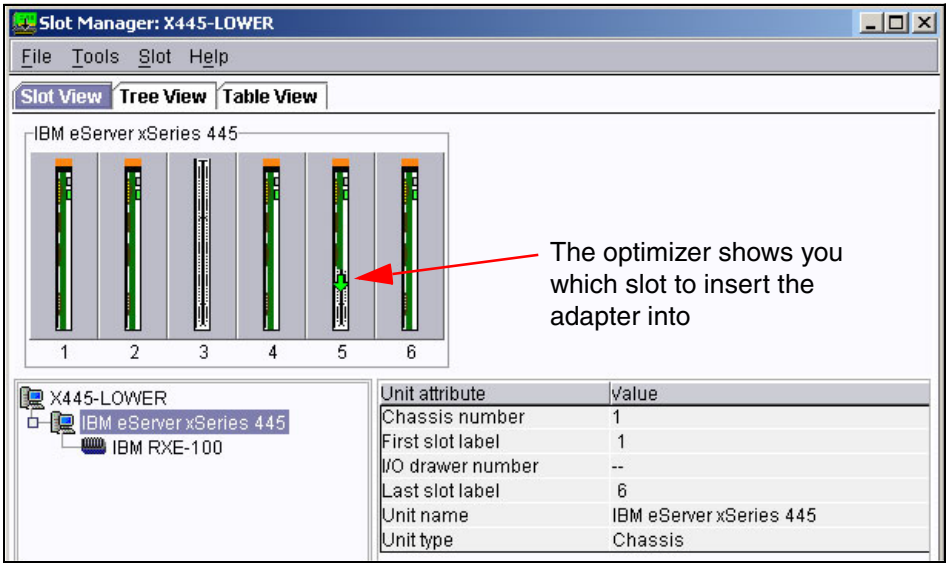

Figure 6-17 Indication of the slot for adding the adapter

11. Insert the adapter and close the latches.
12. Once all the recommended actions are complete, restart the server and re-run the optimizer to confirm that there are no performance issues.

By using the Slot Manager, you are assured that the I/O subsystem is configured for optimal performance. You can use this tool to assist you in configuring the system, or reconfiguring it at a later time.

6.2 Scalable Systems Manager

Scalable Systems Manager (SSM) is a plug-in to IBM Director that you can use to create and manage static hardware partitions.

Before you can use SSM on a 16-way x445 configuration, you will first need to cable the systems and configure static IP addresses on the Remote Supervisor Adapters. These steps are described in 5.5, “16-way x445 setup instructions” on page 179.

Tip: The steps in 5.5, “16-way x445 setup instructions” on page 179 describe how to configure the 16-way system completely without the use of SSM. You may consider following those steps, then, once the 16-way configuration is running, run SSM and import the 16-way configuration into SSM to further manage it.

6.2.1 Configuring the Remote Supervisor Adapters for SSM

SSM uses out-of-band connections to the Ethernet ports of the Remote Supervisor Adapter (RSA) for communication with the x445 systems. SSM will only be able to communicate with the x445 servers if the following conditions are met:

- The RSA in each server must have a static IP address assigned. Alternatively, if you cannot avoid using DHCP, you will need to configure DHCP reservations for the RSA.
- The IP addresses assigned to the RSA of each server must not change once the servers are discovered in IBM Director.
- All RSA alerts related to partitioning are sent to the IBM Director Server. You should only send alerts to one IBM Director management server.

To configure alerting, do the following:

1. From the RSA Web interface, click the **Alerts** item in the menu on the left.

Figure 6-18 Configuring Remote Alert Recipients

You can configure up to 12 recipients and specify various notification methods. Only send alerts to one IBM Director management server.

2. Click one of the unused positions. The panel shown in Figure 6-19 appears.

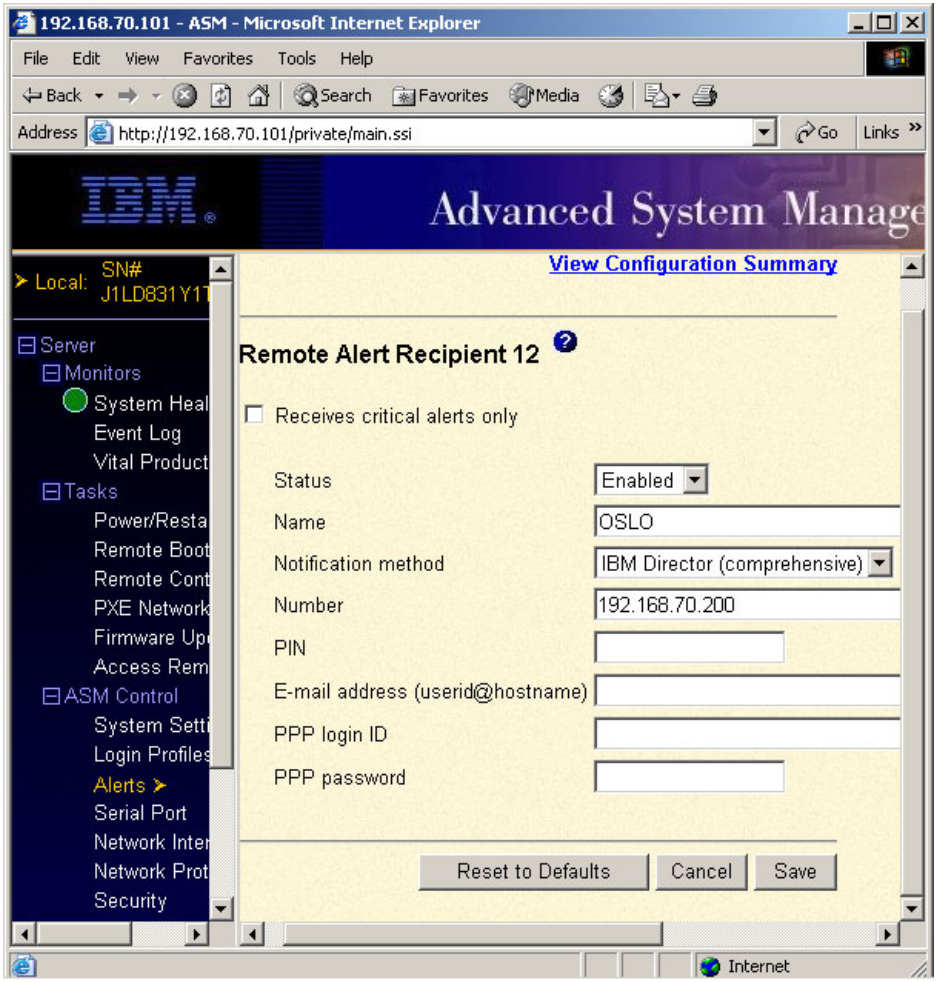

Figure 6-19 Remote Alert Recipient parameters

3. Enter the following parameters to enable alerts to be sent to the IBM Director Server.
   - Status: Select **Enabled**
   - Name: Type the name that will identify the recipient.
   - Notification method: Select **IBM Director (comprehensive)**
   - Number: Enter the IP address of the IBM Director Server

   The remaining fields are not used here. Selecting IBM Director (comprehensive) here results in *all* Remote Supervisor Adapter alerts will be sent to the IBM Director server.

4. Click **Save**.
5. You are returned to the Remote Alert Recipient window (Figure 6-18 on page 212). Scroll down to the bottom and click the **Save** button.

Tip: You do not need to check any of the Alert checkboxes because the notification method “IBM Director (comprehensive)” means that *all* alerts will be sent.

6.2.2 User interface

The SSM uses the following objects and icons:

Table 6-3 SSM objects and icons

Icon	Object description
	**Scalable node** is an individual x445 chassis. You can group up to two scalable nodes into a scalable system and then create scalable partitions on such a system.
	**Scalable partition** spans across one or two scalable nodes. It also includes an RXE-100 enclosure, if one is attached to the nodes. Scalable partition can run an instance of an operating system (it appears as a physical server to the operating system). When you create the scalable partition, you have to designate one of the nodes as the primary. This is required even if you use just one node.
	**Scalable system** contains one or two scalable nodes and scalable partitions.

Figure 6-20 shows the appearance of IBM Director Console with the SSM installed.

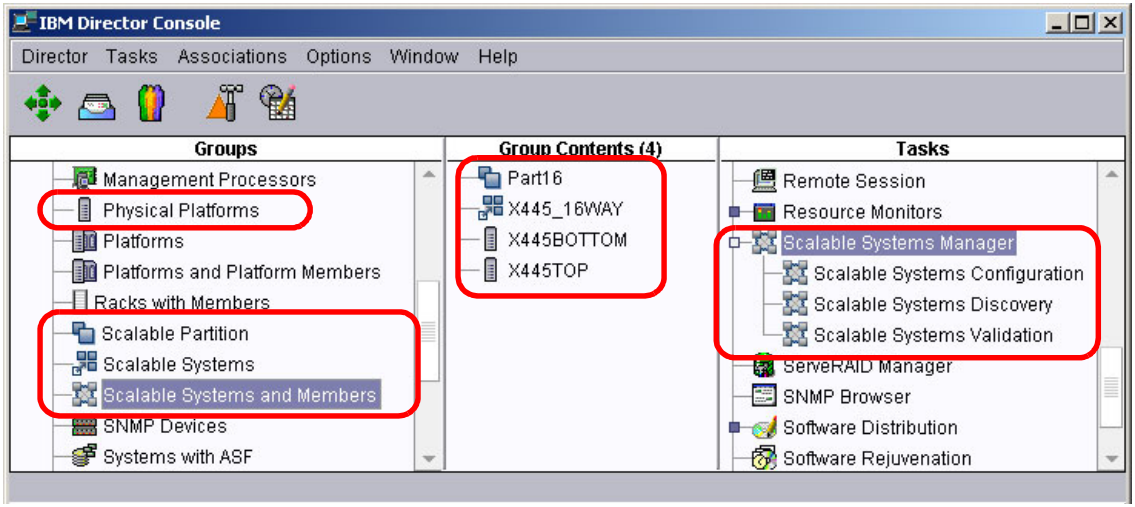

Figure 6-20 SSM groups, objects and tasks in IBM Director Console

There are four groups that are of relevance to SSM:

- Physical Platforms — the systems that are or could be a part of a scalable system. Not all Physical Platforms can be put into a Scalable System.
- Scalable Partition — existing scalable partitions.
- Scalable Systems — existing scalable systems.
- Scalable Systems and Members — This group only contains Physical Platforms and Remote I/O Enclosures that are part of a Scalable System. Physical Platforms (stand-alone) systems are not included in this group.

The middle pane of Figure 6-20 (Group Contents) shows the members of the Scalable Systems and Members group:

- A scalable partition named Part16
- A scalable system named X445_16WAY
- Two scalable nodes named X445BOTTOM and X445TOP

The right pane includes the SSM tasks. We discuss them in the following sections:

- “Scalable Systems Configuration task” on page 216
- “Scalable Systems Discovery task” on page 223
- “Scalable Systems Validation task” on page 224

We will now explain these three tasks.

Scalable Systems Configuration task

Figure 6-21 shows the Scalable Systems Configuration user interface. The window has several parts:

- **Topology** displays a tree-like structure of existing scalable systems and their scalable nodes and partitions.
- **Unassigned** shows nodes not yet assigned to any scalable system.
- **Details** and **Properties** show more information on the currently selected object in the topology tree.

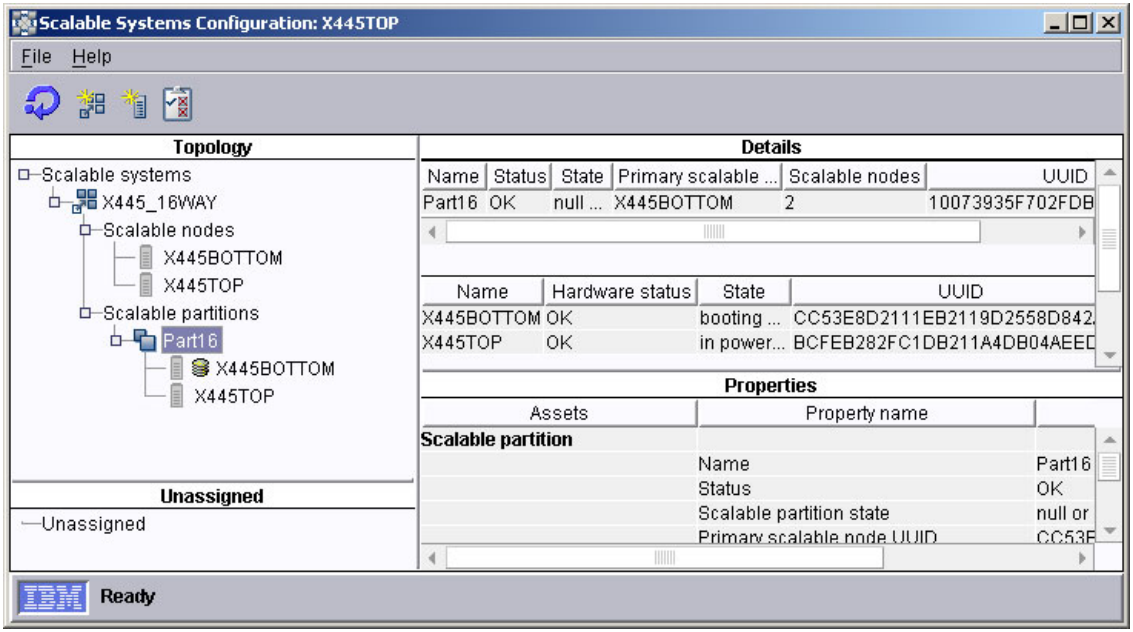

Figure 6-21 Scalable Systems Configuration

Figure 6-21 shows a scalable system that is already configured. The following sections explain the steps we took to configure the nodes.

Creating scalable nodes

You can use the IBM Director discovery capability to identify the scalable nodes. Alternatively, you can create the nodes manually using the Scalable Systems Configuration window. You do this by clicking on the Create Scalable Node icon, as indicated in Figure 6-22.

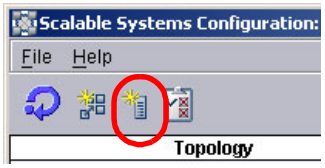

Figure 6-22 Scalable node creation icon

SSM prompts (Figure 6-23) you for the name of the new scalable node, the RSA host name or IP address, and the RSA user ID and password.

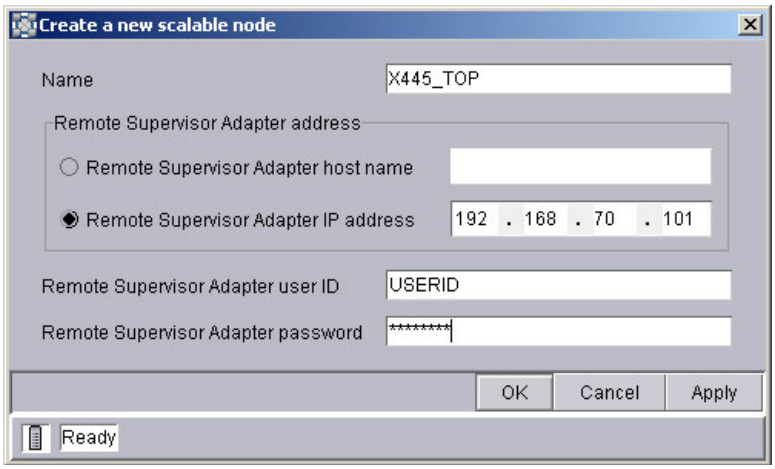

Figure 6-23 Creating a new scalable node

After performing this step for both our systems, X445_TOP and X445_BOTTOM, they appear in the Scalable Systems Configuration window as unassigned nodes. Figure 6-24 shows this.

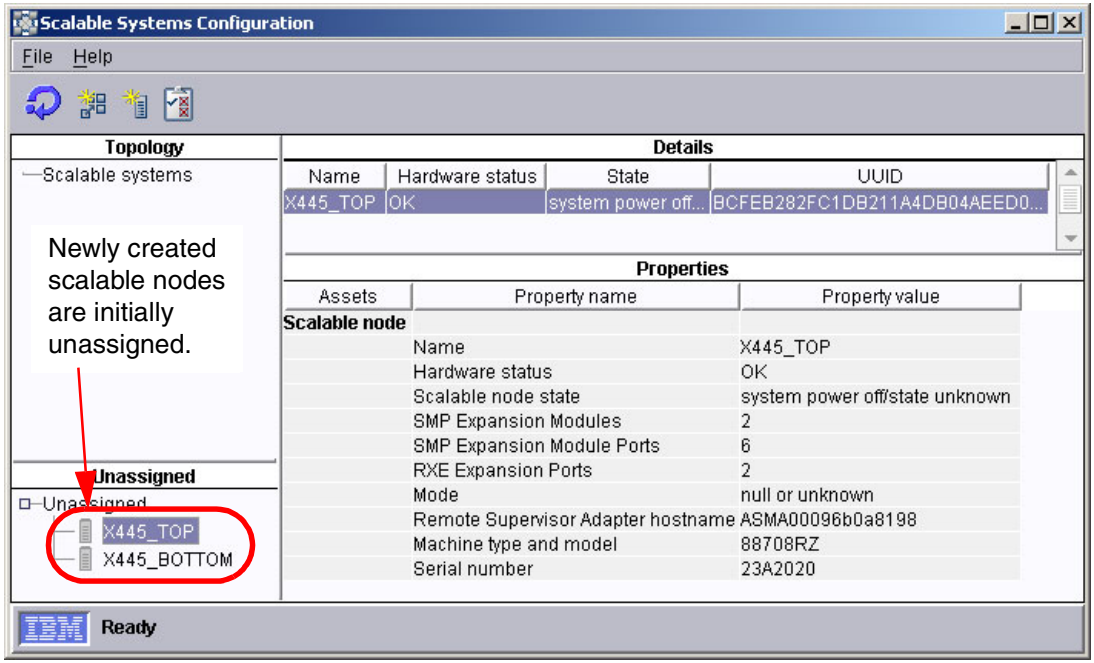

Figure 6-24 Unassigned scalable nodes

Creating a scalable system

Now that the unassigned nodes are created and available, the next step is to create a scalable system and add the nodes to it. Launch the wizard by clicking on the icon indicated in Figure 6-25.

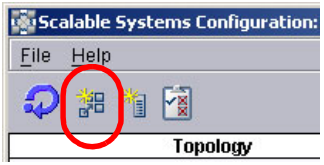

Figure 6-25 Scalable system creation icon

You will have to specify the name of the scalable system (we selected X445_16WAY) and the next step is to add the scalable nodes, as shown in Figure 6-26.

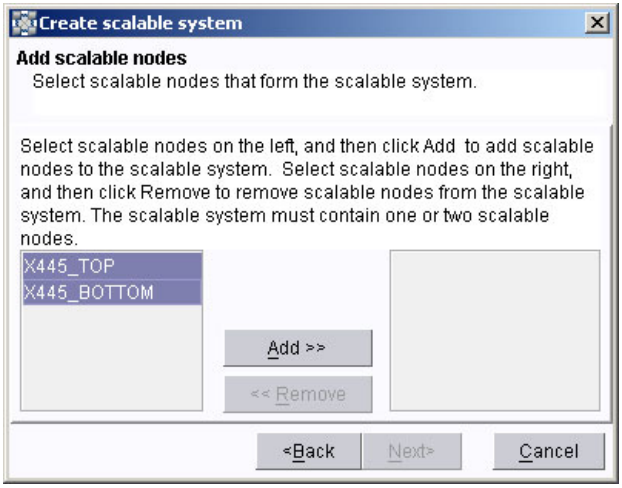

Figure 6-26 Adding the nodes

Figure 6-27 displays the summary window, which shows the diagram and the list of scalable nodes in this new scalable system.

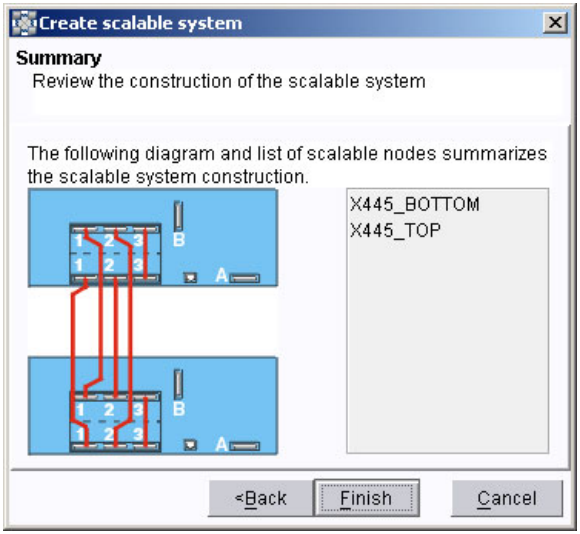

Figure 6-27 Summary review of the new scalable system

The new scalable system X445_16WAY appears in the topology tree structure, as shown in Figure 6-28. However, there are still no scalable partitions defined. We do this in the next section.

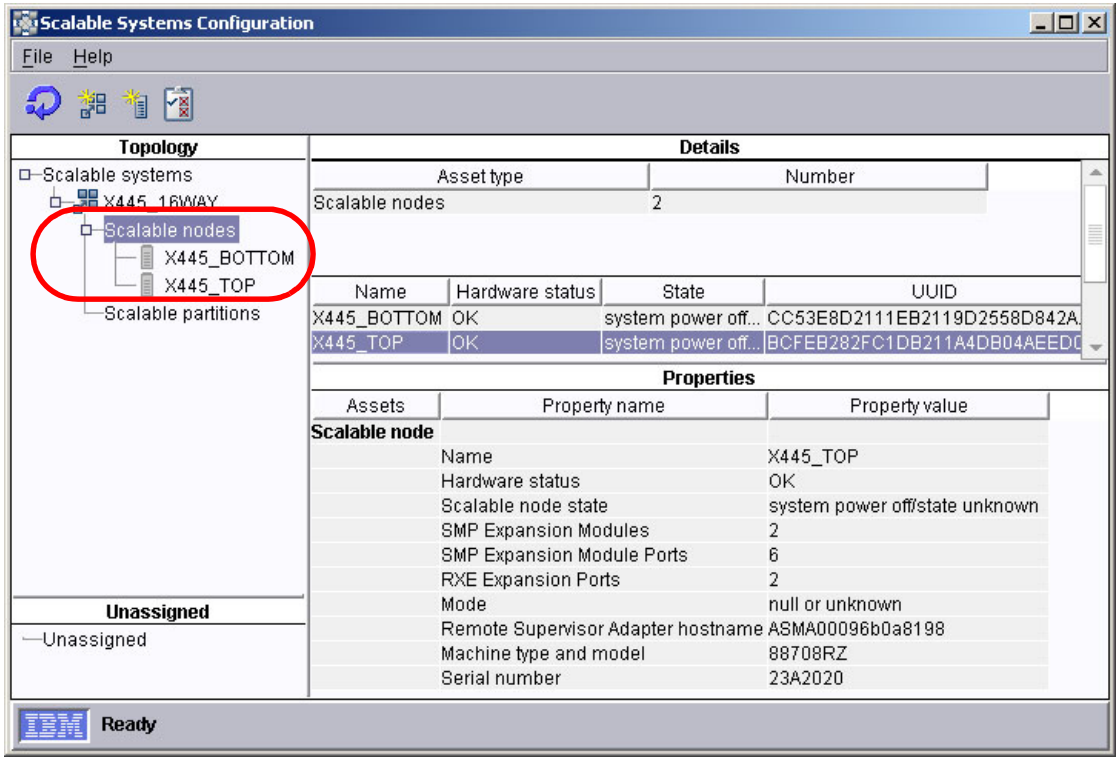

Figure 6-28 New scalable system

Creating a scalable partition

To create scalable partition do the following:

1. Right-click the **Scalable partitions** object in the Topology tree and click **New scalable partition**. You will be prompted for the name of the new scalable partition.
2. Name the partition. We used PART16 for the name. The new partition appears in the topology tree. We selected PART16 for the name.
3. Now add the scalable nodes to the partition. Right-click the partition name (for us, **PART16**) and click **Insert scalable node** → ***node***, where *node* is each of the nodes you want to add (see Figure 6-29). We need to perform this step for both nodes.

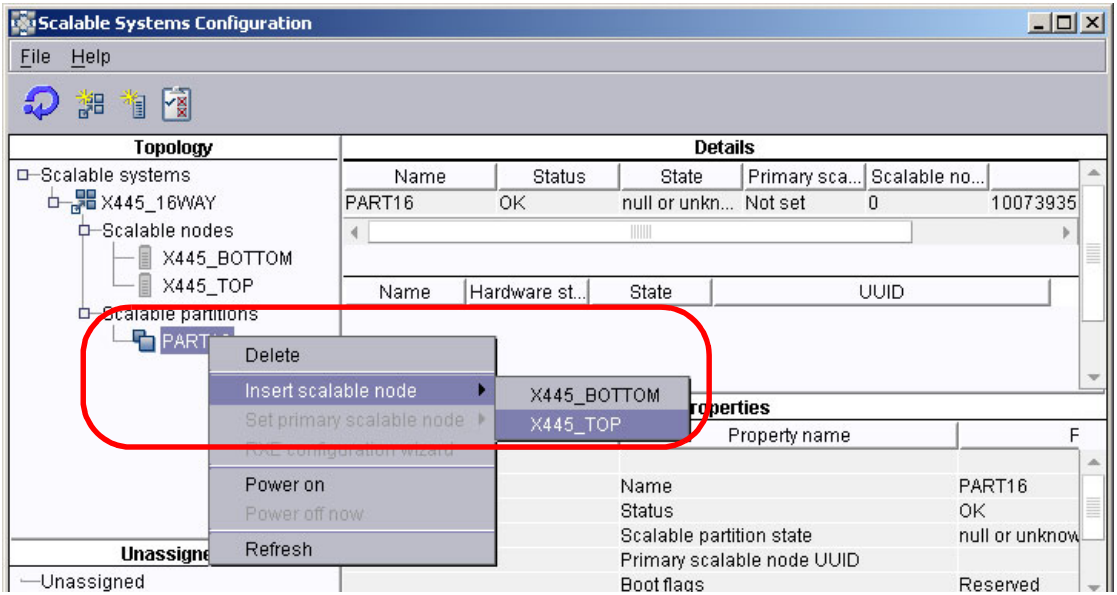

Figure 6-29 Inserting the scalable nodes

4. Define which is the primary node in the partition. The primary node is the one that will boot the operating system. It also will have the access to the local diskette and the DVD drive.

   Right-click the node that you want to be the primary, and set its role appropriately, using the pop-up menu shown in Figure 6-30.

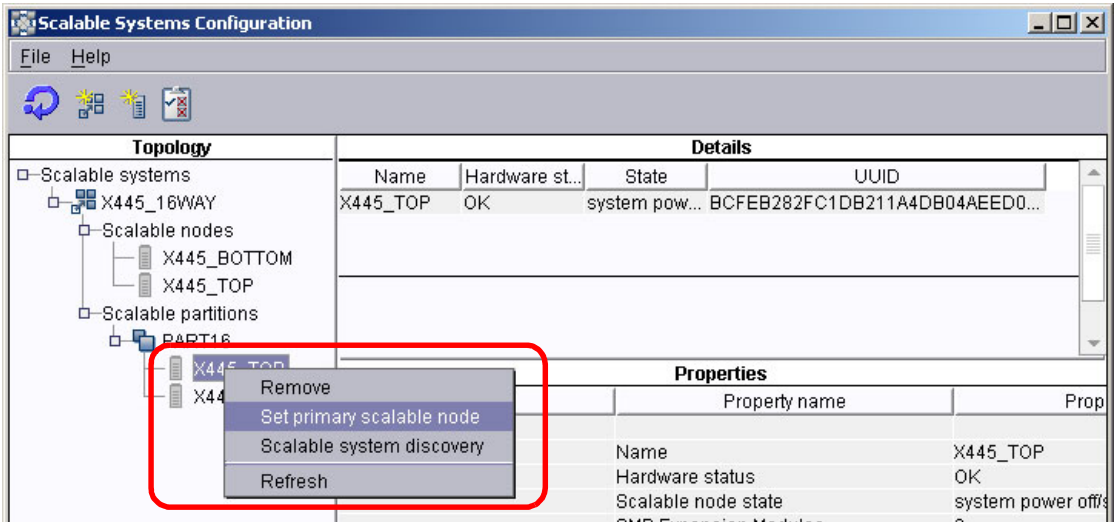

Figure 6-30 Setting the primary scalable node

Now the scalable partition PART16 is configured and, as you can see in Figure 6-31, the hard disk drive icon ( ) next to one of the nodes indicates the primary node.

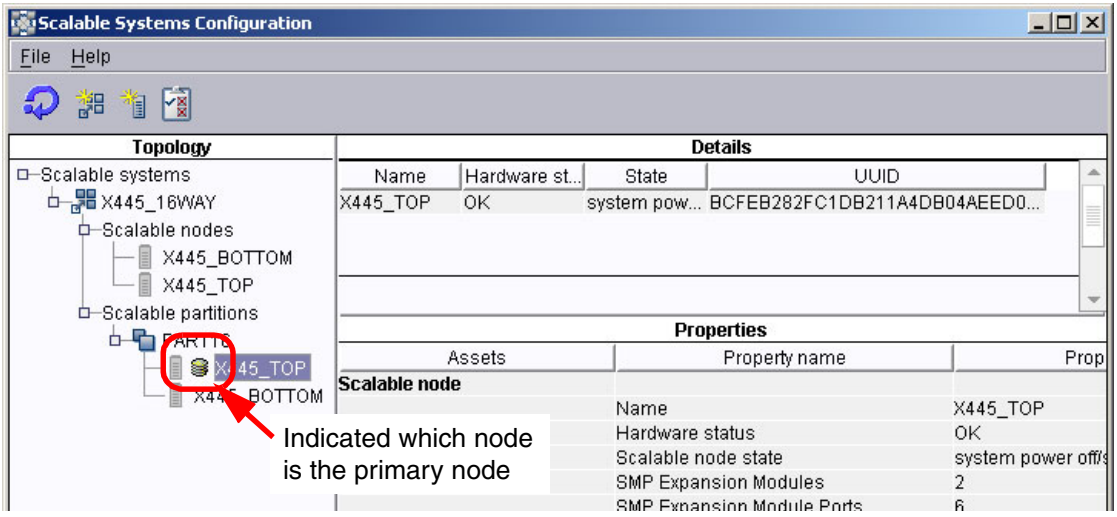

Figure 6-31 The primary scalable node

Scalable Systems Discovery task

The Scalable Systems Discovery process checks the unassigned scalable nodes and tries to determine how they are cabled together. It does this by issuing ping commands to the scalability ports. The results of these queries will lead to one of the following automatic actions:

- The discovery process automatically creates a new scalable system for each pair of unassigned nodes that it identifies as cabled together.
- It also does the same for each unassigned node that is not cabled to any other node.
- The process also verifies the configurations defined through the BIOS Setup utility (see 5.5, “16-way x445 setup instructions” on page 179), to see if all the defined nodes are present. It verifies how many nodes are cabled together (by applying pings to the scalability ports), and then compares this with the number of nodes in the partition descriptor (PD) definition.
- The process also verifies the existing scalable systems, to see if all the defined nodes are present. It verifies how many nodes are cabled together (by applying pings to the scalability ports), and then compares this with the number of nodes in the scalable system definition.

The process can be time consuming, especially in an environment with a lot of unassigned nodes and with complex cabling between the nodes. For better performance, the process does not issue the ping command to every port; it uses only enough pings to understand the cabling of the nodes. If you want to check all the ports and cables, use the validation process, described in “Scalable Systems Validation task” on page 224.

To run the discovery process, do the following:

1. Select the scalable nodes and systems that you want to participate in the process.
2. Drag and drop the Scalable Systems Discovery task to the selected nodes and systems. Alternatively, right-click on the selected nodes and systems and select the **Scalable Systems Discovery** task in the menu.

Tip: The discovery task will only run on unlocked nodes. Therefore, if any of the nodes are locked, you will have to request access to them first.

When you run the discovery task, a window will open, showing you the progress of the process on each node. The lower part of the window will list the nodes that either completed or failed the discovery. Some nodes will also be skipped:

- Unsupported nodes (for example, the x440).
- Scalable nodes already assigned to a scalable system.

- Scalable systems that are already completely configured.

After the discovery process completes, you can use the Scalable Systems Configuration task to create scalable partitions on the successfully discovered scalable systems.

Scalable Systems Validation task

This task validates the wiring between the nodes in a scalable system, which can prove extremely useful. The validation task is an implementation of Copper Diagnostics, a feature of Enterprise X-Architecture.

You can start the task by dragging-and-dropping it to the scalable system object. This will open a window shown in Figure 6-32.

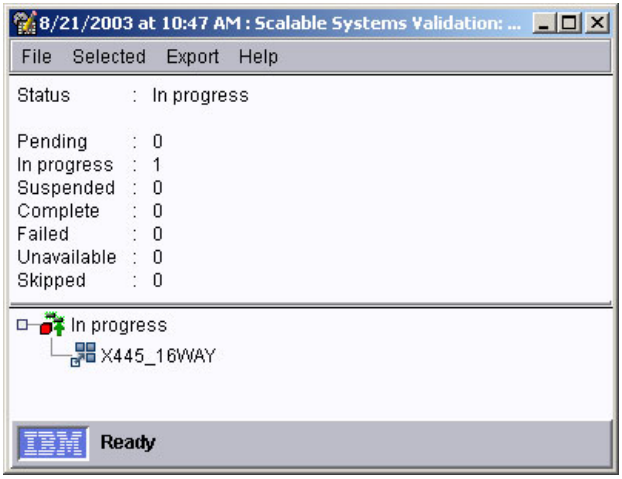

Figure 6-32 Scalable Systems Validation window

You can watch the progress of the task and see the final result: complete or failed. However, for more details select **File** → **View Log**. This action will open another window, Figure 6-33.

This window shows the results of each step of the validation process as it progresses:

1. The validation process begins by issuing a ping from each scalability port on each node. Each SMP Expansion Module contains three scalability ports, so each node has six ports. If all ports can successfully issue a ping, the node will pass the test.
2. Next step is to check for the response to each ping. This test can actually verify if the 16-way wiring between the SMP Expansion Module ports is

correct. The UUID of the node is sent as a reply to the ping. If all UUIDs are successfully returned and correct, the wiring is successfully verified.

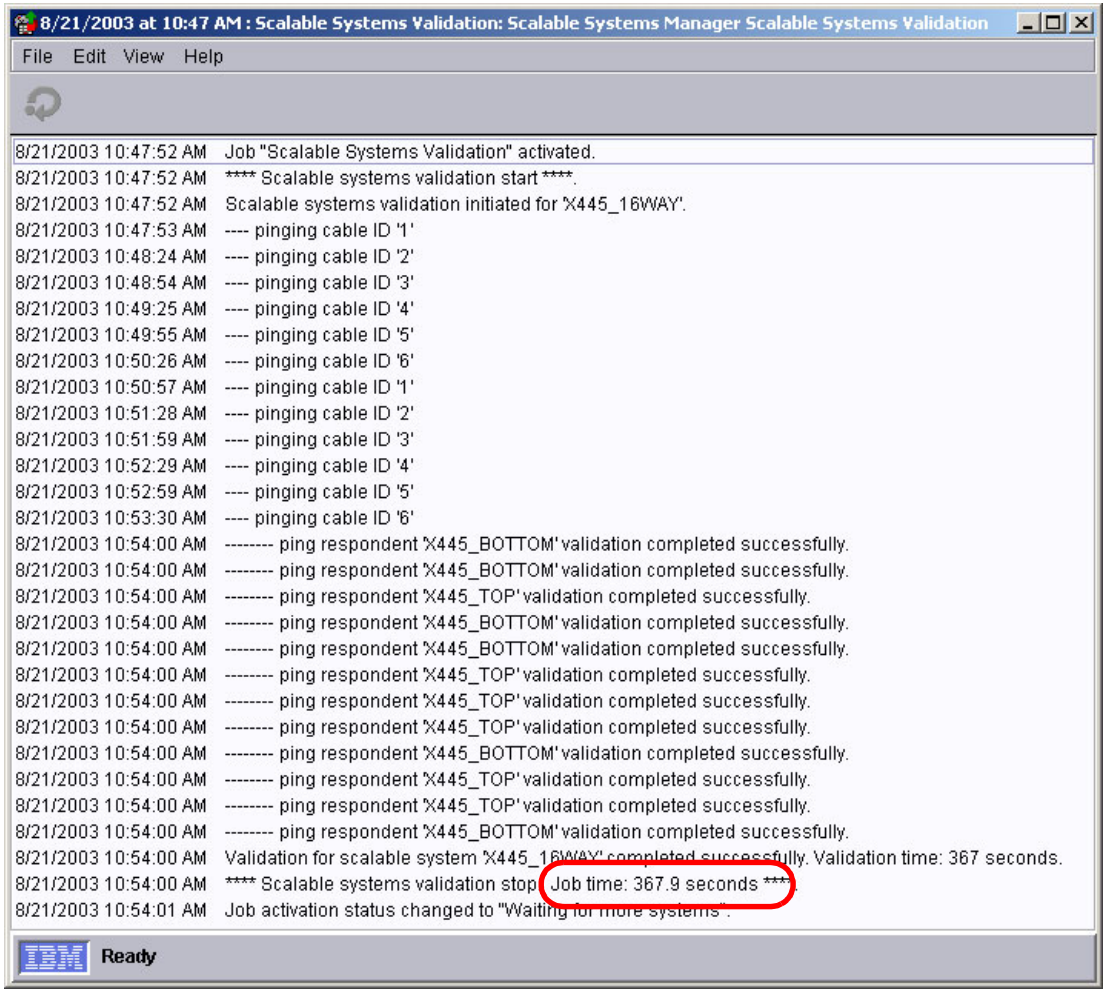

Figure 6-33 Scalable System Validation - log window (Note the task took 6 minutes to run)

The whole validation process usually takes around six to seven minutes, as you can see from the reported job time in the log window.

We recommend you to customize the display of the validation log. First, check the **Dynamic Update** option, as seen in Figure 6-34. The new log messages will show up automatically, without the need to manually refresh the window display.

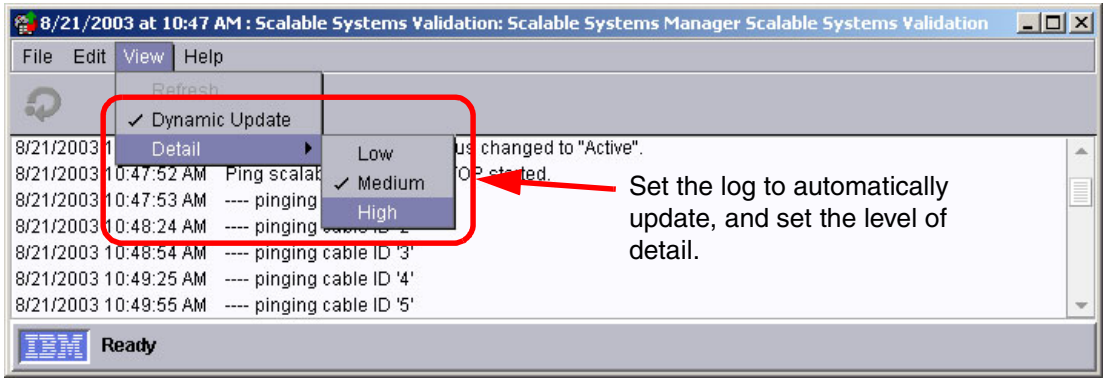

Figure 6-34 Setting the validation log view options

You can also choose the level of details to display. The three options are **Low**, **Medium** and **High**. Figure 6-33 shows low level of details, while the high level display is shown in Figure 6-35.

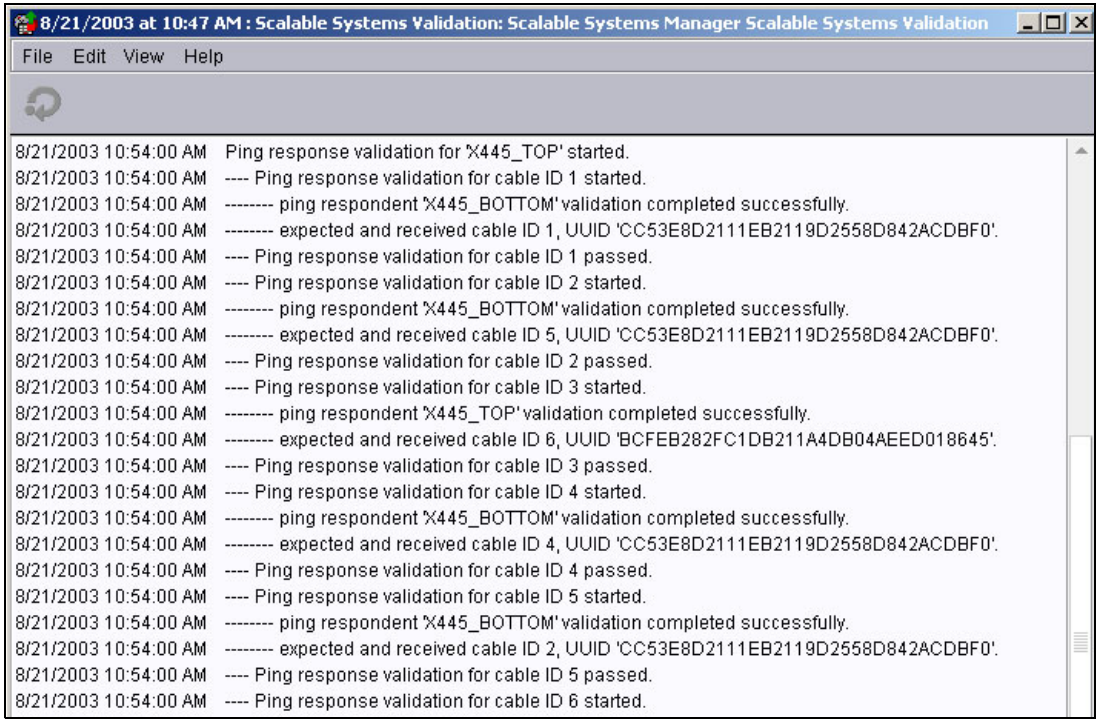

Figure 6-35 Validation log - highly detailed display

Note the level of ping response information: we can see which node responded, using which SMP Expansion Module port and also what the UUID was.

As you can see, the validation task can be very useful in troubleshooting scenarios, whenever you are experiencing connectivity issues with the 16-way systems. The validation log will tell you exactly which ports have communication problems.

Be aware that the communication faults discovered by the validation process do not necessarily mean the scalability ports are at fault. For example, if the RSA remote alert recipients are not configured correctly, the validation task ping between some scalability ports will fail, even though the ports might be fully functional.

We therefore also recommend to use the scalability ports test of the Server Diagnostics (F2 at boot time), which will identify whether the ports are faulty or not.

6.3 Application Workload Manager

Application Workload Manager (AWM) is a very useful tool for application resource usage management and control. It is a fee-based plug-in to IBM Director 4.1 and it works with applications running on Windows 2000 or Windows Server 2003.

AWM gives you the ability to control the allocation of CPU time and memory resources to the applications. This can be especially important in server consolidation scenarios, where we consolidate a number of servers onto a smaller number of more capable machines, such as the x445.The applications that likely had exclusive access to the resources before consolidation took place now run on the same server, competing for the CPU and memory.

AWM provides much better control of the resource allocation than the operating system utilities. With proper use, you can create suitable workspaces for each application, prevent them from consuming too much of the resources, and also ensure that the applications will not experience lack of resources.

The AWM plug-in contains the following components:

- AWM Scheduler

  This is the core resource management service that you have to install on each server you want to manage. This is not an IBM-Director component: it is the scheduler that will actually handle the resource allocation within the defined parameters. Other components are basically an interface between the user, the IBM Director and this service.

- AWM Agent

  This component is the interface between the AWM scheduler and the IBM Director Agent. You have to install it on the same system as the scheduler and the Director Agent.

- AWM Server

  This is the component that must be installed on the IBM Director Server machine. It provides communication between the AWM Console and the AWM Agent, and also integrates the AWM into the IBM Director environment.

- AWM Console

  The AWM Console is the user interface for managing the resource allocation policies, and integrates into the IBM Director Console, which must be installed on the same machine.

You can use the AWM to optimize the usage of the following resources:

- CPU
- Real (physical) memory
- Virtual memory

6.3.1 Using AWM

After installation, AWM will appear as a task in the IBM Director Console, as shown in Figure 6-36. Run the task by dragging and dropping it onto the desired system.

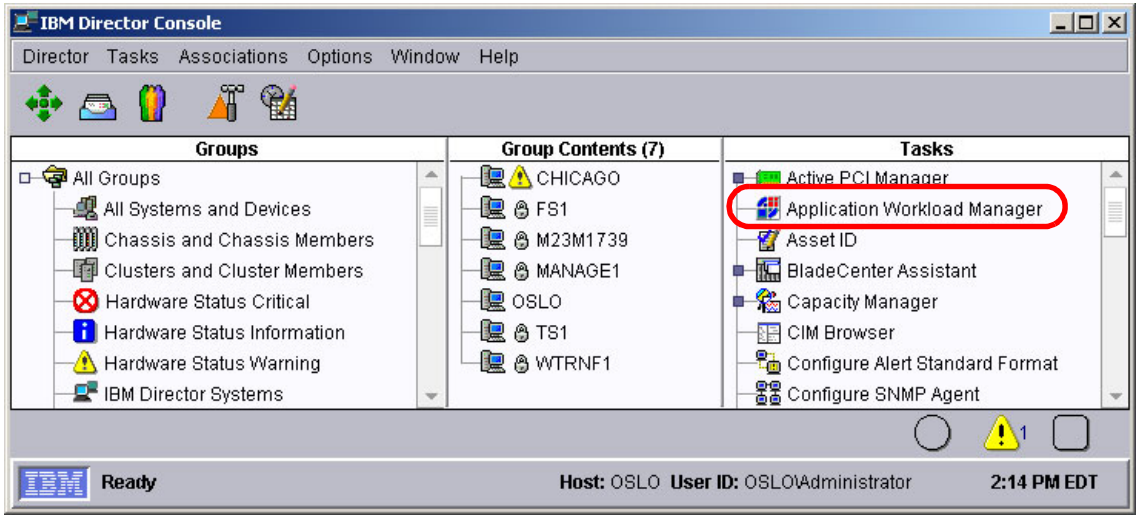

Figure 6-36 The Application Workload Manager task

When the task starts, it will open the AWM console and display the name of the system (OSLO in our case), current CPU usage and a list of applications defined on that system. The Default Application contains all processes that have not been attached to any other application.

Figure 6-37 shows the AWM console. Note the CPU share column: you can use these CPU share values to optimize the way the applications use the CPU. We will explain this in the following sections.

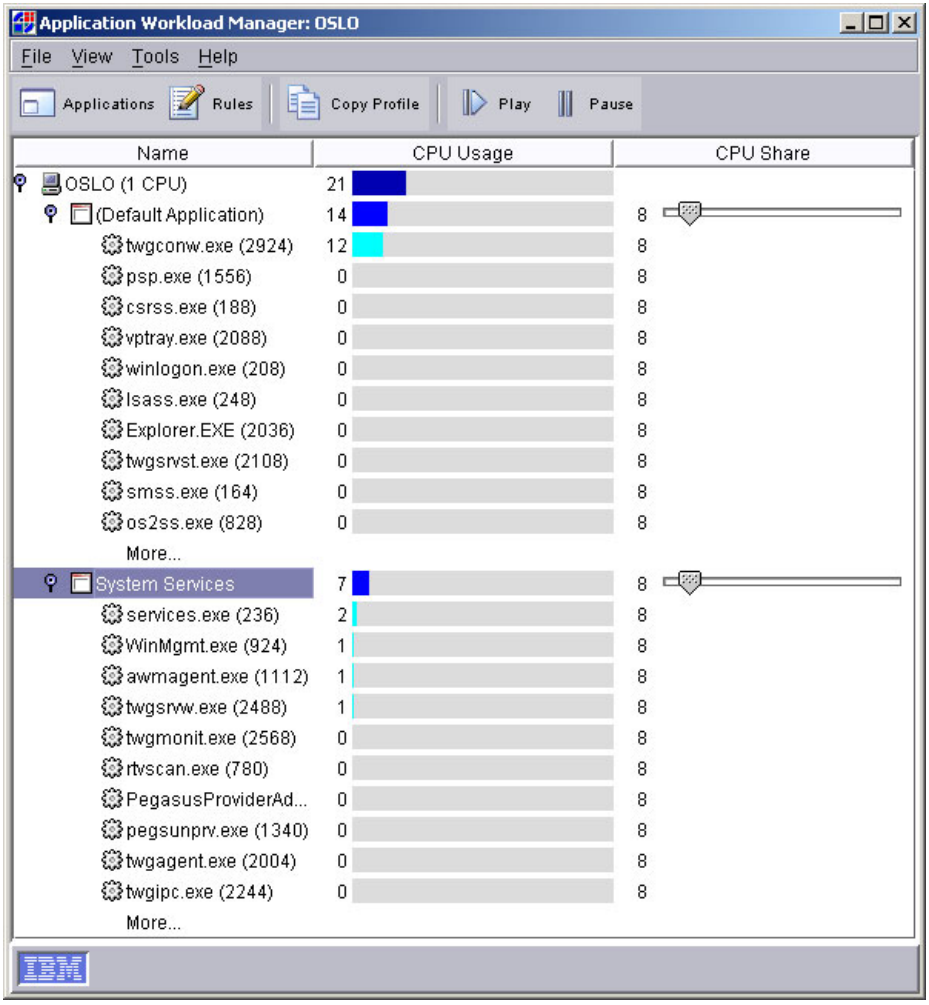

Figure 6-37 AWM Console

Besides the CPU share, it is also possible to display other columns. Select **Tools** → **Options** and the window that opens will let you configure the display of

the AWM console. The window has two tabs: General and Column. Use the Column tab to configure the visibility of various columns, as shown in Figure 6-38. The columns relate to the CPU, Real and Virtual Memory resource settings and usage.

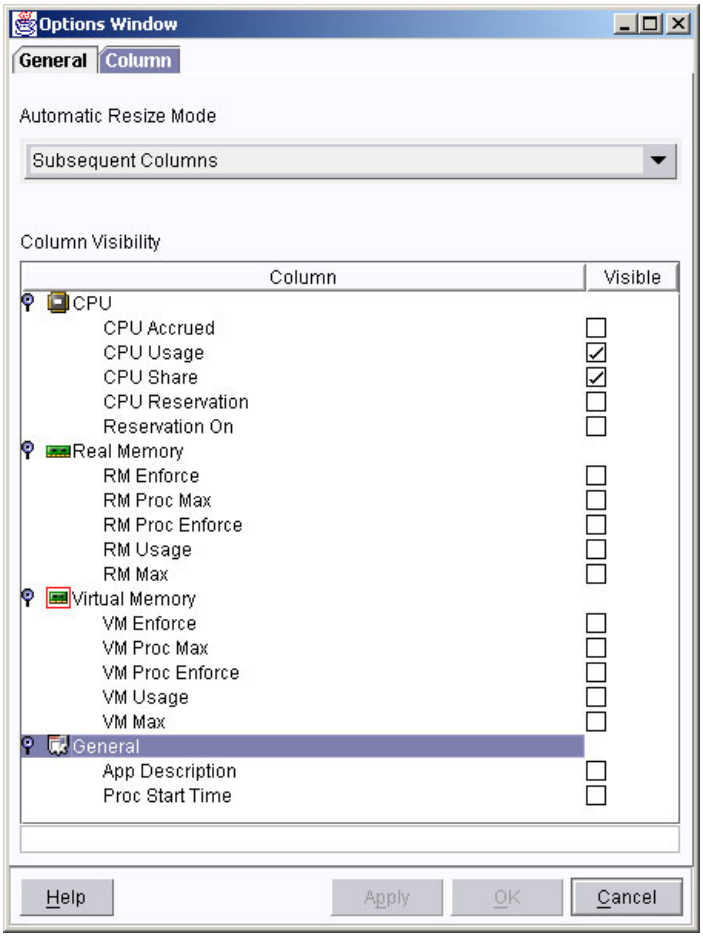

Figure 6-38 Configuring the visibility of columns

Working with applications

To be able to work with specific applications, you will need to add their application definitions. The next step will be to set up the rules that will link individual processes to these applications.

To work with application settings, click on the **Applications** icon in the icon bar (see Figure 6-37). The Applications window opens, showing a list of defined applications. You can add new definitions, and delete or edit the existing ones.

Adding an application will launch the Application Wizard. Follow the steps to complete the application definition.

Editing the existing application opens the Application Properties window. The window has four tabs: General, CPU, Virtual Memory and Real Memory. Figure 6-39 shows the CPU tab page.

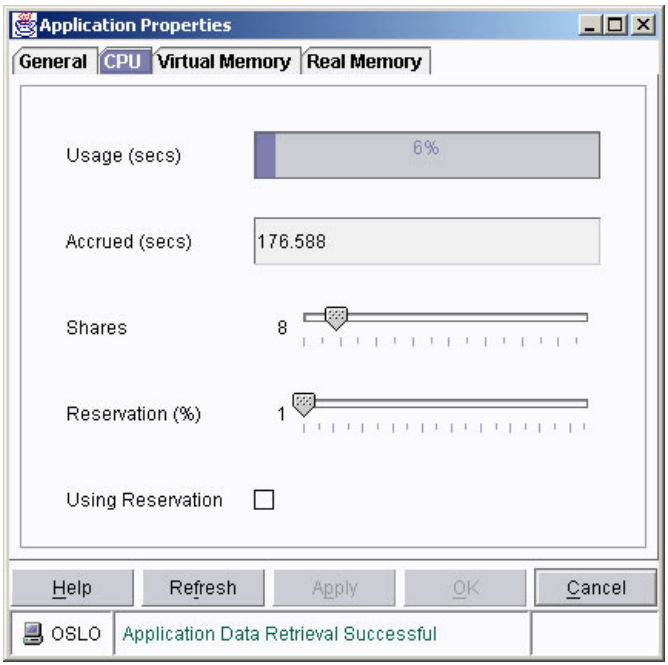

Figure 6-39 Editing the Application Properties - CPU allocation

CPU allocation control

As you can see, you can control the CPU time allocation in two ways:

- Using shares

  Shares are dynamic and relative CPU allocations. The default number of shares for each application is 8. To allow more CPU usage for a certain application, you have to increase its share value.

  The share value you associate with an application is the maximum CPU time it can have, regardless of whether or not there is more CPU capacity available.

  The total number of shares is not constant. Generally, it increases with each new application definition added. For example, if you define four applications and leave the default share value of eight for each, the total number of shares

will be 32. Each application will be entitled to 8 shares out of the total 32 (that is, 25%) of CPU.

If you increase the shares for one of the applications to 16, the total shares number will increase to 40. The prioritized application will be entitled to 40% of CPU (16 shares out of 40), and the other three applications will be entitled to 20% each (8 shares out of 40).

Because shares use relative allocation, you can use them in cases when you need to ensure the applications will never starve of the CPU usage.

- Using reservation

  This is another approach to allocating the CPU resources to the applications. Reservation is expressed in percentage and is the absolute minimum of the CPU allocated to a certain application. Whenever the application requires this minimum, it will get it. It can also use more CPU than the specified percentage, when available. On the other hand, when the application does not need the specified percentage of CPU, it will be allocated to other processes.

  Use the Reservation to assure the CPU resources to the vital, mission-critical applications.

  When using reservations, remember that the sum of all reservations must be below 99%. This will ensure certain CPU allocation to the essential operating system processes.

The applications can be set up with a mixture of shares and reservations. In such case, reservation takes precedence over shares.

Memory allocation control

Figure 6-40 shows the Virtual Memory tab page. The page for Real Memory is very similar. You can specify the Virtual or Real Memory Maximum to restrict the memory available to the application. These maximums can be useful as a protection from memory leaks, however you need to be very cautious. If a Real Memory maximum is set too low, then this can cause a high amount of paging. This will significantly impact the performance. Similarly, if you set the Virtual Memory Maximum too low, the process that reaches it might terminate or crash.

Memory Maximums should normally not be used. We recommend to use them only in special cases, for example when dealing with memory leaks and seeking ways to prevent the crash of the entire operating system.

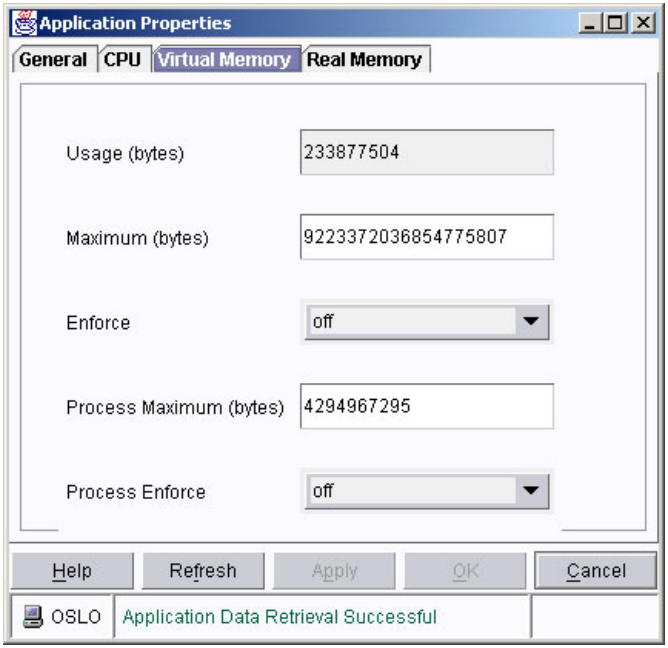

Figure 6-40 Editing the application properties - Virtual Memory allocation

Working with rules

So far, we configured the applications and defined the resource usage for them. Now we need to link the individual processes to these applications. The AWM uses *rules* to determine these links.

To set up the rules, click the Rules icon in the console:

Figure 6-41 The Rules icon

A window opens, showing a list of configured rules. You can add new rules, which launches the Rule Wizard. Alternatively, you can also clone the existing rules and edit them. If you choose to edit the rule, the Rule Properties window opens, as shown in Figure 6-42.

This window has five tabs, General, Program, Library, Command, and Parent, that allow you to specify the relationship between the processes and the applications.

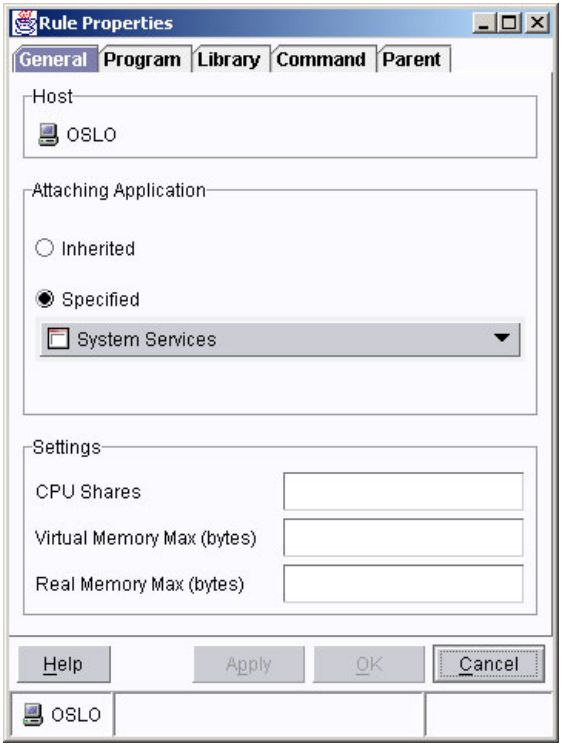

Figure 6-42 The Rule Properties window

The processes which do not have any rules specified to link them to any of the defined applications will act as follows:

If their parent process is linked to a certain application, they will also be linked to the same application. Their CPU, Virtual and Real Memory resources will be controlled by the Shares, Reservation and Memory Maximum settings for that application.

If their parent process is not linked by rules to any of the defined applications, they will be linked to the Default Application.

During the installation of AWM, an application named “System Services” is created automatically. A rule to attach the process SERVICES.EXE to this application is also set up. Therefore, if you keep the default rule, all the processes spawned by the SERVICES.EXE will be linked to the “System Services” application.

Abbreviations and acronyms

ACPI	Advanced Configuration and Power Interface
ADMT	Active Directory Migration Tool
APIC	advanced programmable interrupt controller
AS	advanced server
ASM	advanced system management
ASP	Application Service Provider
AWM	Application Workload Manager
BASP	Broadcom Advanced Server Program
BI	business intelligence
BIOS	basic input/output system
CAL	client access license
CAT	Configuration Audit Test
CDAT	Consolidation Discovery and Analysis Toolset
CEC	Central Electronics Complex
CEE	Central and Eastern Europe
CLR	common language runtime
CMAK	Connection Manager Administration Kit
CNA	Certified Novell Administrator
CPU	central processing unit
CRM	Customer Relationship Management
DC	Domain Controller, Datacenter
DDR	Double Data Rate
DDS	Digital Data Storage
DHCP	Dynamic Host Configuration Protocol
DIMM	dual inline memory module
DIV	Datacenter Infrastructure Vendor
DLT	Digital Linear Tape
DOS	disk operating system
DSP	Datacenter Service Provider
ECC	error checking and correcting
EFS	Encrypting File System
EIDE	Enhanced Integrated Drive Electronics
EMEA	Europe, Middle East, Africa
EQP	Enterprise Qualification Program
ER	enterprise rack
ERP	Enterprise Resource Planning
ESMT	Enterprise Services for Microsoft Technologies
ESS	Enterprise Storage Server
EXA	Enterprise X-Architecture
FC	fibre channel
FRS	File Replication Service
FRU	field replaceable unit
FTMI	Fault Tolerant Management Interface
HA	high availability
HAL	hardware abstraction layer
HARQ	High Availability Resolution Queue
HBA	host bus adapter
HCL	Hardware Compatibility List
HCT	Hardware Compatibility Test
I/O	input/output
IAS	Internet Authentication Service
IBS	International Business Systems
ICF	Internet Connection Firewall
ICS	Internet Connection Sharing
IE	Internet Explorer
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IIS	Internet Information Services

ILS Internet Locator Service
IPX Internetwork Packet eXchange
ISP Internet Service Provider
ISV Independent Software Vendor
JSQ Joint Support Queue
LAN local area network
LDAP Lightweight Directory Access Protocol
LED light emitting diode
LTO Linear Tape Open
MAPS Microsoft Authorized Premier Support
MBR master boot record
MCP Microsoft Certified Professional
MMC Microsoft Management Console
MMS Metadirectory Services Support
MSCS Microsoft Cluster Services
NCMCON Novell Configuration Manager Console
NEMA National Electrical Manufacturers Association
NIC network interface card
NLB network load balancing
NTFS NT File System
NUMA Non-Uniform Memory Architecture
OEM original equipment manufacturer
OLAP Online Analytical Processing
OSPF Open Shortest Path First
PAE Physical Address Extension
PCD Personal Computer Division
PCI Peripheral Component Interconnect
PD partition descriptor
PFA Predictive Failure Analysis
PME power management events
POST power-on self test
PPTP Point-to-Point Tunneling Protocol
RAC Oracle Real Application Cluster
RADIUS Remote Authentication Dial-in User Service
RAID redundant array of independent disks
RAM random access memory
RETAIN Remote Technical Assistance Information Network
RFC request for comments
RHN Red Hat network
RISC reduced instruction set computer
RPM revolutions per minute
RSA Remote Supervisor Adapter
RXE Remote eXpansion Enclosure
SAK Server Appliance Kit
SAN Storage Area Network
SAR Solution Assurance Review
SCM supply chain management
SCON server consolidation
SCSI small computer system interface
SDRAM synchronous dynamic random access memory
SLES SuSE Linux Enterprise Server
SMB small-to-medium business
SMP symmetric multiprocessing
SMTP Simple Mail Transfer Protocol
SNMP Simple Network Management Protocol
SP service processor
SQL structured query language
SR standard rack
SRAT Static Resource Affinity Table
SRPM Source Red Hat Package Manager
SSM Scalable Systems Manager
SVGA super video graphics array
TAR tape archive
TCO total cost of ownership
TCP/IP Transmission Control Protocol/Internet Protocol
UPS uninterruptable power supply
URL Uniform Resource Locator

USB	Universal Serial Bus
UUID	universally unique identifier
VDS	Virtual Disk Service
VHDCI	very high density connector interface
VLAN	virtual local area network
VM	virtual machine
VMFS	virtual machine file system
VPN	Virtual Private Network
VRM	voltage regulator module
VSS	Volume Shadow-Copy Service
WAN	wide area network
WHQL	Windows Hardware Quality Labs
WLBS	Windows Load Balancing Service
WMI	Windows Management Instrumentation
WSRM	Windows System Resource Manager
ZIF	zero insertion force

Related publications

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

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 244. Note that some of the documents referenced here may be available in softcopy only.

- *Windows Server 2003, Datacenter Edition on the IBM @server xSeries 445*, REDP3700
- *Netfinity Tape Solutions*, SG24-5218
- *Tuning IBM @server xSeries Servers for Performance*, SG24-5287
- *Integrating IBM Director with Enterprise Management Solutions*, SG24-5388
- *S/390 Server Consolidation - A Guide for IT Managers*, SG24-5600
- *The IBM LTO Ultrium Tape Libraries Guide*, SG24-5946
- *Implementing Systems Management Solutions using IBM Director*, SG24-6188
- *Server Consolidation with the IBM @server xSeries 440 and VMware ESX Server*, SG24-6852

Other publications

These publications are also relevant as further information sources. These are available from:

`http://www.pc.ibm.com/support`

- *IBM @server xSeries 445 Installation Guide*
- *IBM @server xSeries 445 Option Installation Guide*
- *IBM @server xSeries 445 User's Guide*
- *IBM @server xSeries 445 Hardware Maintenance Manual*
- *IBM @server xSeries 445 Rack Installation Guide*
- *IBM Remote Expansion Enclosure Installation Guide*

Online resources

These Web sites and URLs are referenced in this redbook:

IBM product information

- Enterprise Storage Server:
  http://www.storage.ibm.com/disk/ess/index.html
- ServeRAID adapter comparison:
  http://www.redbooks.ibm.com/abstracts/tips0054.html
- Configurators: http://www.pc.ibm.com/europe/configurators and
  http://www.pc.ibm.com/us/eserver/xseries/library/configtools.html
- Solution Assurance: http://w3.ibm.com/support/assure and
  http://www.ibm.com/partnerworld/techsupport
- SCSI product information:
  http://ibm.com/pc/ww/eserver/xseries/scsi_raid.html
- FAStT Fibre Channel information:
  http://www.storage.ibm.com/disk/fastt/index.html
- Tape product information:
  http://ibm.com/pc/ww/eserver/xseries/tape.html
- VMware information (IBM employees only)
  http://ats.greenock.uk.ibm.com/vmware_new.htm
- VMware information http://www.ibm.com/eserver/xseries/vmware
- IBM Product Customization Services
  http://www.pc.ibm.com/us/eserver/xseries/systems_management/pcs.html
- IBM Datacenter Solution Program
  http://www.pc.ibm.com/ww/eserver/xseries/windows/datacenter.html
- IBM Advanced Support offering
  http://www.ibm.com/services/its/us/mus89d1.html

IBM support information

- xSeries support home http://www.pc.ibm.com/support
- xSeries operating system installation instructions (IBM employees only)
  http://nosinfo.raleigh.ibm.com
- ServerProven home page: http://www.pc.ibm.com/compat
- ServerProven LAN adapter support matrix:
  http://www.pc.ibm.com/us/compat/lan/matrix.html

- ServerProven operating system support matrix:
  http://www.pc.ibm.com/us/compat/nos/matrix.shtml
- ServerProven x445 support matrix:
  http://www.pc.ibm.com/us/compat/machines/x445.html

Operating system installation instructions

- Complete list of OS installation instructions for the x445:

  http://www.ibm.com/pc/support/site.wss/search.do?free_text=Install&qtxbrand=IBM+PC+Server&qtxfamily=xSeries+445&qtxnav=es&qtxdoctype=OS%20installation
- Installing SCO UnixWare
  http://www.ibm.com/pc/support/site.wss/MIGR-52460.html
- Installing NetWare 6:
  http://www.ibm.com/pc/support/site.wss/MIGR-52367.html
- Installing Microsoft Windows Server 2003:
  http://www.ibm.com/pc/support/site.wss/MIGR-51935.html
- Installing SuSE Linux Enterprise Server 8:
  http://www.ibm.com/pc/support/site.wss/MIGR-52209.html

IBM support documents

- IBM Active PCI Software
  http://www.ibm.com/pc/support/site.wss/MIGR-4J2QEQ.html
- x445 driver matrix:
  http://www.ibm.com/pc/support/site.wss/MIGR-52102.html
- NetWare LSI Logic driver:
  http://www.ibm.com/pc/support/site.wss/MIGR-44870.html
- NetWare Broadcom driver:
  http://www.ibm.com/pc/support/site.wss/MIGR-52019.html
- Windows Server 2003 installation hangs at final shutdown:
  http://www.ibm.com/pc/support/site.wss/MIGR-52480.html
- RSA View Blue Screen not supported Windows Server 2003:
  http://www.ibm.com/pc/support/site.wss/MIGR-51942.html
- Windows Server 2003 only supports NUMA configurations with PAE enabled:
  http://www.ibm.com/pc/support/site.wss/MIGR-52036.html
- Cannot remote install Microsoft Windows Server 2003:
  http://www.ibm.com/pc/support/site.wss/MIGR-52038.html
- Blue screen with ServeRAID 6.xx, Director 4.1 and Windows Server 2003:
  http://www.ibm.com/pc/support/site.wss/MIGR-52387.html

- Hot-adding memory under Windows 2003 results in slow performance:
  http://www.ibm.com/pc/support/site.wss/MIGR-52034.html
- System Management tests unavailable under diagnostics version 1.00:
  http://www.ibm.com/pc/support/site.wss/MIGR-52030.html
- NetWare restarts when NCMCON is loaded:
  http://www.ibm.com/pc/support/site.wss/MIGR-52043.html
- Broadcom 5704 Gigabit Ethernet device must be updated to V6.67 or later:
  http://www.ibm.com/pc/support/site.wss/MIGR-52039.html
- Windows Server 2003 numerates ServeRAID disks in reversed order:
  http://www.ibm.com/pc/support/site.wss/MIGR-52389.html
- Tape operations result in Windows Server 2003 Blue Screen:
  http://www.ibm.com/pc/support/site.wss/MIGR-52368.html
- Processor VRM LEDS always on when processor VRM not present:
  http://www.ibm.com/pc/support/site.wss/MIGR-52037.html
- Fans increase speed after hot swapping the power supply:
  http://www.ibm.com/pc/support/site.wss/MIGR-52031.html
- Remote video corrupted via RSA with 8-bit color modes:
  http://www.ibm.com/pc/support/site.wss/MIGR-52040.html
- IBM FAStT Storage Manager:
  http://www.ibm.com/pc/support/site.wss/MIGR-50176.html
- Token ring adapter not detected or not configured:
  http://www.ibm.com/pc/support/site.wss/MIGR-42139.html
- Behavioral differences between LSI and ServerRAID RAID-1:
  http://www.ibm.com/pc/support/site.wss/MIGR-46174.html
- NetWare 5.1: Reboots itself after the server shuts down:
  http://www.ibm.com/pc/support/site.wss/MIGR-43679.html
- Netware 6.0 - Utilization percentage is fluctuating in extreme values:
  http://www.ibm.com/pc/support/site.wss/MIGR-43532.html
- HAL for Microsoft Windows NT 4.0:
  http://www.ibm.com/pc/support/site.wss/MIGR-42067.html
- HAL for Microsoft Windows 2000:
  http://www.ibm.com/pc/support/site.wss/MIGR-42325.html
- IBM warranty upgrades:
  http://www.ibm.com/pc/support/site.wss/MIGR-4CMSC7.html

Microsoft Web pages

- Hardware Compatibility List:
  http://www.microsoft.com/whdc/hcl/default.mspx

- Windows and Hyper-Threading Technology:

  http://www.microsoft.com/windows2000/server/evaluation/performance/reports/hyperthread.asp
- VMWare Support Policy and Support Boundaries:
  http://support.microsoft.com/default.aspx?scid=kb;en-us;Q273508
- Static Resource Affinity Table:
  http://www.microsoft.com/whdc/hwdev/platform/proc/sratdwn.mspx
- Windows Datacenter High Availability FAQ:
  http://www.microsoft.com/windowsserver2003/datacenter/dcprogramfaq.mspx
- Windows Server 2003 Technical Overview:
  http://www.microsoft.com/windowsserver2003/techinfo/overview
- Windows Server 2003 TCP/IP Implementation Details:

  http://www.microsoft.com/technet/prodtechnol/windowsserver2003/plan/TCPIP03.asp
- Windows Server Catalog:
  http://www.microsoft.com/windows/catalog/server

Red Hat Web pages

- Red Hat Network: http://rhn.redhat.com
- Red Hat Enterprise Linux AS product information:
  http://www.redhat.com/software/rhel/as

Novell Web pages

- NetWare Consolidated Support Packs:
  http://support.novell.com/tools/csp

Intel Web pages

- Xeon technology overview:
  http://www.intel.com/eBusiness/pdf/prod/server/xeon/ds020903.pdf
- Hyper-Threading overview: http://www.intel.com/technology/hyperthread
- Xeon MP product briefing:
  http://www.intel.com/design/Xeon/xeonmp/prodbref
- Xeon DP product briefing: http://www.intel.com/design/xeon/prodbref
- Hyper-Threading for Servers:
  http://www.intel.com/ebusiness/hyperthreading/server

SuSE Web pages

- SuSE SLES product information:
  http://www.suse.com/us/business/products/server/sles/x86.html

VMware pages

- VMware ESX Server 2.0 documentation
  http://www.vmware.com/support/esx2/doc

How to get 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

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

B

C

D

E

F

N

O

P

R

S

T

U

V

W

X

Z

IBM @server xSeries 445 Planning and Installation Guide

(0.5” spine)
0.475”<->0.875”
250 <-> 459 pages

IBM @server xSeries 445 Planning and Installation Guide

Learn the technical details of the x445 models

Discover key IBM Director management tools

Prepare for and perform an installation

The IBM @server xSeries 445 is the IBM flagship industry standard server and is the second generation implementation of the 32-bit IBM XA-32 chipset as part of the Enterprise X-Architecture strategy. The x445 provides new levels of high availability and price performance, and offers scalability from two-way to 32-way SMP, from 2 GB to 256 GB of memory, and up to 36 PCI slots, all in one single system image.

This IBM Redbook is a comprehensive resource on the technical aspects of the server, and is divided into six key subject areas:

Chapter 1, “Technical description” introduces the server and its subsystems and describes the key features.

Chapter 2, “Positioning” examines the types of applications that are used on the x445 server.

Chapter 3, “Planning” describes the aspects of planning to purchase and install the x445.

Chapter 4, “Windows Server 2003, Datacenter Edition” describes the IBM Datacenter program.

Chapter 5, “Installation” goes through the process of installing supported operating systems.

Chapter 6, “Management” describes how to use the key IBM Director extensions designed for the x445 including the new Scalable Systems Manager.

SG24-8870-00 ISBN 073849920X

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