



IBM *e*server xSeries 455 Planning and Installation Guide

Describes the technical details of the
new 64-bit server

Covers supported Windows and
Linux 64-bit operating systems

Helps you prepare for and
perform an installation



David Watts
Aubrey Applewhite
Yonni Meza



International Technical Support Organization

**IBM @server xSeries 455 Planning and Installation
Guide**

February 2004

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

First Edition (February 2004)

This edition applies to the IBM @server xSeries 455, machine type 8855.

© Copyright International Business Machines Corporation 2004. All rights reserved.

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

Contents

Notices	vii
Trademarks	viii
Preface	ix
The team that wrote this redbook	ix
Become a published author	xi
Comments welcome	xii
Chapter 1. Technical description	1
1.1 Features	2
1.1.1 Comparing the x455 with the x450	2
1.1.2 Features not supported	3
1.2 The x455 base models	4
1.2.1 Front and rear views	4
1.3 System assembly	6
1.4 IBM XA-64 chipset	7
1.4.1 The processor-board assembly	9
1.4.2 The memory-board assembly	13
1.4.3 PCI-X board assembly	17
1.5 Remote Supervisor Adapter	19
1.6 RXE-100 Expansion Enclosure	21
1.7 Multinode scalable partitions	22
1.7.1 RXE-100 connectivity	24
1.7.2 Multinode configuration	25
1.7.3 Integrated I/O function support	26
1.7.4 Error recovery	26
1.8 Redundancy	27
1.9 Light path diagnostics	27
1.10 Extensible Firmware Interface	29
1.10.1 GUID Partition Table disk	31
1.10.2 EFI System Partition	33
1.10.3 EFI and the reduced-legacy concept	34
1.11 Operating system support	35
1.12 Enterprise X-Architecture	35
1.12.1 NUMA architecture	36
Chapter 2. Positioning	39
2.1 Migrating to a 64-bit platform	40
2.2 Scalable system partitioning	41

2.2.1	RXE-100 Expansion Enclosure	42
2.3	Operating system support	42
2.4	Server consolidation	43
2.5	ServerProven®	44
2.6	IBM Datacenter Solution Program	45
2.7	Application solutions	46
2.7.1	Database applications	46
2.7.2	Business logic	48
2.7.3	e-Business and security transactions	50
2.7.4	In-house developed compute-intensive applications	51
2.7.5	Science and technology industries	51
2.8	Why choose x455	52
Chapter 3. Planning		55
3.1	System hardware	56
3.1.1	Processors	56
3.1.2	Memory	57
3.1.3	PCI-X slot configuration	61
3.1.4	Broadcom Gigabit Ethernet controller	64
3.2	Cabling and connectivity	65
3.2.1	SMP Expansion connectivity	67
3.2.2	Remote Expansion Enclosure connectivity	70
3.2.3	Remote Supervisor Adapter connectivity	76
3.2.4	Serial connectivity	78
3.3	Storage considerations	79
3.3.1	xSeries storage solutions	79
3.3.2	Tape backup	84
3.4	Rack installation	85
3.5	Power considerations	87
3.6	Operating system support	87
3.6.1	Clustering	88
3.7	IBM Director support	89
3.8	Solution Assurance Review	90
3.8.1	Trigger Tool	90
3.8.2	Electronic Solution Assurance Review (eSAR)	90
Chapter 4. Installation		91
4.1	Using The Extensible Firmware Interface	92
4.1.1	EFI Firmware Boot Manager	93
4.1.2	The EFI shell	95
4.1.3	Driver Setup	103
4.1.4	Flash update	112
4.1.5	Configuration/Setup utility	118

4.1.6 Diagnostic utility	121
4.1.7 Boot Option Maintenance	122
4.2 Configuring scalable partitions	131
4.2.1 Creating a scalable partition	131
4.2.2 Booting a scalable partition	134
4.2.3 Multiple Monitors	135
4.2.4 Deleting a scalable partition	136
4.3 Installing Windows Server 2003	136
4.3.1 Important information	136
4.3.2 Preparing to install	138
4.3.3 Installation	139
4.3.4 Post-setup phase	144
4.4 Installing Linux.	151
4.4.1 Linux IA-64 kernel overview	152
4.4.2 Choosing a Linux distribution	153
4.4.3 Installing SUSE LINUX Enterprise Server 8.0.	156
4.4.4 Installing Red Hat Enterprise Linux AS.	160
4.4.5 Linux boot process	162
4.4.6 Information about the installed system	163
4.4.7 Using the serial port for the Linux console	171
4.4.8 RXE-100 Expansion Enclosure.	172
4.4.9 Upgrading drivers	173
Chapter 5. Management	175
5.1 IBM Director	176
5.1.1 Scalable Systems Manager	177
5.2 The Remote Supervisor Adapter.	177
5.2.1 Connecting via a Web browser.	179
5.2.2 Connecting via the ASM interconnect.	184
5.2.3 Installing the device driver.	185
5.2.4 Configuring the remote control password	186
5.3 Management using the Remote Supervisor Adapter	187
5.3.1 Configuring which alerts to monitor.	188
5.3.2 Configuring SNMP	190
5.3.3 Sending alerts directly to IBM Director	191
5.3.4 Creating a test event action plan in IBM Director	193
5.4 Windows System Resource Manager	196
5.4.1 WSRM description	197
5.4.2 WSRM features.	198
5.4.3 WSRM in the x455	199
Abbreviations and acronyms	201
Related publications	203

IBM Redbooks	203
Other publications	203
Online resources	204
How to get IBM Redbooks	208
Help from IBM	208
Index	209

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

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

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

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

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

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


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

COPYRIGHT LICENSE:

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

Trademarks

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

Chipkill™	IBM®	ServerGuide™
DB2 Connect™	ibm.com®	ServerProven®
DB2 Universal Database™	iSeries™	ServeRAID™
DB2®	LANClient Control Manager™	ThinkPad®
DRDA®	Notes®	Tivoli®
Enterprise Storage Server®	OnForever™	TotalStorage®
ESCON®	Predictive Failure Analysis®	Wake on LAN®
Eserver™	PS/2®	X-Architecture™
@server™	pSeries®	xSeries®
eServer™	Redbooks™	zSeries®
FlashCopy®	Redbooks (logo)  ™	
FICON™	RETAIN®	

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

Rational®

The following terms are trademarks of other companies:

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

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

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

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

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

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

Preface

The IBM eServer xSeries® 455 is the second generation Enterprise X-Architecture™ server using the 64-bit IBM® XA-64 chipset and the Intel® Itanium 2 processor. Unlike the x450, its predecessor, the x455 supports the merging of four server chassis to form a single 16-way image, providing even greater expandability and investment protection.

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

- ▶ Chapter 1, “Technical description” on page 1, introduces the server and its subsystems and describes the key features and how they work. This includes the Extensible Firmware Interface, which provides a powerful replacement to the BIOS facility found on the IA-32 platform.
- ▶ Chapter 2, “Positioning” on page 39, examines the types of applications that would be used on a server such as the x455.
- ▶ Chapter 3, “Planning” on page 55, describes the considerations when planning to purchase and planning to install the x455. It covers such topics as configuration, operating system specifics, scalability, and physical site planning.
- ▶ Chapter 4, “Installation” on page 91, covers the process of installing Windows® Server 2003, SUSE LINUX Enterprise Server, and Red Hat Enterprise Linux AS on the x455.
- ▶ Chapter 5, “Management” on page 175, describes how to use the Remote Supervisor Adapter to send alerts to an IBM Director management environment.

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 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 IBM Redbooks and redpapers; his most recent books include *Implementing Systems Management Solutions Using IBM Director*, SG24-6188. 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.

Aubrey Applewhaite is a Senior IT Specialist working for the IBM Systems Group in the United Kingdom. He is a member of the Server Implementation Team and specializes in xSeries hardware, Microsoft® Windows, clustering and VMware. He has worked in the IT industry for over 16 years and been at IBM for eight years. He currently works in a customer-facing role providing consultancy and practical assistance to help IBM customers implement new technology with particular emphasis on xSeries hardware. He holds a Bachelor of Science Degree in Sociology and Politics from Aston University and is an MCSE for both Windows NT® and Windows 2000, IBM eServer™ Certified Systems Expert, Cisco CNA, and Compaq Proliant ASE.

Yonni Meza is an xSeries Specialist in Peru who also works as the country's PCI Instructor, teaching courses on xSeries and Personal Computer Division (PCD) products. Additionally, he supports the PCD team with pre- and post-sales technical support, conducting demos and presentations on a regular basis. He has four years of experience in personal computing systems as well as Intel servers. Furthermore, he has implemented several xSeries solutions such as clustering in Windows and Linux with SCSI, SAN and ESS. Yonni studied Systems Engineering at the University of Lima.



The redbook team (l-r): David, Aubrey, Yonni

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

Henry Artner, Service Education Curriculum Manager, Raleigh
Pat Byers, Program Director, Linux xSeries Alliances & Marketing
Alex Candelaria, IBM Center for Microsoft Technologies, Seattle
Greg Clarke, IBM Advanced Technical Support, Dallas
Rufus Credle, International Technical Support Organization, Raleigh
Gary Hade, IBM Linux Technology Center, Beaverton
Jim Hanna, xSeries development, Austin
Cecil Lockett, Senior Engineer, Engineering Software, Raleigh
Gerry McGettigan, Advanced Technical Support, EMEA
Michael L Nelson, IBM @server Solutions Engineering, Raleigh
Lubos Nikolini, Systems Engineer, HT Computers
Charles Perkins, Course Developer, Service and Support Education, Raleigh
Steve Powell, Service and Support Education Team, Raleigh
Ken Rauch, Delivery Project Manager, Markham
Jose Rodriguez Ruibal, Advanced Technical Support, EMEA
Steve Russell, EMEA ATS xSeries Product Introduction Center, Hursley
Bob Zuber, x455 World Wide Product Manager, Raleigh
Julie Czubik, Technical Editor, ITSO, Poughkeepsie

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



Technical description

The IBM @server xSeries 455 is the latest IBM top-of-the-line server and is the second implementation of the 64-bit IBM XA-64 chipset, code named “Summit”, which forms part of the Enterprise X-Architecture strategy. The x455 completes the xSeries product family, leveraging the proven Enterprise X-Architecture to deliver robust and reliable 64-bit systems.

- ▶ “Features” on page 2
- ▶ “The x455 base models” on page 4
- ▶ “System assembly” on page 6
- ▶ “IBM XA-64 chipset” on page 7
- ▶ “Remote Supervisor Adapter” on page 19
- ▶ “RXE-100 Expansion Enclosure” on page 21
- ▶ “Multinode scalable partitions” on page 22
- ▶ “Redundancy” on page 27
- ▶ “Light path diagnostics” on page 27
- ▶ “Extensible Firmware Interface” on page 29
- ▶ “Operating system support” on page 35
- ▶ “Enterprise X-Architecture” on page 35

1.1 Features

The following are the key features of the x455:

- ▶ One-way or two-way Intel Itanium 2 models, upgradable to 4-way in a single node and 16-way in a 4-node partition.
- ▶ 64 MB Xcel4 Server Accelerator Cache providing an extra level of cache, upgradable to 256 MB in a 4-node partition.
- ▶ 1 GB or 2 GB RAM standard, upgradable to 56 GB in a single node and 224 GB in a 4-node partition. Available options are 512 MB, 1 GB, and 2 GB ECC DDR SDRAM RDIMMs.
- ▶ Memory enhancement such as memory mirroring, Chipkill™, Memory ProteXion, and hot swap.
- ▶ Dual channel Ultra320 SCSI/RAID controller.
- ▶ Six 64-bit Active PCI-X slots: Two 133 MHz, two 100 MHz, and two 66 MHz, upgradable to 24 x 64-bit PCI-X slots in a four-node partition.
- ▶ Scalable system partitioning in two-node and four-node configurations via three scalability ports.
- ▶ Connectivity to an RXE-100 external enclosure for an additional 12 PCI-X slots, upgradable to 24 additional PCI-X slots in a 4-node partition.
- ▶ Two hot-swap 1-inch drive bays, upgradable to eight in a four-node partition.
- ▶ Support for major storage subsystems, including SCSI and Fibre Channel.
- ▶ Light path diagnostics for troubleshooting.
- ▶ Remote Supervisor Adapter (RSA) for systems management and remote diagnostics.
- ▶ Integrated dual 10/100/1000 Mbps Ethernet controller.
- ▶ Integrated ATI Rage XL with 8 MB video RAM.
- ▶ Three USB ports and one serial port.
- ▶ Two 1050 W hot swap power supply.
- ▶ 24x combination DVD/CD-RW drive.
- ▶ 4U x 26" rack drawer design.

1.1.1 Comparing the x455 with the x450

The x455 builds on the proven and popular x450 and brings a number of enhancements. Table 1-1 on page 3 summarizes the differences and enhancements between the two servers.

Table 1-1 Comparing the differences between the x455 and x450

Component	x450	x455
Maximum memory (GB)	40	56
Active memory with hot-swap support	No	Yes
Multi chassis support	No	Yes (one, two or four chassis)
Shared RXE-100 between 2 machines	No	Yes
Redundant cabling to RXE-100 (only from a single independent machine)	No	Yes
Enterprise X-Architecture	First generation chipset	Second generation chipset

1.1.2 Features not supported

Due to its 64-bit architecture many existing 32-bit applications are no longer supported. These include:

- ▶ 32-bit and 16-bit operating systems
- ▶ ServerGuide™
- ▶ Remote Deployment Manager (RDM)
- ▶ LANClient Control Manager™ (LCCM)
- ▶ UpdateExpress
- ▶ Access Support

64-bit versions of some of these tool will be made available in the future.

The following functions are also not supported:

- ▶ More than one RXE-100 connected to a single node
- ▶ Hot add/remove of an RXE-100
- ▶ Physical partitioning within a single node
- ▶ Partial mirroring of memory
- ▶ Hot add/swap of the CD-ROM or diskette drive (if installed)
- ▶ Inter-Process Communications (IPC) over scalability ports
- ▶ Hot adding memory (hot swap is supported)
- ▶ PS/2® keyboard and mouse
- ▶ Parallel port

Important: The x455 does not have PS/2 ports for a keyboard and mouse. Either a USB keyboard and mouse are required or the appropriate cables to connect to a KVM switch.

1.2 The x455 base models

Powered by XA-64 Enterprise X-Architecture and the 64-bit Itanium 2 “Madison” processors, the x455 server brings the future of 64-bit processing and production-level reliability to your data centers today. Featuring mainframe-inspired advanced mission-critical functions, you can depend on these 16-way-capable enterprise servers to run your complex business applications around the clock.

The initial models of the x455 are listed in Table 1-2.

Table 1-2 Initial x455 base models

Base model	8855-1RX	8855-2RX	8855-3RX
Itanium 2 processors	1 x 1.3 GHz	2 x 1.4 GHz	2 x 1.5 GHz
Max SMP	4-way	4-way	4-way
Memory	1 GB	2 GB	2 GB
L1 cache	32 KB	32 KB	32 KB
L2 cache	256 KB	256 KB	256 KB
L3 cache	3 MB	4 MB	6 MB
Xcel4 Accelerator Cache	64 MB	64 MB	64 MB

The base models can also be connected together to form two-node (eight CPUs) and four-node (16 CPUs) configurations. See “Multinode scalable partitions” on page 22 for details.

1.2.1 Front and rear views

Figure 1-1 on page 5 shows the front view of the x455 showing the system components.




Figure 1-1 Front panel of the xSeries 455

Figure 1-2 shows the rear view of the x455 showing the system connectors.




Figure 1-2 Rear view of the x455

1.3 System assembly

The x455 has a similar internal design to the x450. The midplane board (viewed from the front) interfaces with three major assemblies:

- ▶ The processor-board assembly

This is located to the right of the midplane board and under the memory-board assembly. It houses the Itanium 2 processors, Cache and Scalability Controller, and the 64 MB of Xcel4 cache.

- ▶ The memory-board assembly

This is located to the right of the midplane board and above the processor-board assembly. It houses the memory and memory controller.

- ▶ The PCI-X board assembly

This is located to the left of the midplane board. It houses all the PCI-X slots and all other I/O components.




Figure 1-3 Memory-board and processor-board assembly locations

1.4 IBM XA-64 chipset

The IBM XA-64 chipset is the product name describing the chipset developed under the code name “Summit” and implemented on the IA-64 platform. A product of the IBM Microelectronics Division, the XA-64 chipset leverages the proven Enterprise X-Architecture chipset used initially in the x440 and applies the same technologies to the IA-64 architecture. The XA-64 chipset comprises the following components:

- ▶ Itanium 2 processor(s)
- ▶ Cache and Scalability Controller

A single controller, code named “Tornado”, located within the processor-board assembly.

- ▶ Memory controller

A single memory controller, code named “Cyclone”, located within the memory-board assembly.

- ▶ Two PCI-X bridges

Two PCI-X bridges, code named “Winnipeg”, one located on the PCI-X board and the other on the I/O board. These control both the PCI-X and Remote I/O.

Figure 1-4 shows the various IBM XA-64 components in an x455 configuration.

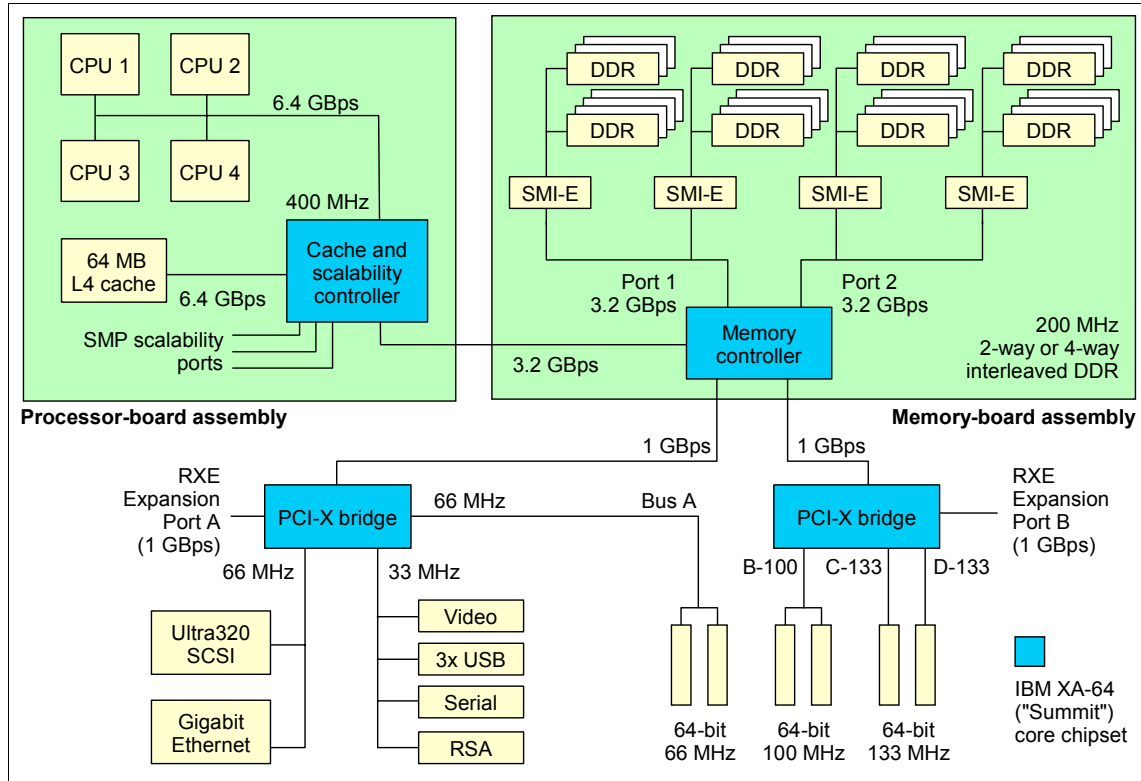


Figure 1-4 xSeries 455 system block diagram

Table 1-3 shows how the bandwidths in Figure 1-4 are calculated.

Table 1-3 Bus speeds

From	To	Bandwidth	Calculation
CPUs	Cache controller	6.4 GBps	400 MHz x 128-bit data path
L4 Cache	Cache controller	6.4 GBps	400 MHz x 128-bit data path

From	To	Bandwidth	Calculation
SDRAM	Memory controller	3.2 GBps	400 MHz x 64-bit data path
Cache controller	Memory controller	3.2 GBps	400 MHz x 64-bit data path

1.4.1 The processor-board assembly

The processor-board assembly is located below the memory board. It is held in place by retaining levers, an EMC shield and a retention bracket. For instructions to remove or install please refer to the *Option Installation Guide*.

Warning: Be careful when removing or installing the processor-board assembly or the memory-board assembly. It is possible to damage the midplane if not done correctly.




Figure 1-5 The processor-board assembly

The power modules shown in Figure 1-5 supply power to the processors and are equivalent to VRMs in other systems.

Processors should be installed in the order 1, 2, 3, 4. The bootstrap processor (BSP) may not necessarily be the processor located in processor socket 1. The Intel Itanium Architecture processors are initialized and tested in parallel. The first processor to complete initialization becomes the BSP.

The CPUs are connected together with a 200 MHz frontside bus, but supply data at an effective rate of 400 MHz using the “dual-pump” design of the Intel Itanium 2 architecture is described in “Intel Itanium 2 processors” on page 10.

The processor-board assembly is also equipped with LEDs for light path diagnostics for the following components:

- ▶ Each processor
- ▶ Each power module (“pod”)

In addition, a “remind” button is located on the upper side of the processor-board assembly. Pressing this button while the processor-board assembly is not attached to AC power will illuminate any light path LEDs for 10 seconds that had been lit while the system was under power.

Intel Itanium 2 processors

The Itanium 2 processor used in the x455 is code named “Madison”. It uses a ZIF socket design and its small form factor is what allows up to four processors in a 4U dense machine.

Table 1-4 outlines some of the differences between the Itanium and Itanium 2 processors (both the “Madison” and the earlier “McKinley” processor).

Table 1-4 Itanium v Itanium 2 processors

Feature	Itanium	Itanium 2 “McKinley”	Itanium 2 “Madison”
Processor core speed	733 or 800 MHz	900 MHz or 1.0 GHz	1.3, 1.4 or 1.5 GHz
L3 Cache	2 or 4 MB	1.5 or 3 MB	3, 4 or 6 MB
Frontside bus	266 MHz	400 MHz, 128 bit	400 MHz, 128 bit
Frontside bus bandwidth	2.1 GBps	6.4 GBps	6.4 GBps
Pipeline stages	10	8	8
Issue ports	9	11	11
on-board registers	328	328	328
Integer units	3	6	6
Branch units	3	3	3
Floating point units	2	2	2
SIMD units	2	1	1
Load and store units	2 (total)	2 load and 2 store	2 load and 2 store

The Itanium 2 processor has three levels of cache, all of which are on the processor die:

- ▶ Level 1 cache, 32 KB

This is new and the “closest” to the processor, and is used to store micro-operations. These are decoded executable machine instructions. It serves these to the processor at rated speed. This additional level of cache saves decode time on cache hits.

- ▶ Level 2 cache, 256 KB

This is equivalent to L1 cache on the Pentium® III Xeon.

- ▶ Level 3 cache 3–6 MB

This is equivalent to L2 cache on the Pentium III Xeon or the L3 cache on the Pentium Xeon MP processor. Unlike the design of the original Itanium processor, this L3 cache is now on the processor die, greatly improving performance, up to 2 times greater than that of the original Itanium.

The x455 also implements a Level 4 cache as described in “IBM Xcel4 Accelerator Cache” on page 12.

Intel has also introduced a number of features associated with its Itanium micro-architecture. These are available in the x455, including:

- ▶ 400 MHz frontside bus

The Pentium III Xeon processor had a 100 MHz frontside bus that equated to a burst throughput of 800 MBps. With protocols such as TCP/IP, this had been shown to be a bottleneck in high-throughput situations. The Itanium 2 processor improves on this by using a single 200 MHz clock but using both edges of each clock cycle to transmit data. This is shown in Figure 1-6.




Figure 1-6 Dual-pumped frontside bus

This increases the performance of the frontside bus. The end result is an effective burst throughput of 6.4 GBps (128-bit wide data path running at 400 MHz), which can have a substantial impact, especially on TCP/IP-based LAN traffic. This is opposed to the Itanium processor, which had a burst throughput of only 2.1 GBps (64-bit wide data path running at 266 MHz).

- ▶ Explicitly Parallel Instruction Computing (EPIC)

EPIC technology, developed by Intel and HP, leads to more efficient, faster processors because it eliminates numerous processing inefficiencies in current processors and attacks the perennial data bottleneck problems by increasing parallelism, rather than simply boosting the raw “clock” speed of the processor.

Specifically, in today's 32-bit processors, much of the instruction scheduling (the order in which computing instructions are executed) is done on the chip itself, leading to a great deal of overhead and slowing down overall processor performance. Moreover, today's processors are plagued by instruction flow problems since the processor often has to stop what it is doing and reconstruct the instruction flow due to inherent inefficiencies in instruction handling.

EPIC makes the instruction scheduling more intelligent and handles much of the scheduling off-chip, in the compiler program, before feeding "parallelized" instructions to the Itanium 2 processor for execution. The parallelized instructions allow the chip to process a number of instructions simultaneously, increasing performance.

The Itanium 2 architecture is based on EPIC technology and has the following features:

- Provides faster online transaction processing
- Has the capability to execute multiple instructions simultaneously
- Enables faster calculations and data analysis
- Allows for faster storage and movement of large models (CAD, CAE)
- Speeds up simulation and rendering times

For more information about the features of the Itanium 2 processor, go to:

<http://www.intel.com/design/itanium2>

IBM Xcel4 Accelerator Cache

Integrated into the processor-board assembly is 64 MB of Level 4 cache, which is shown in Figure 1-4 on page 8. This Xcel4 Server Accelerator Cache provides the necessary extra level of cache to maximize CPU throughput by reducing the need for main memory access under demanding workloads, resulting in an overall enhancement to system performance.

Cache memory is two-way interleaved 200 MHz DDR memory and is faster than the main memory because it is directly connected to the Cache and Scalability Controller and does not have additional latency associated with the large fan-out necessary to support the 28 DIMM slots. Since the data interface to the controller is 400 MHz, peak bandwidth for the Xcel4 cache is 6.4 GBps.

The Xcel4 Accelerator Cache has been designed with commercial workloads in mind that tend to have high cache hit rates. This effectively boosts performance and compensates in part for the 3.2 GBps bandwidth between the Cache and Scalability Controller and Memory Controller.

1.4.2 The memory-board assembly

The x455 memory-board assembly is installed in the top of the server and mounts to the side of the midplane using two retaining levers on the top. This location allows for easy access to all memory DIMMs without having to remove any components from the system.

The memory-board assembly houses 28 DIMM slots. All DIMM slots can be used with 512 MB, 1 GB or 2GB RDIMMs for a maximum of 56 GB. Memory can be hot swapped but not hot added. This function, however, is as much a part of the operating system as the hardware. Check with your operating system Help for details.




Figure 1-7 The memory-board assembly, showing the two memory ports

The memory-board assembly is also equipped with LEDs for light path diagnostics for each DIMM. In addition, the assembly is equipped with LEDs for the following:

- ▶ Power to memory port 1
- ▶ Power to memory port 2
- ▶ Hot-plug memory enabled

Memory used in the x455 is standard PC2100 ECC DDR SDRAM RDIMMs. The memory is 2-way interleaved; however, 4-way interleaving is also supported when both ports are engaged. Interleaving requires DIMMs to be installed in matched pairs and in specific DIMM sockets (see “Memory” on page 57).

System memory

DIMMs must be installed in matched pairs, since the DIMMs are two-way interleaved. However, if memory is installed in matched fours (a matched pair in each port), the system automatically detects this and will enable 4-way interleaving. With this, memory access is performed simultaneously from both ports (two separate paths into the memory controller as shown in Figure 1-4 on page 8), leading to improved memory performance.

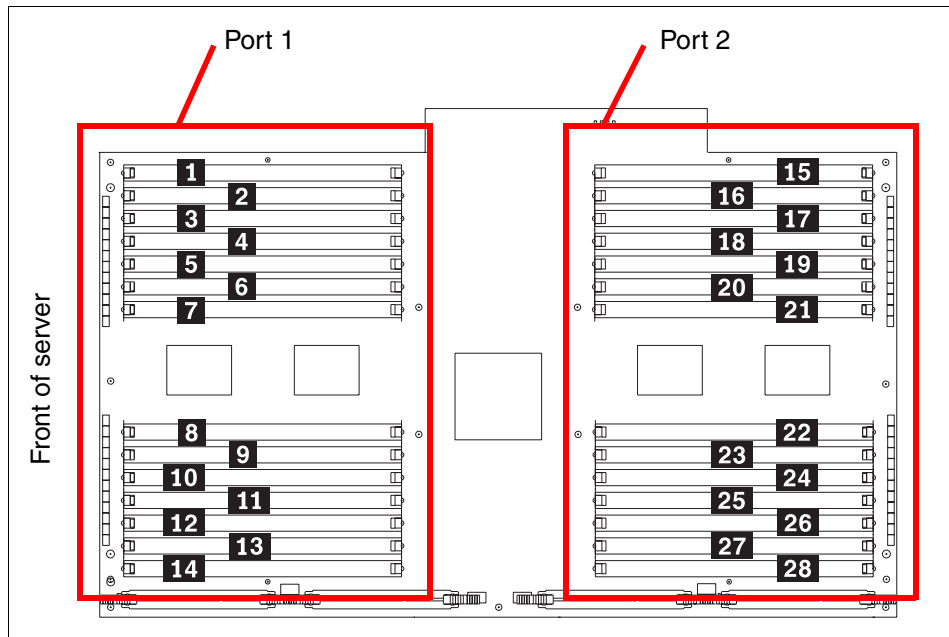


Figure 1-8 Memory DIMMs are divided into two ports

There are a number of advanced features implemented in the x455 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.

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 x455 (and several other xSeries servers) use an advanced ECC algorithm that is based not on bits but on memory symbols. Symbols are groups of multiple bits, and in the case of the x455, each symbol is 4 bits wide. With two-way interleaved memory, the algorithm needs only three symbols to perform the same ECC functions, thus leaving one symbol free (2 bits on each DIMM).

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 symbol (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 corrects soft errors and reports recoverable errors. An excessive rate of recoverable errors reported triggers Memory ProteXion to replace the failing locations.

- ▶ Memory mirroring

Memory mirroring is equivalent to RAID-1 in disk arrays, in that memory is divided in two ports and one port is mirrored to the other (see Figure 1-8 on page 14). If 8 GB is installed, then the operating system sees 4 GB once memory mirroring is enabled (it is disabled by default). All mirroring activities are handled by the hardware without any additional support required from the operating system.

When memory mirroring is enabled the data that is written to memory is stored in two locations. One copy is kept in the port 1 DIMMs, while a second copy is kept in the port 2 DIMMs.

During the execution of the read command, the data is read simultaneously from both ports, and error-free data from either port is forwarded. This provides an extra level of error recovery capability.

When an unrecoverable memory error from one of the memory ports is encountered, good data from the non-failing memory port is forwarded to the system. The failing DIMM is reported and indicated with light path. The failing memory port is then disabled.

Certain restrictions exist with respect to placement and size of memory DIMMs when memory mirroring is enabled. These are discussed in “Memory mirroring” on page 58.

► Chipkill memory

Chipkill is integrated into the XA-64 chipset and does not require special Chipkill DIMMs. Chipkill corrects multiple single-bit errors to keep a DIMM from failing. When combining Chipkill with Memory ProteXion and Active Memory, the x455 provides very high reliability in the memory subsystem. Chipkill memory is approximately 100 times more effective than ECC technology, providing correction for up to 4 bits per DIMM, whether on a single chip or multiple chips.

If a memory chip error does occur, Chipkill is designed to automatically take the inoperative memory chip offline while the server keeps running. The memory controller provides memory protection similar in concept to disk array striping with parity, writing the memory bits across multiple memory chips on the DIMM. The controller is able to reconstruct the “missing” bit from the failed chip and continue working as usual.

Chipkill support is provided in the memory controller and implemented using standard RDIMMs, 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 DIMM and continue operating with reduced memory capacity. You can manually re-enable the memory bank after the problem is corrected via the Setup/Configuration option in the EFI Firmware Boot Manager menu. EFI is *Extensible Firmware Interface*, the replacement to BIOS as described in “Extensible Firmware Interface” on page 29.

Memory ProteXion, memory mirroring, and Chipkill provide multiple levels of redundancy to the memory subsystem. Combining Memory ProteXion with Chipkill enables up to two memory chip failures per memory port (14 DIMMs). Both memory ports could sustain up to four memory chip failures. Memory mirroring provides additional protection with the ability to continue operations with memory module failures.

1. The first failure detected by the Chipkill algorithm on each port does not generate a light path diagnostics error, since Memory ProteXion recovers from the problem automatically.
2. Each memory port could then sustain a second chip failure without shutting down.
3. Provided that memory mirroring is enabled, the third chip failure on that port would send the alert and take the DIMM offline, but keep the system running out of the redundant memory bank.

The combination of these technologies provides the most reliable memory subsystem available.

1.4.3 PCI-X board assembly

Strictly speaking the PCI-X board assembly does not truly exist in the same way that the processor board and memory board assemblies do. The term is used here to loosely describe the PCI-X board and I/O subsystem that together comprise the other components of the machine.

The two PCI-X bridges in the XA-64 chipset provide support for 33, 66, 100, and 133 MHz devices using 4 PCI-X buses (labelled A–D in Figure 1-4 on page 8).

The PCI-X bridges also have two 1 GBps bi-directional Remote Expansion I/O (RXE) ports for connectivity to the RXE-100 Expansion Enclosure. The RXE-100 provides up to an additional 12 PCI-X slots and can be connected by a single cable to port A or to both ports A and B to provide redundancy.

The rear panel of the x455, which indicates the location of the RXE Expansion Ports, is shown in Figure 1-2 on page 5.

PCI-X subsystem

There are six available PCI-X slots in four buses. These are shown in Figure 1-9.




Figure 1-9 PCI-X slot information

Only 3.3V cards are supported, but all slots will accept the following:

- ▶ 32-bit or 64-bit cards
- ▶ 33MHz, 66MHz, 100MHz or 133MHz cards

You should bear in mind that the overall throughput will be the lowest common denominator. As such it is important that you carefully consider the placement of cards to get the best throughput.

The following guidelines will help you gain the best performance from the PCI-X buses:

- ▶ Put cards in a slot that will give the best performance but not in a slot rated higher than the card.
- ▶ If a 133 MHz card is in bus B then avoid using the other slot in the bus if possible so the card will operate at its rated speed.
- ▶ Use buses B, C and D in preference to A. Bus A is connected to the PCI-X south bridge with legacy devices and competes with these “slower” devices when sending and receiving data to and from the memory controller.

Important: The rating of bus B at 100 MHz is for a fully populated bus. If a single 133 MHz card is installed the bus increases to 133 MHz. This is part of the PCI-X specification.

See “PCI-X slot configuration” on page 61 for details on what adapters are supported and in what combinations.

The PCI-X subsystem also supplies these I/O devices:

- ▶ Dual channel Ultra320 SCSI/RAID controller with one internal port and one external port using the LSI 53C1030 chipset

The internal port supports both single disks and RAID-1 mirrored pairs of disks. The external port will support additional RAID configurations but this has a number of limitations and is not recommended.

Note: The LSI 53C1030 is a PCI-X to Dual Channel Ultra320 SCSI/RAID Multifunction Controller. Internally, however, each SCSI channel is managed separately and controlled by two separate PCI-X configuration spaces. As a result it may be viewed by software such as LSI configuration, ServeRAID™, and the operating system as two separate controllers, each with a single channel.

- ▶ Dual Gigabit 10/100/1000 Ethernet ports using the Broadcom 5704 chipset.
The BCM5704 supports full- and half-duplex performance at all speeds (10/100/1000 Mbps, auto-negotiated) and includes integrated on-chip memory for buffering data transmissions to ensure the highest network performance. It has dual onboard RISC processors for advanced packet parsing and backwards compatibility with today's 10/100 networks. It includes

software support for failover, layer-3 load balancing, 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 that your network infrastructure is capable of the necessary throughput to match the server's I/O capacity.

- ▶ SVGA with 8 MB video memory (ATI RageXL chipset)
- ▶ Three USB ports: One on the front panel and two on the rear
All ports are USB 2.0 compliant.
- ▶ One RS-232 serial port, located on the rear of the machine
- ▶ Remote Supervisor Adapter with:
 - 10/100 Ethernet port management port
 - RS-485 ASM interconnect bus
 - Serial port
 - External power port

There are no PS/2 keyboard or mouse ports on the x455. USB keyboard and mice are supported. If you require KVM support, the 1.5 M USB Conversion Option (UCO) (part number 73P5832) enables the x455 to be attached to one of the Advanced Connectivity Technology (ACT) switches for common management within the rack. This smart cable is plugged into the USB and video ports on the server. It converts KVM signals to CAT5 signals for transmission over a CAT5 cable to either a Remote Console Manager (RCM) or Local Console Manager (LCM). USB servers can be managed on the same set of switches as legacy PS/2-based or C2T-based KVM servers.

1.5 Remote Supervisor Adapter

The x454 includes a Remote Supervisor Adapter (RSA), which is positioned horizontally in a dedicated PCI-X slot beneath the PCI-X adapter area of the system.




Figure 1-10 Remote Supervisor Adapter connectors

The Remote Supervisor Adapter allows you to provide remote management both out-of-band and in-band. Out-of-band refers to managing the server without the use of the operating system. This would be done via the Ethernet or serial port via a Web or telnet session. In-band refers to managing the server through the operating system. This would typically be via IBM Director and/or SNMP.

The features of the Remote Supervisor Adapter include:

- ▶ In-band and out-of-band remote server access and alerting through IBM Director
- ▶ Full Web browser support with no other software required
- ▶ Enhanced security features
- ▶ Graphics/text console redirection for remote control
- ▶ Dedicated 10/100 Ethernet access port
- ▶ ASM interconnect bus for connection to other service processors
- ▶ Serial dial in/out
- ▶ E-mail, pager and SNMP alerting
- ▶ Event log
- ▶ Predictive Failure Analysis on memory, power, hard drives, and CPUs
- ▶ Temperature and voltage monitoring with settable threshold
- ▶ Light path diagnostics
- ▶ Automatic Server Restart (ASR) for operating system and POST
- ▶ Remote firmware update
- ▶ LAN access
- ▶ Alert forwarding
- ▶ Ability to manage and monitor an RXE-100

See the IBM Redbook *Implementing Systems Management Solutions using IBM Director*, SG24-6188, for more information on the Remote Supervisor Adapter.

1.6 RXE-100 Expansion Enclosure

The attachment of an RXE-100 Remote Expansion Enclosure is also supported. These are supported in a number of configurations with some configurations offering redundancy. The RXE-100 is connected to the x455 via one or two remote I/O cables and has a throughput of 1 GBps. It is available as standard with six PCI-X slots available but is upgradable to 12 PCI-X slots, giving you up to a total of 12 PCI-X or 18 PCI-X slots, respectively. In a multinode configuration an RXE-100 can be shared by two machines. This allows a total of one RXE-100 in a 2-node partition and two RXE-100s in a 4-node partition.

For systems configured for RXE failover, one of the RXE Expansion Port connections may fail at either runtime or boot time and the system will continue to operate properly. Failover capability requires the RXE-100 to be configured with 12 slots, be configured in unified (that is, non-shared) mode, and have both RXE Expansion Ports connected. Runtime failures will be logged in the System Error Log. The broken connection is not re-established at runtime. A Reboot is required to recover failover capability.

An RXE-100 connection also requires the connection of the RXE management port(s). There is no failover capability for RXE management port connections.

Important: Although other combinations of attaching RXE-100s in a multinode partition will work, they are not supported configurations.

The following connectivity options are supported:

- ▶ Connecting one 6-slot RXE to one x455
- ▶ Connecting one 12-slot RXE to one x455 using one or (recommended) two data cables
- ▶ Sharing one 12-slot RXE between two independent x455s
- ▶ Connecting one RXE to a two-node x455, configured so as to share the slots between the nodes
- ▶ Connecting one RXE to a 4-node x455 (one RXE connected to two of the nodes, either nodes 1 and 2, or nodes 3 and 4)
- ▶ Connecting two RXEs to a 4-node x455 (one RXE connected to two of the nodes), configured so as to share the slots between nodes 1 and 2, and between nodes 3 and 4

See “Remote Expansion Enclosure connectivity” on page 70 for details.

1.7 Multinode scalable partitions

The x455 can be configured as part of a multinode partition. The partition can comprise either two nodes or four nodes. The partition is a hardware partition, which is invisible to the operating system. From an operating system point of view it sees all the hardware as a single machine.

All nodes in a scalable partition must have the following:

- ▶ Same base model
- ▶ Processors that have the same core speed and cache
- ▶ Memory mirroring that is configured the same on all nodes, that is, either all enabled or all disabled

When you have a multi-node configuration, all nodes must be fully populated with processors. In other words, only eight-way and 16-way CPU configurations are supported.

Conceptually two-node and four-node partitions are depicted in Figure 1-11 and Figure 1-12 on page 23. These also show the cable requirements.

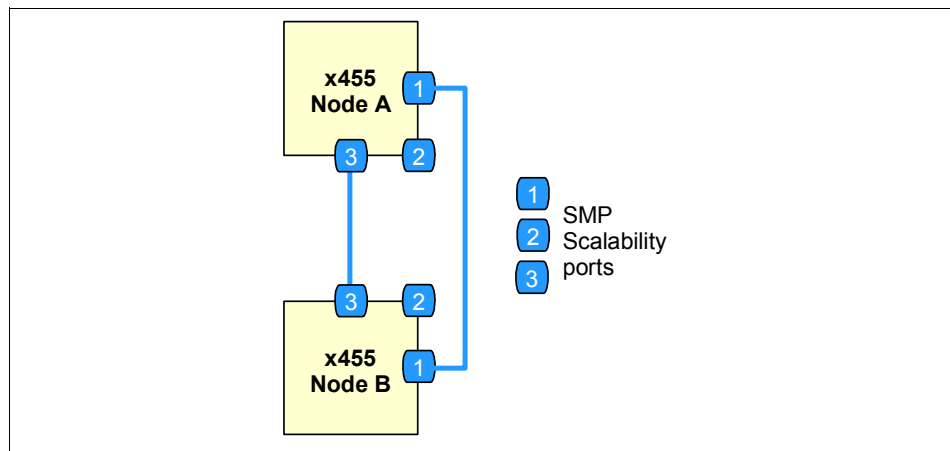


Figure 1-11 Two-node partition

Component	1 node	2 nodes	4 nodes
Memory	56 GB	112 GB	224 GB
Xcel4 Cache	64 MB	128 MB	256 MB
Internal hard disks	2	4	8
Active 64-bit PCI-X slots	2 x 66 MHz 2 x 100 MHz 2 x 133 MHz	4 x 66 MHz 4 x 100 MHz 4 x 133 MHz	8 x 66 MHz 8 x 100 MHz 8 x 133 MHz
RXE-100 support	1	1	2
Active 64-bit PCI-X slots with RXE	2 x 66 MHz 2 x 100 MHz 14 x 133 MHz	4 x 66 MHz 4 x 100 MHz 16 x 133 MHz	8 x 66 MHz 8 x 100 MHz 32 x 133 MHz
Dual channel integrated Ultra320 SCSI support	1	2	4
Dual port integrated Ethernet controller	1	2	4
Keyboard and mouse	1	2	4
USB ports	3	6	12
Video ports	1	2 ¹	4 ¹
Wake on LAN® Ethernet cards	2	2	2
1. This is supported for Windows Server 2003 Enterprise Edition only			

1.7.1 RXE-100 connectivity

RXE-100s may be attached to x455 systems. They can be connected using either optional 3.5m Scalability Port Cables or optional 8m RXE Cables. The RXE cable length allows placement of RXE-100s in a rack adjacent to the rack containing the x455 system.

Conceptually, two-node and four-node configurations with RXE-100s are depicted in Figure 1-13 on page 25 and Figure 1-14 on page 25. These also show the cable requirements.

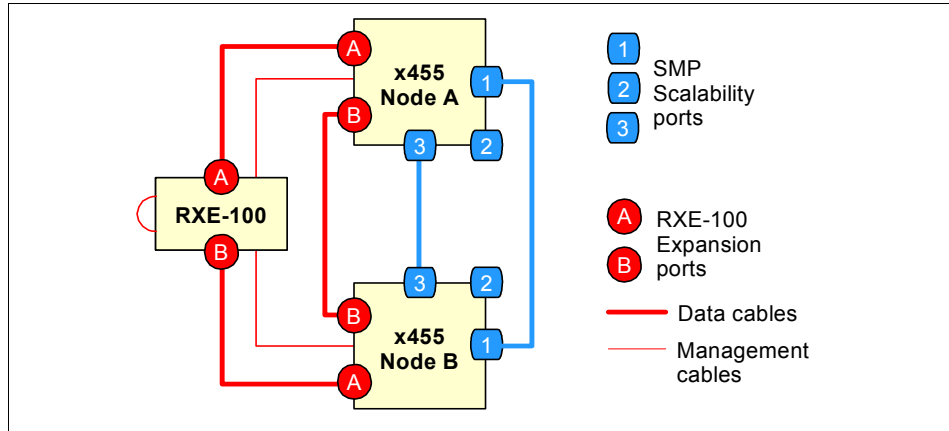


Figure 1-13 Two-node configuration with one RXE-100

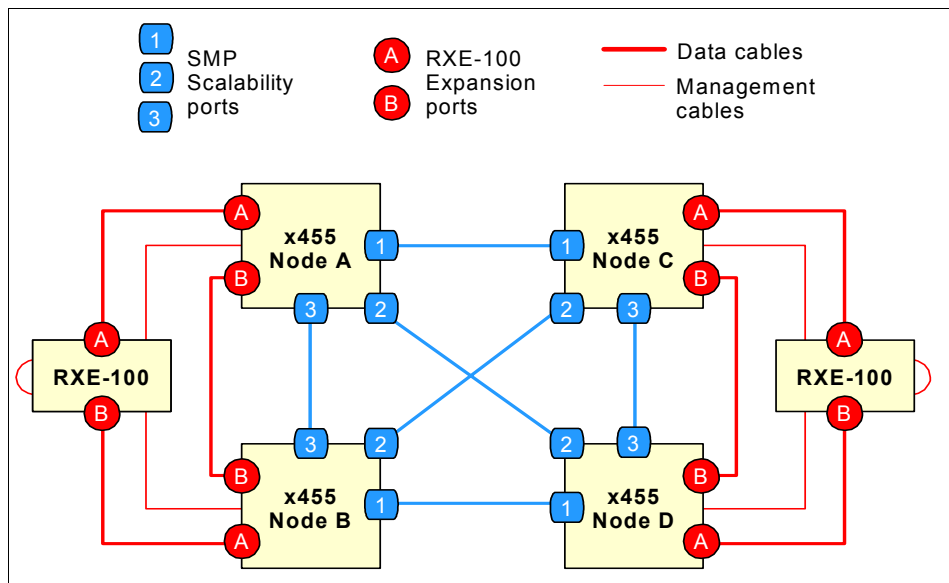


Figure 1-14 Four-node configuration with two RXE-100s

1.7.2 Multinode configuration

Multinode system configurations are defined to each x455 using the EFI Configuration/Setup menus. These options are fairly self explanatory and allow you to configure the following:

- Partition size (two-node or four-node)

- Partition ID
- IP address of member nodes (IP address of the RSA)
- Additional shared resources (video and/or CD-ROM)

All system node participants in a multinode system must be connected to the network via the Service Processor Ethernet port.

1.7.3 Integrated I/O function support

The following integrated I/O functions are available in multinode x455 systems.

Table 1-6 Integrated I/O function support

Function	Primary node	Secondary nodes
USB ports	Yes	Yes ¹
Serial port	Yes	No
Disk drive bays (SCSI)	Yes	Yes
SCSI port	Yes	Yes
Ethernet ports	Yes	Yes
Video ports	Yes	Yes ^{2/3}
Media bays (IDE)	Yes	Yes
Wake on LAN	Yes	No
1. Keyboard or mouse can be used on any machine. 2. Windows Server 2003 Enterprise Edition only. 3. Windows Server 2003 Datacenter Edition does not support multiple monitors as there currently is no certified driver for the ATI RageXL chipset.		

1.7.4 Error recovery

With a multinode configuration, in the event there is a major problem with one of the nodes, the remaining nodes will boot to EFI as single node partitions. This is done to facilitate diagnostics and reconfiguration.

1.8 Redundancy

The x455 has the following redundancy features to maintain high availability:

- ▶ Four hot-swap multi-speed fans.

With four hot-swap redundant fans, the x455 has adequate cooling for each of its major component areas. There are 2 fans located at the front of the server that direct air through the memory-board assembly and processor-board assembly. 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.

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 and will compensate for the reduction in air flow. In general, failed fans should be replaced within 24 hours following failure.

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.

For large configurations, redundancy is achieved only when connected to a 220 V power supply. See 3.5, “Power considerations” on page 87 for details.

To ensure adequate power, a UPS with a rating of RMB 5000 or more is recommended.

- ▶ Two hot-swap hard disk drive bays. Using either the onboard LSI chipset or a ServeRAID adapter, these drives can be configured to form a RAID-1 disk array for the operating system.
- ▶ The memory subsystem has a number of redundancy features, including memory mirroring, as described in “System memory” on page 14.

The layout of the front panel of the x455, showing the location of the four fans, two drive bays, and two power supplies is shown in Figure 1-1 on page 5.

1.9 Light path diagnostics

To limit the need to slide the server out of the rack to diagnose problems, a light path diagnostics panel has been incorporated in the front of the x455, as shown in Figure 1-15. This panel can be ejected from the server to view all light path

diagnostics monitored server subsystems. In the event that maintenance is 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 1-15, light path diagnostics are able to monitor and report on the health of CPUs, main memory, hard disk drives, PCI-X slots, fans, power supplies, power modules and the internal system temperature.




Figure 1-15 Light path diagnostics panel on the x455

Important: If a light path diagnostics LED has been illuminated and system power is removed, there is no way to redisplay the LEDs on the system tray without re-applying AC power. If the fault has not been rectified when power is restored, the LED will re-light.

The light path diagnostics on the x455 has four levels:

- ▶ The first level is the front panel fault LED.
- ▶ Level 2 is the pop-out panel, as shown in Figure 1-15.
- ▶ 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.
- ▶ For the fourth level of diagnostics, LEDs on major system components indicate the component causing the error.

As the processor-board assembly is not visible during normal operation, a light path diagnostics button has been incorporated into it to assist with diagnosing errors. You can light up the LEDs for a maximum of two minutes. After that time, the circuit that powers the lights is exhausted.

The pop-out panel (Figure 1-15 on page 28) also has a remind button. This places the front panel system-error LED into remind mode, which means it flashes briefly every two 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 and no longer blink. 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.

Tip: The remind button on the pop-out LPD panel does not function when AC power has been removed from the system. The button is just used to acknowledge a system error as described above.

1.10 Extensible Firmware Interface

The *Extensible Firmware Interface* (EFI) specification describes an interface between the operating system and platform firmware as shown in Figure 1-16 on page 30. The interface offers platform-related information to the operating system as well as boot and runtime service calls that are available to the operating system and OS loader. Together, this makes a well-defined environment for booting the operating system and running pre-boot applications, such as diagnostics, system setup, and driver setup utilities.

In comparison to a BIOS-based, legacy system, the EFI is an additional layer between the operating system and the firmware. In a legacy system, the OS loader calls BIOS functions directly. Consequently, to provide a stable boot environment, changes in the OS loader and the platform firmware must go hand-in-hand. Figure 1-16 on page 30 shows a conceptual view of the EFI.

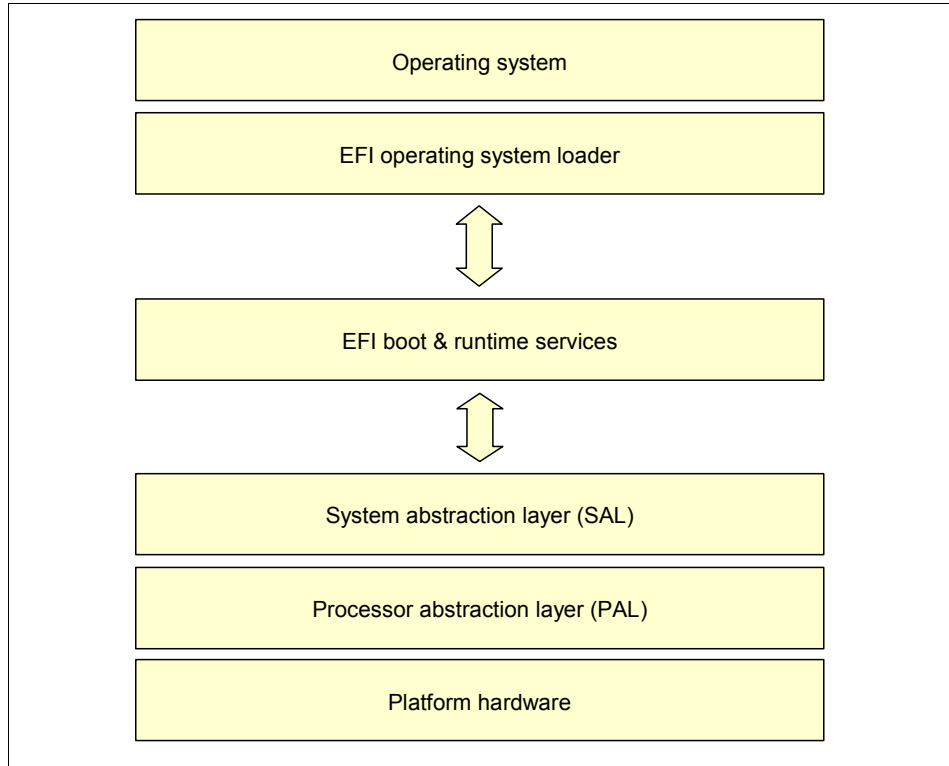


Figure 1-16 The EFI concept

The primary goal of this specification is to provide an abstract model for both operating system and hardware developers. With such a model in a place, OS loader customizations are not required if there are changes in the platform hardware or firmware such as added new boot or input devices, for instance. The EFI breaks up a tight dependency between the operating system and the firmware, thus speeding up the process of releasing the new products and introducing the new features and functionality to the hardware and/or operating system.

Consider, for example, the situation where a new type of boot device, say a USB key, is to be implemented. First the legacy BIOS would have to offer an option to choose this new device for booting, then new USB key-specific functions would have to be added to the firmware to support booting from a USB device, and finally, the OS loader would have to be modified to use these functions.

The same situation with the EFI would be dramatically simplified. The OS loader calls unified (not vendor-specific) EFI API functions for booting. These functions are not dependent on the boot device used, so when a new boot device type is

added to the platform and the firmware is modified to recognize it, the operating system can immediately boot.

Although EFI was originally introduced with the Itanium architecture, it is not restricted to 64-bit platforms. The EFI architecture is modular, extensible, and offers backward compatibility for the older systems by default. This means that it is possible for non-EFI-aware operating systems to communicate directly with system BIOS/firmware; however, this depends on how a particular manufacturer implements the EFI on its products. The x455 has a small portion of legacy BIOS code used only for running the IA-32 video driver; however, it will only support EFI-aware operating systems. From a long-term perspective, for IA-32 systems there is expected to be a gradual transition from BIOS to the EFI.

EFI does offer an extensive development interface that enables it to be customized to load additional drivers, provide corporate identity, or provide authentication and security features. Tools are available from Intel's Web site to assist with this.

If you have not used an EFI system before you will come across a number of new acronyms. These are listed in the table below with a brief description.

Table 1-7 EFI acronyms and descriptions

Acronym	Description
ACPI	Advanced Configuration and Power Interface - This is an industry standard interface for OS-directed configuration and power management for servers, desktops, and laptops.
EFI	Extensible Firmware Interface - A new interface between the system firmware and operating system.
PAL	Processor Abstraction Layer - This provides various processor functions that vary between different implementations of a particular platform.
SAL	System Abstraction Layer - This is a firmware layer and provides various system functions that vary between different implementations of a particular platform.
SMBIOS	Systems Management BIOS - A specification that defines how systems management information is presented in a standard format on Intel architecture systems.

1.10.1 GUID Partition Table disk

The *GUID Partition Table* (GPT) was introduced as part of the EFI initiative. Every disk is assigned a global unique identifier (GUID) to allow self-identification

of the disks. GPT replaces the older Master Boot Record (MBR) partitioning scheme that has been common to PCs.

There are several reasons for introducing a new partitioning scheme:

- ▶ MBR disks support only four partition table entries and a volume size of 2 TB (terabytes). If more partitions are required then one of these partitions must be an extended partition. Only one extended partition is allowed per disk drive. Extended partitions are then subdivided into one or more logical disks.
- ▶ In theory, GPT disks support an unlimited number of partitions. The number of partitions' is limited only by the amount of space reserved for making partitions entries. GPT disks can grow to a very large size. This can be up to 2^{64} logical blocks in length (logical blocks are typically 512 bytes). This equates to 18 EB (Exabytes). In practice, the maximum is less, as there are practical limitations with other parts of the system accessing this amount of disk space efficiently.
- ▶ GPT disks use primary and backup partition tables for redundancy and CRC32 fields for improved partition data structure integrity.

For backward compatibility with legacy MBR disk tools, all GPT disks contain a *protective MBR*. The protective MBR, beginning in sector 0, precedes the GUID Partition Table on the disk and contains only one partition entry that appears to span the entire disk. The legacy tools are not aware of GPT and do not know how to properly access a GPT disk. The benefit of a protective MBR is that these tools will view a GPT disk as having a single encompassing (possibly unrecognized) partition, rather than mistaking the disk for one that is unpartitioned. That is why a GPT-partitioned disk appears to have an MBR.

GPT disks can be converted to MBR disks and vice versa only if all existing partitioning is first deleted, with associated loss of data. To a legacy tool, the structure of any GPT disk appears to look like Figure 1-17. The protective MBR is followed by a theoretically unlimited number of data (or possibly unrecognized) partitions.



Figure 1-17 General GPT disk structure

Currently, only 64-bit operating systems have the ability to read, write, and boot from GPT disks. 32-bit operating systems do not have built-in support for GPT disks.

The specification for GPT disk partitioning can be found in Chapter 16 of the Extensible Firmware Interface (EFI) specification. This document is available at:

<http://developer.intel.com/technology/efi/download.htm>

1.10.2 EFI System Partition

A special partition on the GPT disk is the *EFI System Partition (ESP)*. It contains the OS loader files of all installed operating systems. These files are stored in the EFI directory. The ESP may also contain other files necessary to boot the system, such as drivers. The EFI System Partition is shareable among all installed operating systems. To support multiple operating system installations, create multiple data partitions.

Warning: Windows Server 2003 Enterprise and Datacenter Editions require a Microsoft Reserved Partition (MSR), which must come after the ESP. If Windows and Linux need to reside on the same machine then Windows must be installed first.

An example directory structure for an EFI System Partition present on a hard disk with Linux and Windows Server 2003 installed is as follows:

```
\EFI
  \Microsoft
    \WINNT50
    \EFIDrivers
  \SuSE
  \RedHat
\MSUtils
```

There can be only one ESP on a single disk. The size of the ESP is determined using the following algorithm:

$$\text{ESP} = \max(100 \text{ MB}, \min(1\% \text{ of physical disk}, 1\text{GB}))$$

In other words, the size of the ESP must be the larger of these two numbers, 100 MB or 1 percent of the physical disk size (up to 1 GB). For example, for an 18 GB disk, the size of the ESP is 184 MB. The value 1 percent of the physical disk is calculated at the time that the ESP is created and does not change if the disk is extended later (for example, via RAID).

Each bootable GPT disk must contain an EFI System Partition and it should be the first partition on the disk, right after protective MBR, as shown in Figure 1-18.




Figure 1-18 Boot GPT disk structure

The EFI specification supports only FAT or FAT32 on the ESP partition. The ESP is not visible to the operating system's users by default, but can be accessed for read/write operations from within the operating system by special commands. For Windows-specific information, see "Accessing EFI System Partition from Windows" on page 147. For Linux information, see "Partitions on IA-64 Linux" on page 166.

Note that a non EFI disk partitioning tool would not understand Figure 1-18 and will see it as Figure 1-17 on page 32.

1.10.3 EFI and the reduced-legacy concept

The EFI is *reduced-legacy*, which means that it eliminates certain hardware and firmware elements of the original PC architecture while advancing the PC's stability and usability. These I/O components have been part of the PC architecture for a long time. In the x455, there is no support in EFI for a parallel port, ISA slots, mouse and keyboard ports, or a diskette drive.

Tip: A diskette drive can be used if the 64-bit operating system supports the device, but the EFI does not support the use of a diskette drive. We recommend that you use a USB memory key instead.

EFI replaces this support with the support for new technology such as bootstrap loading from USB devices and the use of USB mouse and keyboard interfaces.

For more information about the EFI specification, see the following:

<http://www.intel.com/technology/efi/index.htm>

<http://www.microsoft.com/hwdev/platform/firmware/EFI/default.asp>

1.11 Operating system support

The following operating systems are supported

- ▶ Windows Server 2003 Enterprise Edition
- ▶ Windows Server 2003 Datacenter Edition
- ▶ Red Hat Linux AS 2.1
- ▶ Red Hat Enterprise Linux AS 3.0
- ▶ SUSE LINUX Enterprise Server 8.0

These are further discussed in “Operating system support” on page 87. This list will increase as vendors port the operating systems to 64-bit architecture. For the latest operating system support information, go to:

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

1.12 Enterprise X-Architecture

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

Enterprise X-Architecture technology enables the following capabilities:

- ▶ XpandOnDemand scalability
- ▶ System partitioning
- ▶ PCI-X I/O subsystem
- ▶ Active PCI-X
- ▶ Remote I/O
- ▶ Active memory
- ▶ High-speed (DDR) memory
- ▶ Memory ProteXion
- ▶ Chipkill memory
- ▶ Memory mirroring
- ▶ Hot-add/hot-swap memory
- ▶ XcelL4 server accelerator cache

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

unmatched “economies of scale” and new levels of server availability and performance.

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

The Enterprise X-Architecture paradigm sets IBM xSeries servers apart in the industry while maintaining the advantages of compliance with industry standards for processors, memory, I/O, storage and software. Enterprise X-Architecture also establishes a revolutionary new economic model for servers through its flexible modular design. The x455 Enterprise X-Architecture-based server allows customers to XpandOnDemand, the ability to grow when you need to, using the existing infrastructure.

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

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

1.12.1 NUMA architecture

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

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

- ▶ Caches have lower latency and faster access time than a memory access. Because a cache can be accessed more quickly than memory, the processor waits less time before receiving a response from a request.

- ▶ Caches also improve performance because they reduce queuing time of accesses that miss the caches and require a physical memory access.

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

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

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

Non-uniform Memory Access (NUMA) is an architecture designed to improve performance and solve latency problems inherent in large (greater than four processors) SMP systems. The x455 implements a NUMA-based architecture and can scale up to 16 processors using multiple servers.

The servers each contain up to four CPUs and 28 memory DIMMs. Each server also has a dedicated local Cache and Scalability Controller, Memory Controller, and 64 MB Xcel4 Level 4 cache. The additional fourth level of cache greatly improves performance for the four processors in the server because it is able to respond to a majority of processor-to-memory requests, thereby reducing the load on the memory controller and speeding up average memory access times.

As shown in Figure 1-12 on page 23, each server is connected to another server using three independent 3.2 GBps scalability cables. These scalability cables mirror front-side bus operations to all other servers and are key to building large multiprocessing multinode systems.

By mirroring transactions on the front-side bus across the scalability links to other processors, the x455 is able to run standard SMP software. All SMP systems must perform processor-to-processor communication (also known as “snooping”) to ensure that all processors receive the most recent copy of requested data. Since any processor can store data in a local cache and modify that data at any

time, all processor data requests must first be sent to every processor in the system so that each processor can determine if a more recent copy of the requested data is in that processor cache.

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

Its easy to see that increasing the number of processors and using faster processors results in greater communication overhead and memory controller bottlenecks. But unlike traditional SMP designs, which send every request from every processor to all other processors, greatly increasing snooping traffic, the x455 has a more optimal design. The XceL4 cache in the x455 improves performance because it filters most snooping operations.

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

The majority of data requests not found in the XceL4 cache can be sent directly to the memory controller to perform memory look-up without any remote processor-to-processor communication overhead. The directory-based coherency protocol used by the XceL4 cache greatly improves scalability of the x455 because processor-to-processor traffic is greatly reduced.

By using this cache and providing fast local memory access to all processors, typical bottlenecks that are associated with large SMP systems using a single memory controller design are eliminated.



Positioning

In this chapter we discuss topics that help you to understand how the x450 can be useful for your business. The topics covered are:

- ▶ “Migrating to a 64-bit platform” on page 40
- ▶ “Scalable system partitioning” on page 41
- ▶ “Operating system support” on page 42
- ▶ “Server consolidation” on page 43
- ▶ “ServerProven®” on page 44
- ▶ “IBM Datacenter Solution Program” on page 45
- ▶ “Application solutions” on page 46
- ▶ “Why choose x455” on page 52

2.1 Migrating to a 64-bit platform

With the addition of 64-bit computing to the xSeries platform you can now run large scale, mission-critical applications with the x455. The Itanium 2 processor is suitable for the next generation of high-end applications that will take advantage of its performance benefits particularly with regards to its flat memory addressable space, with a theoretical addressable memory limit of 16 EB (Exabytes).

Migrating existing 32-bit applications can be a time-consuming and complex process particularly for large-scale mission-critical applications. It is also acknowledged that some applications will never need the power that 64-bit processing offers. To assist with this transition and offer backwards compatibility, the Itanium 2 offers support for IA-32 applications.

In addition, Microsoft has introduced an IA-32 Execution Layer software driver to further enhance 32-bit support on Windows Server 2003. The driver works by translating IA-32 code into native Itanium architecture code before it is executed. It is available from:

<http://www.microsoft.com/windowsserver2003/64bit/ipf/ia32el.msp>

A similar driver for Linux is also planned.

To fully take advantage of 64-bit computing, it is important that your existing 32-bit applications are re-written and/or recompiled correctly. Using the 32-bit compatibility mode of the Itanium 2 or simply recompiling will not yield the potential benefits. In some cases it will even lead to poorer performance.

32-bit binaries (that is, those that run on Pentium and Xeon processor-based systems) should execute on Itanium 2 without modification. In fact, both 32-bit and 64-bit binaries can run simultaneously on the same processor. This allows a system that is only partially converted to 64 bits to run on the x455. However, unless the applications are recompiled, 32-bit applications will run slower on the Itanium 2 because of the emulation overhead and an inability to take advantage of all the benefits the 64-bit platform offers.

The advantages that 64-bit offers over 32-bit are realized most for applications that can take advantage of its large addressable flat memory, ability to execute instructions in parallel, and raw computational ability particularly for applications that use floating point calculations.

Before migrating, you need to consider many factors. Below are some of these considerations:

- ▶ Does your current operating system support 64-bit hardware?

- ▶ Do your applications support 64-bit hardware?
- ▶ Will your application migrate to provide true 64-bit performance?
- ▶ How long will it take to migrate your applications and how complex will it be?
- ▶ Will your 64-bit platform integrate seamlessly with your existing 32-bit platforms?

2.2 Scalable system partitioning

With the power of Enterprise X-Architecture, the x455 sets the pace in the market for 8-way and 16-way multi-node partitions. It blends XpandOnDemand processor and I/O scalability with OnForever™ availability for mission-critical applications.

The x455 has a modular and flexible design, which allows you to start with a configuration that meets your current needs and expand when your needs increase. This forms part of IBM's pay-as-you-grow strategy that allows customers to only pay for what they need. With the x455 you can start with a minimum configuration and scale up through several stages up to the maximum configuration, which is designed to meet the most demanding needs. These are listed below.

Minimum configuration:

- ▶ 1 x server chassis
- ▶ 1 x Itanium 2 processor
- ▶ 1 GB RAM
- ▶ 6 x PCI-X slots
- ▶ 1 x 18.2 GB internal hard disk

Maximum configuration:

- ▶ 4 x server chassis
- ▶ 16 x Itanium 2 processors
- ▶ 224 GB RAM
- ▶ 48 x PCI-X slots with the use of two RXE-100
- ▶ 8 x 146.8 GB internal hard disk = 1175 GB (587 GB in RAID-1 configuration)

With the 64-bit architecture, IBM's X-Architecture, and its scalability options, the x455 has moved into territory that was traditionally the preserve of mini and mainframe systems. This offers the following advantages:

- ▶ Customers who have invested in Microsoft technologies can now scale up their applications.

- ▶ Investment protection for customers who have skills in xSeries (or other PC hardware).
- ▶ Provides a cost-effective alternative for customers who need to invest in large-scale Linux/UNIX® systems and would have traditionally considered a mini or mainframe solution.

Partitioning is invisible to the operating system and therefore does not require special installation techniques or configuration from a software point of view. The operating system is only required on the primary node (Node 1), which means simple but effective management.

At the time of writing the x455 supported hardware partitioning only. Software partitioning using products such as those from VMware were not supported.

2.2.1 RXE-100 Expansion Enclosure

The RXE-100 forms part of IBM's XpandOnDemand technology and allows customers to expand their PCI-X I/O capability as demands increase. With the option to purchase a "6 pack" or "12 pack" of PCI-X slots and to share one RXE-100 between two servers, this provided a flexible and cost-effective solution.

2.3 Operating system support

The x455 will only support 64-bit operating systems. Microsoft and most of the Linux vendors have invested a great deal of time and effort in producing 64-bit versions of their operating systems. At the launch of the x455 the following operating systems are expected to be ready immediately or within three months.

Table 2-1 Operating system support at x455 launch

Operating system	Support at x455 general availability:		
	4-way (1 node)	8-way (2 node)	16-way (4 node)
Windows Server 2003, Enterprise Edition	Supported	Supported	No ¹
Windows Server 2003, Datacenter Edition	Supported	Supported	Supported
Red Hat Linux AS 2.1	Supported	No	No
Red Hat Enterprise Linux 3.0	Supported ²	Supported ²	No
SUSE LINUX Enterprise Server 8.0	Supported	Supported ^{2, 3}	No
Notes: 1. Enterprise Edition supports a maximum of eight processors. 2. Support will be after the general availability (GA) of the x455. See: http://www.pc.ibm.com/us/compat/nos/matrix.shtml 3. Support approximately 30 days after GA.			

2.4 Server consolidation

With its scalability options, the x455 is ideally suited to server consolidation. Enterprise applications such as those used for large databases, enterprise resource planning (ERP), customer relationship management (CRM) and messaging/collaboration have traditionally been distributed across several servers. The x455 now allows you to consolidate these applications onto less hardware.

Consolidation will allow business to centralize their business computing workloads and reduce cost, complexity, network traffic volumes, and management overhead. The x455 protects the investment you have already made by allowing you to expand to two then four nodes. Resources such as processors, memory, hard disks, network cards, monitors, and even CD-ROMs are accessible to the operating system and applications from all the nodes.

Server consolidation has many interpretations and covers many areas. There is no single strategy or methodology that can be applied that will yield an exact outcome, as each environment is different. However, there are broad categories that may have particular importance where benefits can be realized. Below is a summary of some of the different types of server consolidation and the potential benefits.

Table 2-2 Server consolidation strategies

Type of consolidation	Potential benefit
Centralization	Reduced management Lower operational costs Greater security and monitoring Reduced network traffic Reduced staff costs
Physical	Reduced facilities costs Reduced hardware and software costs Lower operational costs Reduced management
Data	Reduced storage costs Better storage management Better data integrity Better data access Improved data recovery Improved disaster recovery
Application	Reduced administration Increased reliability and availability Reduced software and licence costs Lower operational costs

The x455 can play a role in all of these areas and help the IT manager to reduce costs while providing a better level of service.

2.5 ServerProven®

Originally announced in 1996, the IBM ServerProven Program enhances the position of xSeries servers as one of the most open server systems on the market today. The program establishes a relationship between IBM and leading vendors of options and applications often selected for installation into xSeries servers.

The program is designed to:

- ▶ Test leading vendors' products for compatibility with xSeries servers.
- ▶ Improve ease of installation.
- ▶ Reduce total cost-of-ownership.

IBM ServerProven Program participants agree to work closely with IBM engineers during product test cycles and over product life cycles. This means that the IBM ServerProven Program is a commitment to resolve problems before they get to you.

ServerProven participants, in conjunction with IBM, test a set of solution building blocks—specific devices and applications—in one or more operating system environments. This allows customers to build a complete server solution. Compatibility testing by IBM and its program participants will reduce cost-of-ownership by reducing installation and setup time.

When you see the ServerProven logo, you can be assured that the product displaying this emblem has been tested and certified to run reliably on the xSeries hardware.

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

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

For more information on the participants of the IBM ServerProven program, see:

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

2.6 IBM Datacenter Solution Program

Windows Server 2003, 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 that 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>

2.7 Application solutions

With the x455 in the IA-64 environment, you are ready to deploy even larger implementations of enterprise solutions.

As companies' performance demands grow, 64-bit technology becomes an increasingly attractive option, due to increased memory addressability and true parallel architecture. There are a number of ways the x455 can be deployed in specific application solution environments. These include:

- ▶ Database applications
- ▶ Business logic
- ▶ e-Business and security transactions
- ▶ In-house developed compute-intensive applications
- ▶ Science and technology

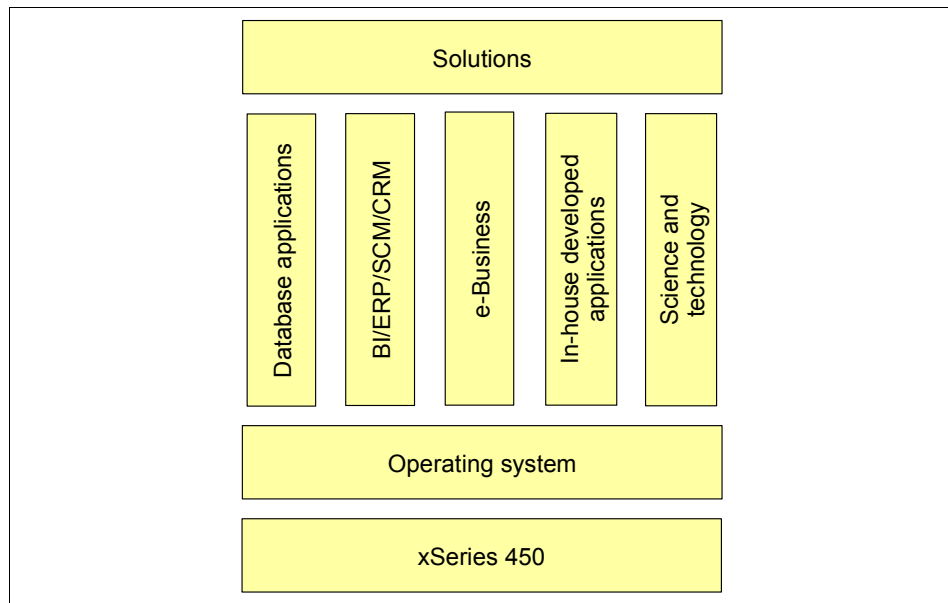


Figure 2-1 xSeries 450-based solutions

2.7.1 Database applications

Sixteen-way x455 configurations can be used as database servers and application servers providing a high-performance reliable platform. These configurations require an external storage enclosure or SAN, depending on the size of the database that is driven by the number of users.

Database applications with memory-intensive workloads that require working data sets larger than 4 GB to be loaded in memory will benefit from the larger memory support of the 64-bit platform.

The following is an example from the field. Microsoft SQL Server Enterprise Edition uses Advanced Windowing Extensions (AWE) memory only for the buffer pool. The AWE API allows applications up to 64 GB of RAM. However, due to the AWE mapping overhead, it is not practical to try to use it for sort areas, procedure cache, or any other type of work area. Many applications do make heavy use of this extra large buffer pool but will not fully exploit its benefits. The most efficient solution in such cases is to move the applications onto a 64-bit database server, which can access memory area above 4 GB as a flat address space without the need to move data in and out of a 4 GB memory area.

Even at the same clock speed a 64-bit processor will move twice as much data as a 32-bit platform. With the improvement Intel has made to the way the data is handled and the additional cache, you should see a noticeable performance increase.

The database server will also benefit from a larger 3 MB third-level and 64 MB XceL4 cache. With such large cache, the need to go to memory or disk for database transaction elements is greatly reduced, and this directly implies a performance increase, faster access to data, and improved throughput. Itanium 2 systems are likely to be able to hold database transaction records in cache during the entire transaction, which enables the I/O portion of the transaction to occur at speeds faster than memory access.

In-memory databases

Architectures with 64-bit addresses can store reasonably large databases in memory and access them with little or no paging overhead. This is often done for databases that are constantly being accessed and for databases that serve as the basis for complex analysis. The theoretical maximum of 16 Exabytes for memory has not yet been tested, but multi-Gigabyte databases are frequently run on 64-bit machines.

A major challenge to providing high-performance access to database information is the time it takes to access disk drives. When disk access is required, disk access times add what can be an intolerable delay to efficient information access and utilization. Access to disk is typically hundreds to thousands of times slower than access to memory.

Today, the disk access time challenge can be overcome. The price of random access memory has come down to affordable levels for many systems. This price reduction means that an entire database can be stored in system memory if the system processors can provide a very large linear address space.

A processor that supports a 64-bit address space may provide access to in-memory databases that range from tens of Gigabytes to thousands of Terabytes. In contrast, traditional 32-bit processors most often only address a maximum of 4 GB.

Multiple databases that required large amounts of memory and previously had to be distributed across several machines can now be consolidated onto fewer machines, taking advantage of the larger addressable linear address space.

Key database software available for Itanium 2 systems include:

- ▶ IBM DB2® Universal Database™ 8.1 (both Linux and Windows versions)
- ▶ Microsoft SQL Server 2000
- ▶ Oracle 9i Database (Linux/Windows)

IBM DB2 Universal Database

The DB2 Universal Database offers increased performance, reliability, and scalability on Windows and Linux platforms by exploiting the IA-64 servers built around the Itanium Processor Family of chips from Intel. The 64-bit version of DB2 delivers significantly higher levels of performance and reliability by exploiting the capabilities of the new systems.

As the adoption of IA-64 architecture continues to move forward with most major server system OEMs using Windows and Linux operating systems, IBM DB2 Universal Database leads the software industry in providing business database application content to the IA-64 server platform. IBM DB2 Universal Database is currently enabling the most robust and varied database software application performance on the IA-64 architecture.

By investing in IBM DB2 Universal Database application software designed for the IA-64 architecture, starting with the Itanium microprocessor, companies will be assured of having a solid foundation for their electronic business far into the future.

Client/server configurations supported in this release are:

- ▶ 64-bit client to 64-bit DB2 UDB engine
- ▶ 64-bit client to 64-bit DB2 Connect™ gateway to DRDA® host
- ▶ 64-bit DB2 Connect Personal Edition to DRDA host

2.7.2 Business logic

More and more enterprise applications such as ERP, SCM, CRM, and BI are released or announced to be released on a 64-bit platform. Such applications process large amounts of data and the large flat memory model means that this

processing will be more efficient. That combined with up to four Itanium 2 processors and a highly efficient cache system make the x450 an ideal choice.

Market leaders offer 64-bit optimized versions of their enterprise applications for use on the x455 now, including SAS Release 9.0 (Windows version) and SAP R/3 4.6C (Windows/Linux), among others.

Independent software vendors such as JD Edwards, Baan/Invensys, i2 Technologies, PeopleSoft, Veritas, Computer Associates, BMC Software, and many others already offer a variety of products available in 64-bit version.

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 x455 brings high I/O bandwidth and performance to handle compute-intensive BI applications. A world-class floating-point engine, capable of addressing much larger amounts of memory, helps speed up data-intensive BI applications that help companies to increase employee productivity.

Enterprise Resource Planning

ERP is an industry term for the broad set of activities supported by multi-module application software that helps manufacturers manage the important parts of their businesses: Product planning, procurement, inventory maintenance, supplier interaction, customer service, and order tracking. 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.

Key server attributes for ERP applications are availability, scalability, and performance. The x455, with its flat memory model, Itanium 2 processors, and Enterprise X-Architecture technology such as Active Memory and Xcel4 Server Accelerator Cache, provides a robust base on which to build and implement successful ERP solutions.

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 x455 is a preferred platform for 64-bit SCM management applications. The x455 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 x455 covers all high-availability features for customers looking for servers to power their SCM solutions.

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.

The x455 provides a performance-based foundation upon which customers can build and deploy CRM solutions in which the x455 will most likely be implemented as an application server and/or a database server.

2.7.3 e-Business and security transactions

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 x455 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 x455. The performance of the server shows promise in Web servers that perform secure e-commerce transactions.

When using x455 as a Web server for your e-commerce solutions, you will benefit from its highly parallel computation that can handle higher volumes of secure data transmissions using complex encryption/decryption techniques.

The Itanium 2 micro-architecture is a perfect match for the significant compute power requirements of Secure Sockets Layer (SSL), providing protection without incurring delay. SSL uses advanced public key encryption technology to safely move sensitive information across the Internet.

In today's complex environments, encryption can occur at two levels simultaneously. For example, the IPSec standard performs encryption on every packet sent over the network. These packets frequently contain data that is

encrypted itself. As a result, enterprises need platforms that can encrypt and decrypt data very quickly. The x455 is suited to this task.

To compare the performance in integer and floating-point operations with other processors on the market, see SPECint_base2000 and SPECfp_base2000 benchmarks at:

<http://www.spec.org>

2.7.4 In-house developed compute-intensive applications

For developers, the 64-bit architecture allows creation of applications using a familiar programming model that encourages the development of a wide-ranging set of enterprise solutions. The experiences gained on the IA-32 platform can be reused when creating applications for Itanium 2-based systems. Thus developers are not required to start from scratch to make their transition to the 64-bit world. In most cases, existing 32-bit code will not require a complete rewrite, but only recompilation. This, however, should not be used as a long-term and easy option if re-writing will yield significant performance benefits.

IA-64 parallelism is managed by the compiler itself. Application development does not require special techniques before it is compiled. Currently available compilers optimized for Itanium 2 include:

- ▶ Intel C++ Compiler
- ▶ Intel Fortran Compiler
- ▶ Microsoft SDK/C/C++ Compiler

2.7.5 Science and technology industries

Science and technology industries (S&TC) require the processing of large and complex calculations to solve challenging problems.

While S&TC industries are characterized by compute-intensive workloads that require special server characteristics to meet their performance needs, each industry, such as aerospace, automotive petroleum, research, or weather, also has its own set of S&TC applications, each demanding different computing solutions.

The Itanium 2-based systems are an ideal platform for compute-intensive applications. It is for this reason that the Itanium and Itanium 2 processors were chosen for the largest American-built supercomputer, the TeraGrid project sponsored by the National Center for Supercomputing Applications. When finally deployed with more than 3,000 Itanium 2 chips, TeraGrid will be capable of more than 14 billion floating-point operations per second (teraflops).

Itanium 2 features, such as increased system memory bandwidth, total number of 328 onboard registers (including 128 floating-point registers), and thus high floating-point performance, speed up calculations and data analysis in S&TC applications.

2.8 Why choose x455

Modern enterprise applications require servers with significant processing power that are able to process large quantities of data and store numerous transactions in cache, and do it all with the highest possible availability and reliability. Systems that cannot deliver at this level are destined to serve as solutions in non-critical settings.

IBM, as a market leader, has been delivering top-of-the-line solutions for the enterprise environment for years, and xSeries 455 is evidence that this will continue to develop with IA-64 technology. With close cooperation between IBM and Intel, the xSeries 455 will continue to be an excellent platform for the most demanding 64-bit enterprise applications.

The x455 combines and builds on a number of technologies developed in previous xSeries models and other IBM developed technology:

- ▶ The mainframe inspired x445 that provides reliable and scalable partitioning.
- ▶ The x380, the first 64-bit xSeries server targeted, tested, and proven by software and hardware developers.
- ▶ The x450, the first xSeries server aimed at large, demanding, mission-critical applications.
- ▶ The IBM pSeries®, iSeries™, and zSeries® servers. Unlike most of our competitors, IBM has many years and a vast amount of experience with 64-bit technology.

All of these advantages make the x455 an enterprise-level server and a viable alternative to RISC-based architectures, while protecting many years of investments and knowledge gained from IA-32 platform development.

The Itanium 2 processor used in the x455 is based on the Explicitly Parallel Instruction Computing (EPIC) architecture, which incorporates a number of new technologies, features, and capabilities that make it ideal for the high-end server and workstation markets. EPIC allows users to take advantage of its large memory addressability and parallel execution capabilities. The chip also supports intelligent prediction and speculation of events to reduce redundancy and improve performance. The Itanium 2's floating point engine enhances

performance for complex computations that are required for data-mining, scientific, and technical applications.

The Itanium 2 processor is the second in a family of Intel 64-bit enterprise-class processors. For more information about the Intel Itanium 2 processor, see:

[http://www.intel.com/products/server/processors/server/Itanium 2/](http://www.intel.com/products/server/processors/server/Itanium_2/)

Intel Itanium Architecture-based microprocessors have the following features:

► Advanced parallelism

High performance requires parallel execution, which is either very limited or hardly achieved in today's architecture. The traditional PC systems are not designed for parallelism, which is critical for current demanding applications (for example, databases and application servers).

Today's processors using limited parallelism are often 60 percent idle. When source code is compiled on today's systems, the result of the compilation is sequential machine code. A regular (non-Itanium processor family) compiler takes sequential code, examines and optimizes it for parallelism, but then has to regenerate sequential code, but in such a way that the processor can re-extract the parallelization from it. The processor will then be required to read this implied parallelism from the machine code, re-build it, and execute it. The parallelism is there, but it is not as obvious to the processor and more work has to be done by the hardware before it can be utilized.

Itanium 2 supports parallelism on multiple levels. Instruction-level parallelism (ILP) is the ability to execute multiple bundles (three instructions in a bundle) at the same time. The Itanium 2 micro-architecture can deliver faster performance by executing multiple bundles per clock cycle. Parallelism, both at the instruction level and at the SMP system level, permits more efficient use of virtually all system resources to enable improved scalability.

The Itanium processor's instruction-level parallelism helps ensure the scalability necessary to manage large data warehouses.

Note: Performance should no longer be measured by just the speed in MHz, but also by the degree of parallelism that the processor achieves.

► Large memory addressability

Another key advantage is that 64-bit operating systems can support far more physical memory than a 32-bit operating system. The limit for directly addressable memory on 32-bit architecture is 4 GB. On 64-bit architecture is it 16 Exabytes (EB) or 2^{64} . The genuinely takes the processor to a level where other hardware and software will take years to match its raw ability to address and handle data.

The Itanium 2 processor cache subsystem has three levels, all of which are on-die cache. In addition to this, IBM provide a fourth level of cache to manage the data and logic between the processors.

The increased physical memory includes the following benefits for applications:

- Each application can support more users. For a comparison of the number of maximum connected users in SAP for various hardware platforms (including Itanium 2), visit the following link:
<http://www.sap.com/benchmark/>
- Each application has better performance and more applications can run simultaneously and remain completely resident in the system's main memory. This reduces or eliminates the performance penalty of swapping pages to and from disk.
- Each application has more memory for data storage and manipulation. Databases can store more of their data in the physical memory of the system. Data access is faster because disk reads are not necessary.
- Applications can manipulate large amounts of data easily and more reliably. Video composition and modeling for scientific and financial applications benefit greatly from memory-resident data structures that are not possible on 32-bit operating systems.

For the enterprise customer, a larger physical memory subsystem provides access to more data more quickly from system memory, since much more data can be held near the processor for faster calculations and data analysis. Large memory addressability allows larger file system caches for read-ahead and write-behind I/O operations, and also allows retention of large amounts of data in memory instead of repeatedly reading the data from disk. Combined with high-memory bandwidth and a variety of performance optimization techniques, this solution provides the performance the enterprise market needs.



Planning

In this chapter we discuss topics that you need to consider before you finalize the configuration of your x455 system and before you begin implementing the system. The topics covered are:

- ▶ “System hardware” on page 56
- ▶ “Cabling and connectivity” on page 65
- ▶ “Storage considerations” on page 79
- ▶ “Rack installation” on page 85
- ▶ “Power considerations” on page 87
- ▶ “Operating system support” on page 87
- ▶ “IBM Director support” on page 89
- ▶ “Solution Assurance Review” on page 90

3.1 System hardware

The x455 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.

- ▶ “Processors” on page 56
- ▶ “Memory” on page 57
- ▶ “PCI-X slot configuration” on page 61
- ▶ “Broadcom Gigabit Ethernet controller” on page 64

3.1.1 Processors

xSeries 455 servers combine copper-based, XA-64 Enterprise X-Architecture technologies with 64-bit, Intel Itanium 2 processors.

Table 3-1 lists the processors standard in each x455 model and the part number of additional processors.

Table 3-1 Processors in each x455 model

Model	Standard CPUs	L2 cache	L3 cache	Max SMP	Extra CPUs
8855-1RX	One 1.3 GHz	256 KB	3 MB	Four-way	73P7076
8855-2RX	Two 1.4 GHz	256 KB	4 MB	Four-way	73P7077
8855-3RX	Two 1.5 GHz	256 KB	6 MB	Four-way	73P7078

The L2 and L3 cache run at the full speed of the processor.

A total of four processors can be installed in the x455. One, two, three, or four processors can be installed. Each processor must be the same speed and cache size. Each processor option includes the processor with heatsink (pre-assembled) and its associated power module.

When you have a multi-node configuration, all nodes must have four processors installed. That is, only 8-way and 16-way are supported.

Processors must be installed in a specific order, as shown in Figure 3-1 on page 57. Special tools are required to install the processors and these are included in the option. We recommend that you update the system abstraction layer/extensible firmware interface (SAL/EFI) code after the installation. To download the most current level of SAL/EFI code for the server, go to the x455 device driver page:

<http://www.ibm.com/pc/support/site.wss/MIGR-53575.html>




Figure 3-1 Processor installation order

3.1.2 Memory

A maximum of 28 DIMMs may be installed in the memory-board assembly of the x455, depending on the size of the DIMMs used. Supported DIMMs are listed in Table 3-2.

Table 3-2 Supported DIMMs

Size	Description	Part number	Max installable
512 MB	PC2100 CL2.5 ECC DDR SDRAM	33L5038	28
1 GB	PC2100 CL2.5 ECC DDR SDRAM	33L5039	28
2 GB	PC2100 CL2.5 ECC DDR SDRAM	33L5040	28

Go to the ServerProven site for the latest information on supported memory modules:

<http://www.pc.ibm.com/us/compat/machines/x455.html>

When installing memory DIMMs, they must be installed in matched pairs (size and technology), in the order shown in Table 3-3.

Table 3-3 DIMM order

Pair	Port 1 DIMM slots	Pair	Port 2 DIMM slots
1	1 and 14	2	15 and 28
3	2 and 13	4	16 and 27
5	3 and 12	6	17 and 26

Pair	Port 1 DIMM slots	Pair	Port 2 DIMM slots
7	4 and 11	8	18 and 25
9	5 and 10	10	19 and 24
11	6 and 9	12	20 and 23
13	7 and 8	14	21 and 22

Figure 3-2 shows the DIMM locations and the way the DIMMs are divided into the two memory ports.

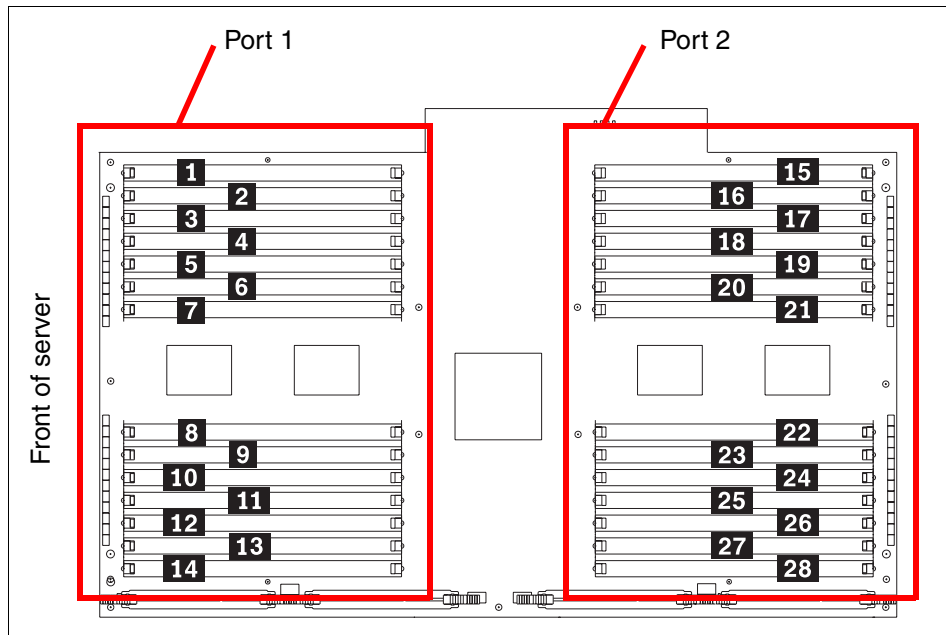


Figure 3-2 DIMM locations on the memory-board assembly

Memory mirroring

Memory mirroring is supported by the x455 for increased fault tolerance and high levels of availability.

Key configuration rules relating to memory mirroring are:

- ▶ Memory mirroring must be enabled in the System Setup in the EFI (it is disabled by default). See “Enabling memory mirroring” on page 118.
- ▶ Enabling memory mirroring halves the amount of memory available to the operating system.

- ▶ Both ports in a memory-board assembly must have the same total amount of memory. Partial mirroring is not supported. If the same total amount of memory is not detected at boot, memory mirroring is automatically disabled by the system.
- ▶ The DIMMs in partner banks across the two ports must be identical.
- ▶ You must install two pairs of DIMMs at a time. These four DIMMs (known as a *bank*) must be identical. Table 3-4 shows the pairs that are in each bank.

Table 3-4 DIMMs that form a bank

Bank	DIMM pairs (see Table 3-3)
1	1 and 2 (DIMMs 1, 14, 15, 28)
2	3 and 4 (DIMMs 2, 13, 16, 27)
3	5 and 6 (DIMMs 3, 12, 17, 26)
4	7 and 8 (DIMMs 4, 11, 18, 25)
5	9 and 10 (DIMMs 5, 10, 19, 24)
6	11 and 12 (DIMMs 6, 9, 20, 23)
7	13 and 14 (DIMMs 7, 8, 21, 22)

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.
- ▶ 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 58.
- ▶ When replacing a failed DIMM, you do not have to replace the other DIMM within the bank. Ensure that the new DIMM is the same type and size as the other DIMM within the bank.

If a problem with a DIMM is detected, light path diagnostics will light the system-error LED on the front of the server, indicating that there is a problem and guiding you to the defective DIMM. When this occurs, first identify the defective DIMM and then remove and replace the DIMM.

Complete the following steps to replace a DIMM in your server with the server turned on:

1. Open the cover and verify that the memory hot-plug enabled LED on the DIMM access door is lit before removing and replacing the DIMM.
2. Open the DIMM access door and verify that the memory port LED is off before replacing a DIMM. Then open the retaining clip on each end of the DIMM connector and remove the DIMM from the server.
3. Install the new DIMM. Take the usual anti-static precautions.
4. Close the DIMM access door.

Note: The memory hot-plug enabled LED flashes to indicate that data is being mirrored on the replacement DIMMs. Wait until the LED stops flashing before you hot-place DIMMs again.

Memory performance considerations

As shown in the server block diagram in Figure 1-4 on page 8, there are two memory ports to the memory controller, each with a throughput of up to 3.2 GBps. These ports correspond to the ports as shown in Figure 3-2 on page 58. The front-side bus of the processors is 6.4 GBps, so maximum performance is achieved when both memory ports are used to access memory simultaneously.

Important: The XceL4 Accelerator Cache has been designed with commercial workloads in mind that tend to have high cache hit rates. This effectively boosts performance and compensates in part for the 3.2 GBps bandwidth between the Cache and Scalability Controller and Memory Controller.

Consequently, for maximum performance, you should install four DIMMs of the same size at a time into a bank (see Table 3-4 on page 59). In this configuration, all memory addresses are spread across all four DIMMs in the bank and, when accessed, both memory ports are used.

Maximum performance can also be achieved with DIMMs of different sizes, as long as the total memory in port 1 matches the total memory in port 2. For example, if you have six DIMMs (four 512 MB DIMMs and two 1 GB DIMMs for a total of 4 GB), install all four 512 MB DIMMs (2 GB) in port 1 and the two 1 GB DIMMs (also 2 GB) in port 2.

If there is a mismatch between the total memory in port 1 and the total memory in port 2, then there will be a delay when accessing the upper memory, for example, if you have two 512 MB DIMMs and two 1 GB DIMMs.

3.1.3 PCI-X slot configuration

As shown in Figure 1-4 on page 8, there are six PCI-X slots internal to the x455. These six slots are implemented using four PCI buses:

- ▶ Bus A (slot 1 and slot 2): Supporting two 64-bit adapters at up to 66 MHz
- ▶ Bus B (slot 3 and slot 4): Supporting two 64-bit adapters at up to 100 MHz, or one adapter at up to 133 MHz, provided the other slot is vacant
- ▶ Bus C (slot 5): Supporting one 64-bit adapter at up to 133 MHz
- ▶ Bus D (slot 6): Supporting one 64-bit adapter at up to 133 MHz

These slots can accept adapters rated at speeds ranging from 33 MHz to 133 MHz.

You should also consider the following:

- ▶ Each adapter has a maximum rated speed. Each bus also has a maximum rated speed.
- ▶ Installed adapters in a single bus will operate at the slowest of three speeds:
 - The rated speed of adapter 1
 - The rated speed of adapter 2 (if the bus the adapter is installed in has two slots)
 - The maximum 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 loss of PCI adapter performance.

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)

Important: A PCI-X and a PCI adapter can be installed in slots on the same bus. However, those 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.

Table 3-5 summarizes the supported adapter speeds. Take into account the speed reductions when there are two adapters installed in a bus, as described above.

Table 3-5 Supported adapter speeds in each slot

Slot	Bus	Width (bits)	Supported adapter speed (MHz)
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 (133 as long as no adapter is in slot 4)
4	B	32 or 64	33, 66, or 100 (133 as long as no adapter is in slot 3)
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 63.




Figure 3-3 PCI-X slots in the x455

Other configuration information:

- ▶ The x455 server supports connection to a RXE-100.
- ▶ Video adapters are not supported.
- ▶ The PCI slots support 3.3 V adapters only.

Important: 5.0 V adapters are not supported.

- ▶ The on-board LSI SCSI/RAID controller and the ServeRAID adapter are supported for connection to and booting from the internal drive bays.
- ▶ 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 slots to assign system resources. The system attempts to start the first device found. The search order is:
 - a. CD-ROM
 - b. Disk drives
 - c. Integrated SCSI devices
 - d. Integrated Ethernet controllers
 - e. x455 PCI-X slots (in the order 1, 2, 3, 4, 5, 6)

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

- a. CD-ROM

- b. Disk drives
- c. Integrated SCSI devices
- d. Integrated Ethernet controller
- e. x455 PCI-X slots (in the order 1, 2, 3, 4, 5, 6)
- f. RXE-100 slots (A1, A2, A3, A4, A5, A6, B1, B2, B3, B4, B5, B6)

Tip: This scan order is different from that of the x450.

3.1.4 Broadcom Gigabit Ethernet controller

The x455 offers a dual Gigabit Ethernet controller integrated standard in the system. The x455 includes a single-port Broadcom BCM5704 10/100/1000 BASE-T controller on a PCI 64-bit 66 MHz bus.

The BCM5704 supports full- and half-duplex performance at all speeds (10/100/1000 Mbps, auto-negotiated) and includes integrated on-chip memory for buffering data transmissions, and dual onboard RISC processors for advanced packet parsing and backwards compatibility with 10/100 devices. The Broadcom controller also includes software support for Wake on LAN, failover, layer-3 load balancing, 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 that your network infrastructure is capable of the necessary throughput to match the server's I/O capacity.

Adapter teaming

The Broadcom controller is capable of participating in an adapter team for the purposes of failover, load balancing, and port trunking. The choice of adapters to team with the onboard controller depends on whether you have a copper-only network or a mixed copper/fiber network. Our recommendations are:

- ▶ If you have a copper Gigabit environment, use the Broadcom-based NetXtreme 1000T Ethernet adapter, part 31P6301. Alternatively, use the Intel PRO/1000 XT Server adapter, part 22P6801.
- ▶ If you have a mixed fiber/copper Gigabit server switch network, use the Broadcom-based 22P7801, NetXtreme 1000 SX Fiber Ethernet adapter.

You can also team any of the onboard Gigabit cards with 10/100 cards such as 06P3601 and 22P4901, but this is not a recommended configuration. You can also team with the older Gigabit fiber card, 06P3701.

Adapter teaming and failover work by using software additional to the adapter driver to provide the failover functionality. This software is operating system

dependent. Detailed instructions for installing the individual driver and failover packages are available with the driver software.

For the latest network adapter drivers and software for the x455 server, go to the xSeries support page:

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

For details about compatibility, see the ServerProven LAN adapter page:

<http://www.pc.ibm.com/us/compat/lan/matrix.html>

3.2 Cabling and connectivity

There are a number of unique factors to consider when cabling the x455 server. These are discussed as follows:

- ▶ “SMP Expansion connectivity” on page 67
- ▶ “Remote Expansion Enclosure connectivity” on page 70
- ▶ “Remote Supervisor Adapter connectivity” on page 76
- ▶ “Serial connectivity” on page 78

We discuss each of these in this section.

The rear panel of the x455 showing the locations of cable connectors is illustrated in Figure 3-4 on page 66. For more details about ports on the Remote Supervisor Adapter, refer to “Remote Supervisor Adapter connectivity” on page 76.




Figure 3-4 x455 rear view

Of note are the following items:

- ▶ There are no PS/2 keyboard or mouse ports on the x455. Only USB devices are supported.

For attachment to an Advanced Connectivity Technology (ACT) KVM switch, a new USB Conversion Option (part number 73P5832) can be used. This smart cable is plugged into the USB and video ports on the server and it converts KVM signals to CAT5 signals for transmission over a CAT5 cable to either a Remote Console Manager (RCM) or Local Console Manager (LCM). The x455 can then be managed on the same set of switches as PS2- or C2T-based KVM servers.

- ▶ The x455 includes a dedicated external serial port. During the boot process, the port acts as an auxiliary console for the EFI. After the operating system has booted, the port is dedicated to the operating system as COM1.
- ▶ The RXE Expansion Ports provide connectivity to a RXE-100. Either a single cable in port A or two cables to both ports are supported. Using two cables provides redundancy as well as additional throughput.
- ▶ There are two RJ-45 ports providing Gigabit Ethernet connectivity, as shown in Figure 3-4. A third RJ-45 Ethernet connector is also on the Remote Supervisor Adapter. This connector is only used to connect to the Remote

Supervisor Adapter for out-of-band management, as described in “The Remote Supervisor Adapter” on page 177.

3.2.1 SMP Expansion connectivity

This section describes multi-node cabling. A node is one of multiple servers in a configuration interconnected through the SMP Expansion Ports to share system resources. The multi-node configuration is described in “Configuring scalable partitions” on page 131.

Note: Only two-node and four-node configurations are supported.

There are two optional SMP Expansion cable kits available to connect the SMP Expansion Port connectors:

- ▶ Four-way to 8-way Scalability Kit, comes with two 2.5 m SMP Expansion cables (part number 73P9911)
- ▶ Eight-way to 16-way Supplemental Kit, comes with three 2.5 m SMP Expansion cables and one 3.5 m SMP Expansion cable (part number 73P9715)

Follow these recommendations in each server before cabling:

- ▶ Update the system abstraction layer/extensible firmware interface (SAL/EFI) firmware.
- ▶ Update the service-processor firmware.

These are available from the x455 driver matrix:

<http://www.ibm.com/pc/support/site.wss/MIGR-53575.html>

Multinode event timing

Multinode configurations have different clock sources driving processors in each node. Therefore, the processor time stamp counter (TSC) should not be used by operating systems or applications for global system event timing. Applications should use an OS API to perform event timing.

Operating systems should use the Cyclone performance counter in node 0 to implement system event timing functions.

Two-node configuration

A two-node configuration requires the two cables from the 4-way to 8-way Scalability Kit option.

1. Connect the cables according to Table 3-6.

Table 3-6 Connecting two nodes

From	To
Node A, Port 1	Node B, Port 1
Node A, Port 3	Node B, Port 3

The cabling will look like Figure 3-5.




Figure 3-5 Connecting two-node configuration

2. Connect the Remote Supervisor Adapter Ethernet connector on each server to a network or to each other with an Ethernet crossover cable. This connection is needed so that the Remote Supervisor Adapters can communicate and manage scalable partitions.

Four-node configuration

A four-node configuration requires one 4-way to 8-way Scalability Kit option and one 8-way to 16-way Supplemental Kit option.

Complete the following steps to cable a four-node configuration:

1. Name the servers, as shown in Figure 3-6 on page 69. We used Nodes A–D, from top to bottom. We also named the SMP Expansion ports 1–3, from left to right.




Figure 3-6 Connecting four-node configuration

2. Label each end of the SMP Expansion cables according to where they will be connected to each server.
3. Connect the cables according to Table 3-7.

Table 3-7 Connecting four nodes

From	To
Node A, Port 1	Node C, Port 1
Node A, Port 2	Node D, Port 2
Node A, Port 3	Node B, Port 3
Node B, Port 1	Node D, Port 1
Node B, Port 2	Node C, Port 2
Node C, Port 3	Node D, Port 3

4. Verify that the final cabling is like Figure 3-6 on page 69.
5. Connect the Remote Supervisor Adapter Ethernet connector on each server to your out-of-band management network (so the Remote Supervisor Adapters can also be accessed for remote management). This connection is needed so that the Remote Supervisor Adapters can communicate and manage scalable partitions.

3.2.2 Remote Expansion Enclosure connectivity

The RXE-100 can be connected to the x455 to provide an additional six or 12 PCI-X slots to the server. Only one RXE-100 can be connected to the x455, although this can be with either one or two data cables.

The RXE-100 has six 133 MHz 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 3-7.

Note: As described in “Connecting the RXE-100” on page 71, when connecting the RXE-100 to an x455 configuration using only one cable, the RXE-100 can have six or 12 PCI-X slots.

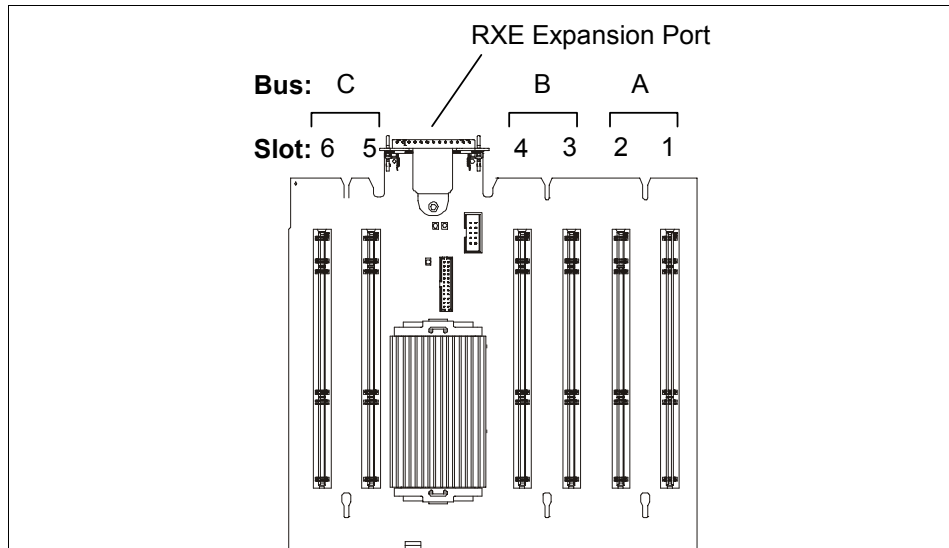


Figure 3-7 RXE-100 PCI-X expansion board (six slots)

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 up to 133 MHz
- ▶ Two 64-bit 3.3 V PCI-X 133 MHz adapters running at up to 100 MHz
- ▶ Two 64-bit 3.3 V PCI or PCI-X, 33 or 66 MHz adapters

Note: The PCI slots support 3.3 V adapters only. 5.0 V adapters are not supported.

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. 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).
- ▶ 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.

Connecting the RXE-100

There are two types of cables used to connect the RXE-100 to the x455:

- ▶ Remote I/O cable, for data. Two lengths are available:
 - 3.5 m Remote I/O cable kit (part number 31P6102), like the one shipped with the RXE-100
 - 8 m Remote I/O cable kit (part number 31P6103)

- ▶ Interconnect management cable, for remote I/O management. The RXE-100 ships with a 3.5 m cable. Two lengths are available:
 - 3.5 m interconnect management cable kit (part number 31P6087)
 - 8 m interconnect management cable kit (part number 31P6088)

Use the 8 m versions of each of these cables if the distances between the two devices warrant the extra length (for example, in separate racks).

Important: This interconnect management cable has standard RJ-45 connectors, but it is not wired the same as an Ethernet or crossover cable. Ensure that 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.

Power to the RXE-100 enclosures is controlled by the servers via the interconnect management cable and under the control of the Remote Supervisor Adapter.

The following connections to the RXE-100 are supported:

- ▶ One server connecting to a six-slot RXE-100
- ▶ One server connecting to a 12-slot RXE-100
- ▶ A two-node server connecting to one shared RXE-100
- ▶ Two standalone servers connecting to one shared RXE-100
- ▶ Four-nodes connecting to one or two RXE-100s

One server connecting to a six-slot RXE-100

When the RXE-100 has only six PCI-X slots installed, connect a single RXE management cable and a single RXE data cable, as shown in Figure 3-8 on page 73.




Figure 3-8 Connecting the RXE-100 to the x455 (six slots in the RXE-100)

One server connecting to a 12-slot RXE-100

If the RXE-100 has 12 PCI-X slots installed, we recommend that you use two data cables to connect the RXE-100 to the server, as shown in Figure 3-9 on page 74.

Tip: The second Remote I/O data cable is optional in this configuration. However, it is recommended because it provides redundancy if the other Remote I/O cable fails. The second cable also improves performance.

The short interconnect management cable to connect Management A (out) Port to Management B (in) Port is supplied with the second set of PCI slots.

Tip: The short interconnect management cable should be *removed* if you will be managing the configuration using Scalable Systems Manager (once it becomes available; see “Scalable Systems Manager” on page 177).

If you are not using Scalable Systems Manager, leave the cable connected as shown.




Figure 3-9 Connecting the RXE-100 to the x455 (12 slots in the RXE-100)

A two-node server connecting to one shared RXE-100

In a two-node (that is, eight-way) x455 configuration, a single RXE-100 is supported in a shared configuration, where six RXE slots are connected to each node. The cabling is shown in Figure 3-10.




Figure 3-10 8-way configuration with a shared RXE-100 (6 PCI-X slots to each node)

In addition, an 8-way configuration can also be connected to a single RXE-100 that only has six PCI-X slots installed, not 12. In this configuration, the data and management cable from the secondary node are not connected.

Two standalone servers connecting to one shared RXE-100

This configuration is very similar to the two-node configuration shown in Figure 3-10 on page 74, except that the cable between the x455s is not present. Each independent server “sees” six slots in the RXE-100.




Figure 3-11 Single RXE shared between two x455s

Four nodes connecting to one or two RXE-100s

In a four node configuration, one or two RXE-100s can be connected. If you have one RXE-100, you can connect it to either nodes 1 and 2, or nodes 3 and 4.

Connectivity of two RXE-100 enclosures is shown in Figure 3-12 on page 76. One RXE-100 is shared by nodes 1 and 2, and the second RXE-100 is shared by nodes 3 and 4.




Figure 3-12 16-way x455 with two shared RXE-100s (6 PCI-X slots to each node)

3.2.3 Remote Supervisor Adapter connectivity

The x455 features an integrated Remote Supervisor Adapter, one of the products in the Advanced System Management (ASM) family. It provides around-the-clock remote access and system management of your server and supports the following features:

- ▶ Remote management regardless of the status of the server
- ▶ Remote control of hardware and operating systems
- ▶ Web-based management with standard Web browsers (no other software is required)
- ▶ Text-based user interface terminal access

The configuration and use of the Remote Supervisor Adapter is discussed in Chapter 5, “Management” on page 175.




Figure 3-13 Remote Supervisor Adapter connectors

The following RSA connections need to be considered when cabling the x455 (see Figure 3-13):

- ▶ External power supply connector. This connector allows the RSA to be connected to its own independent power source. This external power supply is not included with the x455 and will need to be ordered as an option (order a ThinkPad® 56W or 72W AC Adapter with a suitable power cord for your country/region).

If this power supply is not used, the RSA 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. This port is dedicated and can only be used for RSA purposes.
- ▶ Ethernet port, which provides system management functions over the LAN.
- ▶ 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 the IBM Redbook *Implementing Systems Management Solutions using IBM Director*, SG24-6188.

Note: The x455 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.

3.2.4 Serial connectivity

The x455 has an integrated serial port, as shown in Figure 3-4 on page 66. This port has two purposes:

- ▶ During the boot process (before the OS loader starts), the Extensible Firmware Interface (EFI) has control of the port and uses it as an auxiliary console where POST messages are transmitted, many even before the server's video port is enabled. This is especially useful in performing problem determination on the system.
- ▶ Once the operating system loads, it is made available to the operating system as a standard (and dedicated) COM port.

To use the serial port, for use as an auxiliary console, you will need the following:

- ▶ A null modem cable
- ▶ A system running a terminal emulation program, such as HyperTerminal in Windows or minicom in Linux

First, connect the RS-232 cable between the two systems. The cable should be connected to the serial port of the server, not the serial port of the Remote Supervisor Adapter.

Once the cable is connected, start HyperTerminal or your emulation program. HyperTerminal should be set to the following settings:

- ▶ Speed: 115200 bps
- ▶ Data bits: 8
- ▶ Parity: None
- ▶ Stop Bits: 1
- ▶ Flow Control: None

For HyperTerminal, you can leave emulation set to auto detect. If you are using another program that does not have auto detect, you may need to set emulation to ANSIW.

Once your session is configured, click **Connect** and, assuming there is no operating system currently running, you should see POST messages or the EFI Boot Manager menu.

Tip: If you are running Linux, you can redirect the console messages to the serial port. For more information about this feature refer to "Using the serial port for the Linux console" on page 171.

3.3 Storage considerations

When you are planning the storage configuration to accompany the x455, there are important performance and sizing issues that need to be considered.

The two internal hot-swap 1" drive bays will typically be used for operating system installation. We recommend that these drives be configured as a two-drive RAID-1 array to provide a higher degree of system availability. Drives up to 15,000 RPM and the converged tray design are supported. To configure RAID-1, you can use the internal LSI controller that is provided onboard or buy an optional ServeRAID adapter.

Typically the x455 will be attached to an external disk enclosure for data storage requirements. Some of the supported IBM storage options are:

- ▶ SCSI RAID adapters and storage enclosures
- ▶ Fibre Channel adapters and Storage Area Networks (SANs)
- ▶ Network Attached Storage (NAS)
- ▶ SCSI over IP (iSCSI)
- ▶ IBM Enterprise Storage Server® (ESS)
- ▶ ESCON® connectivity to a zSeries server

Note: You can also add as an option an LS-120 diskette drive that can be used to upgrade the firmware or the EFI version. Please note that no diskette drive is provided with the default configuration, and that this solution is not supported.

3.3.1 xSeries storage solutions

This section discusses some of the available xSeries storage solutions and related technologies, as well as tape backup and performance considerations.

All the IBM Storage Solutions are supported under Linux, as well as under Windows. For more information see:

<http://www.storage.ibm.com/linux>

ServeRAID with external storage enclosures

The x455 currently only supports the ServeRAID 4Mx and 6M.

- ▶ The ServeRAID-4Mx is a 66 MHz PCI adapter and features two Ultra160 SCSI channels, 64 MB of battery-backed up ECC SDRAM cache memory, and an Intel 100 MHz i80303 processor. Up to 28 Ultra160 and Ultra2 SCSI devices are supported.

- ▶ The ServeRAID-6M is a 133 MHz PCI-X adapter and features two Ultra320 SCSI channels, according model 128 MB or 256 MB of battery-backed up ECC SDRAM cache memory, and an Intel 600 MHz xScale processor. Up to 28 Ultra320, Ultra160 and Ultra2 SCSI devices are supported.

Each ServeRAID adapter supports up to 14 drives (4Mx have 160 MBps throughput and 6M has 320 MBps throughput) per channel. Multiple adapters can be installed as needs and available slots dictate.

- ▶ The EXP300 and the new EXP400 storage expansion units have a maximum 2 TB of disk storage (14 146.8 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 or EXP400 units can hold an amazing 28 TB). EXP300 and EXP400 provide 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, and the new EXP400 is optimized for Ultra320 SCSI, with a sustained data transfer rate of 320 MBps.

Important: Check the tips in “Notes® on the use of RAID:” on page 108 before doing a installation.

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

<http://ibm.com/pc/ww/eserver/xseries/scsi RAID.html>

For the latest list of SCSI hardware supported by the x455, see ServerProven:

<http://www.pc.ibm.com/us/compat/machines/x455.html>

IBM TotalStorage® FASTT

The IBM TotalStorage FASTT (Fibre Array Storage Technology) family of Fibre Channel storage solutions is designed for high-availability, high-capacity requirements. FASTT solutions can support transfers over distances up to 10 km (6.2 miles) at rates of up to 200 MBps.

The FASTT Storage Server is a RAID controller device that contains 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 FASTT500 Storage Server
- ▶ The IBM TotalStorage FASTT600 Storage Server
- ▶ The IBM TotalStorage FASTT700 Storage Server
- ▶ The IBM TotalStorage FASTT900 Storage Server

For a complete list of supported Fibre Channel hardware, see
<http://www.pc.ibm.com/us/compat/machines/x455.html#Storage>

Note: Check the Microsoft Hardware Compatibility List (HCL) for the currently certified storage solution with FAStT for Microsoft clustered configurations and for updates to certified solutions:

<http://www.microsoft.com/whdc/hcl/search.msp>

The IBM TotalStorage 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 66 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 are supported on the FAStT200. This means that the maximum supported number of disk drives is 60.

Hot-swappable and redundant components provide 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.

The IBM TotalStorage FAStT600 Storage Server

The FAStT600 Storage Server is a 3U rack-mountable Fibre Channel RAID controllers and disk drive enclosure. It targets the entry and midrange segment of

the Fibre Channel storage market. A typical usage of the FASt600 is a two-node cluster directly attached to it. It can support up to 14 internal Fibre Channel disk drives.

The FASt600 features dual 2 Gbps hot-swappable RAID controllers; supports RAID levels such as 1, 3, 5, and 10; and supports global hot spares. The FASt600 has four host ports, two for each controller; this enables a cluster solution without the use of a switch.

There are two models, FASt600 Base and FASt600 Turbo. To understand the differences see Table 3-8.

Table 3-8 Comparison between FASt600 Base and FASt600 Turbo

Feature	FASt600 Base	FASt600 Turbo
Cache size per controller	256 MB	1 GB
Host interface	2 Gb	Auto senses to connect to 1 Gb or 2 Gb
Maximum EXP700 supported	3	7
Maximum drives supported	56	112
Storage capacity	Up to 8.2 TB	Up to 16.4 TB
Storage partitions	Up to 16	Up to 64
FASt Storage Manager	FSM v8.3	FSM V8.4

The IBM TotalStorage FAStT700 Storage Server

The FAStT700 Storage Server would typically implement in high-end cluster and server consolidation environments and where multiple servers are being consolidated onto a multi-node configuration of x455 systems.

The FAStT700 is a 4U rack-mountable Fibre Channel RAID controller device with higher performance controllers. These controllers are 2 Gbps and connect via mini-hubs to the new FAStT FC-2 Host Bus Adapter (HBA) and the 2109 F16 Fibre Channel switch to give full 2 Gbps fabric.

The FAStT700 attaches to up to 220 FC disks via 22 EXP500 expansion units or up to 224 FC disks via 16 EXP700 expansion units to provide scalability for easy growth (18 GB up to 32.6 TB using 146.8 GB drives). 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.

FAStT Storage Manager Version 8.21 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 IBM TotalStorage FAStT900 Storage Server

The IBM TotalStorage FAStT900 Storage Server expands the FAStT family's highly scalable offerings with improved performance capabilities. Built upon a new fourth-generation Fibre Channel RAID controller, the FAStT900 brings high availability, advanced functionality, scalable capacity, and connectivity to a wide range of storage area network (SAN) applications in mission-critical enterprise networks.

The FAStT900 is a 4U rack-mountable high-performance, highly available, Fibre Channel-based storage product for demanding applications. The FAStT900 features a new shared memory bus controller/PCI bridge/RAID assist engine (1.6 GB/s bandwidth), dual 2 Gbps hot swappable RAID controllers with 2 GB of cache memory (1 GB per RAID controller), redundant/hot-swap power supplies, redundant/hot-swap fans, and supports RAID levels such as 0 (implemented as a RAID level 10), 1, 3, and 5 and global hot spares. It also supports four host-side Mini-Hubs for two controllers (each Mini-Hub has two ports); four Host Ports (two on each controller) provide a cluster solution without using a switch.

The FAStT900 attaches up to 224 FC disks via 16 EXP700 expansion units to provide scalability for easy growth (32.6 TB using 146.8 GB drives).

The IBM TotalStorage FAStT EXP700 Expansion Unit

The IBM TotalStorage FAStT EXP700 storage expansion unit expands the highly scalable, high-performance FAStT family with a 14-bay, 2 GBps, rack-mountable Fiber Channel hard disk drive enclosure. Used with the FAStT900, FAStT700, FAStT600, FAStT500, and FAStT200 storage servers, the EXP700 is designed to provide incremental support for over one terabyte of disk storage per unit. When used with the FAStT900, FAStT700, and FAStT600 Turbo storage servers, it provides a full end-to-end 2 GBps Fiber Channel storage solution, extending the FAStT family with the latest Fiber Channel technology.

Note: The new version of FAStT Storage Manager 8.4 supports Windows 2003 Server 64-bit Edition and is only supported on the FAStT600, FAStT700, and FAStT900.

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

<http://www.storage.ibm.com/hardsoft/disk/fastt/index.html>

Enterprise Storage Server (ESS)

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 in excess of 28 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.

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 x455 fit into a heterogeneous environment containing a variety of server architectures. ESS offers several choices of host I/O interface attachment methods, including SCSI and Fibre Channel for xSeries.

For more information on the ESS, go to:

<http://www.storage.ibm.com/hardsoft/products/ess/index.html>

3.3.2 Tape backup

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

- ▶ Currently implemented backup solutions

If you are consolidating a number of servers onto a single x445 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. IBM Tivoli® Storage Manager has 64-bit Windows and Linux versions of their client, and other software vendors have similar products or plans for 64-bit versions.

Disaster recovery procedures may also need to be revised.

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

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

Note: The x455 and RXE-100 support 3.3 Volt PCI adapters only. Make sure that any SCSI adapters you use to connect your tape subsystem are 3.3 V or dual-voltage adapters.

The following IBM Redbooks discuss IBM tape solutions in greater detail:

- ▶ *IBM Tape Solutions for Storage Area Networks and FICON*, SG24-5474
- ▶ *Netfinity Tape Solutions*, SG24-5218
- ▶ *The IBM LTO Ultrium Tape Libraries Guide*, SG24-5946

3.4 Rack installation

The x455 is 4U high and is intended for use as a rack-drawer server. Due to power distribution considerations, it is recommended that no more than eight 4U x455 chassis be installed in a single 42U rack, leaving 10U available for RXE-100 Remote I/O enclosures, disk or tape storage, or other devices.

The x455 is 27.5 inches deep, and is designed to be installed in a 19-inch rack cabinet designed for 28-inch-deep devices, such as the NetBAY42 ER, NetBAY42 SR, or NetBAY25 SR. Although the x455 system is rack optimized, it may be converted into a tower by installing it in a NetBAY11 SR Standard Rack

Cabinet. The NetBAY11 rack supports shipment of fully configured xSeries 455 and other rack-optimized xSeries servers.

Installation considerations include the following:

- ▶ The system is not designed to run vertically, and therefore must always be run in a horizontal position.
- ▶ For thermal considerations, the x455 must be installed with perforated doors on both front and back. Do not install the x455 in a rack with a glass front door.
- ▶ Although installation is supported in non-enterprise racks, it is not recommended, since cable management then becomes an issue.
- ▶ The maximum weight of the system, depending on your configuration, is 50 kg (110 lb). Therefore, this system requires two people to install it in a rack.

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 x455 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. See “Remote Expansion Enclosure connectivity” on page 70 for RXE-100 cabling information.

Since the x455 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**.
 - e. Click **Intel-based servers**.
 - f. Click **Tools**.
 - g. Scroll down to find the Rack Configurator section.

3.5 Power considerations

The x455 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 x455 is populated with more than two processors, memory, and adapters, the power supplies may not be redundant if they are connected to a 110 V power source. Therefore, IBM recommends that the x455 be connected to a 220 V power source to ensure power supply redundancy for large configurations.

Tip: If power is not redundant, the non-redundant LED will be lit in the Light Path Diagnostic panel (see “Light path diagnostics” on page 27).

Two system power-cord connectors are available on the back of the x455, 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 x455 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.

3.6 Operating system support

In line with the overall message of providing application flexibility to meet the varying needs of our enterprise customers, the x455 is optimized for the operating systems listed in Table 3-9 on page 88.

Table 3-9 Operating system support at x455 launch

Operating system	Support at x455 general availability:		
	4-way (1 node)	8-way (2 node)	16-way (4 node)
Windows Server 2003, Enterprise Edition	Supported	Supported	No ¹
Windows Server 2003, Datacenter Edition	Supported	Supported	Supported
Red Hat Linux AS 2.1	Supported	No	No
Red Hat Enterprise Linux 3.0	Supported ²	Supported ²	No
SUSE LINUX Enterprise Server 8.0	Supported	Supported ^{2, 3}	No
Notes: 1. Enterprise Edition supports a maximum of eight processors. 2. Support will be after the general availability (GA) of the x455. See: http://www.pc.ibm.com/us/compat/nos/matrix.shtml 3. Support approximately 30 days after GA.			

Note: Hyper-Threading is not available on systems with the Itanium 2 processor.

3.6.1 Clustering

A cluster is two or more computers (known as nodes) that work together to provide a solution. The idea behind clusters is to join the computing powers of the nodes involved to provide higher scalability, more combined computing power, or to build in redundancy to provide higher availability. The sharing of data between the nodes of a cluster through a common file system is almost always a requirement.

For additional information, refer to the following Web site:

<http://www.pc.ibm.com/ww/eserver/xseries/clustering/info.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/search.aspx>

For information about Windows Server 2003 Enterprise and Datacenter, and changes in the 64-bit version, see:

<http://www.microsoft.com/windowsserver2003/64bit/default.aspx>

3.7 IBM Director support

IBM Director is a systems management application and can be adopted easily into heterogeneous environments. It is based on industry standards and designed for ease of use. The Java™-based GUI helps administrators manage single or large groups of IBM and non-IBM servers, workstations, desktop computers, and notebook computers.

Note: The 64-bit versions of IBM Director Agent and Scalable Systems Manager, for both Windows and Linux, are scheduled to be available in the first quarter of 2004.

The 64-bit IBM Director Agent has the following functions:

- ▶ Event Action Plans
- ▶ Event Log
- ▶ File Transfer
- ▶ Inventory (all except CIM-based information)
- ▶ Management Processor Assistant (requires ASM device driver)
- ▶ Process Management
- ▶ Remote Control
- ▶ Remote Session
- ▶ Resource Monitors
- ▶ ServeRAID Manager
- ▶ Software Distribution (Basic and Premium)
- ▶ SNMP Browser
- ▶ Agent for Windows and Linux

The Scalable Systems Manager is a plug-in of IBM Director that communicates solely out-of-band with the service processor, to create and manage static hardware partitions in the x455. These static partitions are referred to as scalable partitions throughout the Scalable Systems Manager interface. Because Scalable Systems Manager communicates out-of-band with the x455, the primary task you can perform with Scalable Systems Manager is to configure a scalable partition for a specific offline server before starting the operating system for that server.

Scalable Systems Manager is going to be the only way to have two or four operating system instances in a multi-node configuration. For two-node configuration it is possible to have two different instances and for four-node configuration it is possible to have four different instances.

For more information about Scalable Systems Manager, see “Scalable Systems Manager” on page 177.

3.8 Solution Assurance Review

Some level of Solution Assurance Review (SAR) should be performed for 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, a combination of the customer environment risk combined with the complexity of the solution may require that an expert level SAR take place, facilitated by a Quality Assurance practitioner and supported by a team of technical experts.

If a solution contains four or more Enterprise X-Architecture servers (currently x360, x365, x445 and x455), then an expert SAR is mandatory.

For further information on what is required, refer to the Solution Assurance Web sites. EMEA and Americas information is available. Procedures for Asia Pacific countries are currently in development.

3.8.1 Trigger Tool

The SAR Trigger Tool provides a recommendation on the level of Quality and Solution Assurance that will be required. It is available from one of the Solution Assurance Web sites.

The three levels are:

- ▶ Expert
 - For technically challenging, high-risk solutions
 - Process dictates expert personnel's participation
 - Formal, rigorous
- ▶ Peer
 - For low-to-medium risk solutions
 - Informal, inexpensive
- ▶ Self
 - For low-risk solutions
 - Informal, inexpensive

3.8.2 Electronic Solution Assurance Review (eSAR)

There is also an eSAR tool available to help you establish whether you require an expert review. This is available from:

- ▶ For IBM employees: <http://w3.ibm.com/support/assure>, then click **eSAR**
- ▶ For Business Partners: <http://www.ibm.com/support/assure/esar>



Installation

In this chapter we describe both the basic Extensible Firmware Interface (EFI) operations that administrators should know prior to installing an operating system and the operating system installation procedures that are specific to the x455.

The following topics are discussed:

- ▶ “Using The Extensible Firmware Interface” on page 92
- ▶ “Configuring scalable partitions” on page 131
- ▶ “Installing Windows Server 2003” on page 136
- ▶ “Installing Linux” on page 151

A useful document during installation is the paper *IBM @server xSeries 455 FAQ Hints and Tips*, available from:

<http://www.ibm.com/pc/support/site.wss/MIGR-53820.html>

4.1 Using The Extensible Firmware Interface

Intel has introduced the Extensible Firmware Interface (EFI) as a means of giving greater control to the operating system and to the end user as to how the system starts and works. In this chapter we familiarize you with this environment.

The concept of the EFI is described in “Extensible Firmware Interface” on page 29.

Table 4-1 shows the common tasks you would perform in either BIOS or via a DOS diskette and how they are now performed on the x455.

Table 4-1 New ways to do familiar tasks in the EFI

If you want to	On legacy BIOS systems	On EFI systems
Configure the internal SCSI/RAID controller.	Enter adapter's configuration utility during server boot	Select the Driver Setup option from the EFI Firmware Boot Manager menu. See “EFI tips and tricks” on page 102.
Configure and set up the server.	Press F1 during server boot and enter Configuration/Setup utility	Run the Configuration and Setup utility from the EFI Firmware Boot Manager menu. See “Configuration/Setup utility” on page 118.
Configure or flash the Remote Supervisor Adapter.	Boot from Remote Supervisor Adapter CD-ROM or diskette and run the configuration utility	Run the Remote Supervisor utility from the EFI shell. See “Remote Supervisor Adapter firmware update” on page 116.
Run system diagnostics.	Press F2 during server boot and enter the Diagnostic utility	Run the Diagnostic utility from the EFI Firmware Boot Manager menu. See “Error logging” on page 120.
Change the boot sequence.	Press F1 during server boot and make the appropriate changes in Configuration/Setup utility	Enter the Boot Option Maintenance menu from the EFI Firmware Boot Manager. See “Boot Option Maintenance” on page 122.
Flash the firmware, BIOS or diagnostics utilities.	Boot from a Flash diskette with appropriate updates and run flash utility	Run the Flash Update from the EFI Firmware Boot Manager menu. See “System EFI/SAL and Diagnostics” on page 113.
Make changes in disk partitioning before the operating system is installed.	Boot from a DOS diskette and run FDISK.EXE	Run DISKPART from the EFI shell. See “Checking a disk” on page 100.
Basic file operations on the FAT/FAT32 partitions	Boot from DOS diskette and use cd, dir, copy, move, del commands	Enter the EFI shell from EFI Firmware Boot Manager menu and use cd, ls, cp, mv, rm commands. See “File operations in the EFI shell” on page 97.

4.1.1 EFI Firmware Boot Manager

When you boot the x455, it will go through the POST and a number of checks. During this time the various components such as the video port are not fully operational so the information displayed is in basic text mode. Figure 4-1 shows a typical boot up screen with the following information:

- ▶ Build date and time of the EFI/SAL (equivalent to BIOS on IA-32 systems)
- ▶ Partition number
- ▶ Node number within the partition
- ▶ Number of processors
- ▶ Amount of memory
- ▶ Stage the checkpoint has reached
- ▶ Node CPU indicates which node and CPU generated the associated checkpoint entry just above it on the screen

```
IBM                      Eserver                      x455
Build..0YKT12A 2003/10/18 08:52:22      (c) Copyright, IBM Corporation, 2003
Partition.....1 Processors.....[.4|.|.|.|.|.]      Checkpoint.....42-53
Node.....1 of 1 Memory(GB).....[.4|.|.|.|.|.]      Node:Cpu.....1::1
=====
```

Figure 4-1 x455 boot up screen

In a multinode configuration this screen will reflect the information pertaining to the partition and partition node members.

After the checkpoints the hardware components are listed; while this proceeds a number of screens of information will appear. After this you will automatically enter the EFI Firmware Boot Manager menu (as shown in Figure 4-2 on page 94). The menu is the starting point from which you will initiate any actions.

After installing an operating system, the option to start the operating system becomes the first one. By default, when booting the server, the first option is loaded if no key is pressed. This is normally 30 seconds. The menu order and timeout values can be altered as described in “Boot Option Maintenance” on page 122.

```

IBM                               Firmware Boot Manager                               x455
                                Version 2.00 [2.10] (c) copyright IBM Corporation 2003

Boot options: Enter number or ↑/↓ to select, 'Enter' to execute

 1 EFI Shell [Built-in]
 2 Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)
 3 Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)/CDROM (Entry0)
 4 Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)/CDROM (Entry1)
 5 Acpi (PNP0A03,1)/Pci (4|0)/Mac (0002551F809E)
 6 Acpi (PNP0A03,1)/Pci (4|1)/Mac (0002559F809E)

System options: Enter letter ↑/↓ or to select, 'Enter' to execute

  Boot Option Maintenance      Diagnostics      EFI Shell
  Configuration/Setup         Driver Setup     Flash

Messages:

```

Figure 4-2 EFI Firmware Boot Manager menu

In Figure 4-2 the first entry is EFI Shell, a command prompt for the EFI environment. This is discussed in “The EFI shell” on page 95.

In addition to the EFI Shell, there are five other options. The first three represent the CD-ROM, and the other two PXE boot options represent the Broadcom Ethernet controllers. If you choose any of these five options, the system will try to boot from the selected device.

If the boot fails (for example, when trying to boot from CD-ROM without bootable media inserted), the following message appears:

```

Loading.: Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)
Load of Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master) failed: Not Found
Paused - press any key to continue

```

If a bootable CD is in the drive then always choose the first option; the other two will also fail with the message as described above.

When you select one of the network controllers (marked PXE boot), you initiate the PXE remote boot process. If no PXE boot server is discovered on the

network during this process, again a failure message appears. If you need to make the PXE boot the first option, then see “Changing a boot option” on page 126.

The options available from the EFI Firmware Boot Manager are listed below and discussed in further detail.

- ▶ The EFI shell
- ▶ Driver Setup
- ▶ Flash Update
- ▶ Configuration/Setup utility
- ▶ Error logging
- ▶ Boot Option Maintenance

4.1.2 The EFI shell

The EFI shell allows the loading of EFI applications (including booting installed operating systems) from any EFI-defined file system. The EFI environment replaces DOS to perform maintenance operations, such as upgrading firmware or running system diagnostics.

To start the EFI shell, select **EFI Shell [Built-in]** from the EFI Firmware Boot Manager menu. As the shell starts, a device mapping table is displayed, as shown in Figure 4-3 on page 96.

```

Loading.: EFI Shell [Built-in]
EFI Shell version 1.10 [14.61]
Device Mapping Table
fs0  : MemMap(0:FF000000-FFFFFFFF)
fs1  : MemMap(0:FF800200-FFBFFFFFFF)
fs2  : Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)/CDROM (Entry0)
fs3  : Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)/CDROM (Entry1)
fs4  : Acpi (PNP0A03,0)/Pci (5|3)/Usb (0, 0)/HD (Part1,Sig317AD54E)
fs5  : Acpi (PNP0A03,1)/Pci (3|0)/Scsi (Pun0,Lun0)/HD (Part1,Sig3D73CFC0-0000-0000
-507B-9E5F8078F531)
blk0 : MemMap(0:FF000000-FFFFFFFF)
blk1 : MemMap(0:FF800200-FFBFFFFFFF)
blk2 : Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)
blk3 : Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)/CDROM (Entry0)
blk4 : Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)/CDROM (Entry1)
blk5 : Acpi (PNP0A03,0)/Pci (5|3)/Usb (0, 0)
blk6 : Acpi (PNP0A03,0)/Pci (5|3)/Usb (0, 0)/HD (Part1,Sig317AD54E)
blk7 : Acpi (PNP0A03,1)/Pci (3|0)/Scsi (Pun0,Lun0)
blk8 : Acpi (PNP0A03,1)/Pci (3|0)/Scsi (Pun0,Lun0)/HD (Part1,Sig3D73CFC0-0000-0000
-507B-9E5F8078F531)
blk9 : Acpi (PNP0A03,1)/Pci (3|0)/Scsi (Pun0,Lun0)/HD (Part2,Sig3DA00FE0-0000-0000
-F1B3-12714F758821)
blkA : Acpi (PNP0A03,1)/Pci (3|0)/Scsi (Pun0,Lun0)/HD (Part3,Sig2613D040-0001-0000
-D931-F8428177D974)
Shell>

```

Figure 4-3 Entering EFI shell

This mapping table lists both block devices (blk) and file systems (fs). All fixed disks, CD-ROM drives, and USB storage devices will appear as block devices. If the server contains additional controllers and disks, they will appear as additional block devices. The EFI automatically maps known file systems on the block devices to their FSx mappings (FSx is referred to as the *file system identifier*).

Here are some tips for easier identification of devices in the list:

- ▶ Even with no storage devices, the server has at least two default mappings, the MemMap entries listed as fs0: and fs1: in Figure 4-3. These are read-only RAM disks created during POST, so they will always be available. The first MemMap file system is the content of the 4 MB of NVRAM assigned to the EFI and contains a number of drivers. The second MemMap file system is the content of the NVRAM assigned to the server diagnostics.

- ▶ The internal disks contain Scsi and HD strings. See Figure 4-3 on page 96:
 - Scsi (Pun0,Lun0)/HD(Part1,Sig3D73CFC0-0000-0000-507B-9E5F8078F531)
 - Pun is the physical unit number. For non-RAID disks this represents the SCSI ID. For RAID disks this represents the array number as enumerated by the EFI.
 - Lun is the logical unit number
 - Part is the partition number. This starts at 1.
 - Sig is the disk's GUID (discussed in "GUID Partition Table disk" on page 31).
- ▶ The disk partitions are numbered from 1 onwards. The first partition is Part1, the second Part2, and so on.
- ▶ The EFI System Partition is easily identifiable in Figure 4-3 on page 96 if we realize that EFI is always the first partition on the first internal disk (that is, it is identified by the entry with strings SCSI, Pun0, Lun0 and HD Part1). In this case it is fs5. This may be a single disk or RAID-1 array.
- ▶ In Figure 4-3 on page 96, the CD-ROM is represented as fs2 and fs3. The reason for the two entries is because EFI separates the "EFI bootable" part of the CD (Entry1) and the data part (Entry0). The two entries show that this CD is bootable. If the CD was non-bootable, only one entry would exist (Entry 0).

Note: Unless CD-ROM media is inserted, there is no FSx mapping created for the CD-ROM.

There can also be multiple partitions on a CD-ROM, so it is possible to see several entries for a single CD-ROM.

- ▶ The USB memory key devices include the string USB and is represented as fs4. If the USB block device is installed before the machine is booted it will have a lower fs number than the hard disk with the ESP. If it is removed and reinserted or inserted after the machine is booted it will have a higher fs number. This is important for you to remember, as some EFI options will search or save automatically to the lowest suitable fs number.

File operations in the EFI shell

From the EFI shell you may access FAT/FAT32 file systems on existing disk partitions, run EFI executables, and manipulate the disk contents on various media. The EFI executables are text mode utilities similar to DOS or UNIX commands. In general, EFI commands are not case-sensitive.

Tip: All the EFI executable files have an extension of .EFI. The executables may be run only from within EFI, not from within any operating system.

Access to the file system on a device is done through the **FSn:** command (where **n** is the file system number shown in the device mapping table). This is similar in concept to changing drive letters in DOS. You can use the **map** command at any time to remind yourself of the available file systems.

Tip: Adding or removing media may cause arbitrary reassignments of the FSx designations. Always make sure that you access the right media by issuing the **map** command or **map -r** if media has been added or removed.

As an example of the capability of accessing file systems, the following are the steps to access data stored on the USB memory key device:

1. Insert the USB memory key into the USB port (can be hot-added).
2. Start the EFI shell from the EFI Firmware Boot Manager menu.
3. Use the **map** command to identify the appropriate mapping for the USB device the EFI automatically created. Sample output of the command is in Figure 4-4.

If the appropriate mapping does not appear, run **map -r**, where parameter **-r** forces EFI to refresh the device mapping table.

```
Shell> map
Device Mapping Table
fs0 : MemMap(0:FF000000-FFFFFFFF)
fs1 : MemMap(0:FF800200-FFBFFFFF)
fs2 : Acpi(PNPOA03,0)/Pci(5|2)/Usb(0, 0)/HD(Part1,Sig00000000)
blk0 : MemMap(0:FF000000-FFFFFFFF)
blk1 : MemMap(0:FF800200-FFBFFFFF)
blk2 : Acpi(PNPOA03,0)/Pci(5|1)/Ata(Primary,Master)
blk3 : Acpi(PNPOA03,0)/Pci(5|1)/Ata(Primary,Slave)
blk4 : Acpi(PNPOA03,1)/Pci(3|0)/Scsi(Pun0,Lun0)
blk5 : Acpi(PNPOA03,0)/Pci(5|2)/Usb(0, 0)
blk6 : Acpi(PNPOA03,0)/Pci(5|2)/Usb(0, 0)/HD(Part1,Sig00000000)
```

Figure 4-4 Map command

4. Look for the specific device in the FSx section. In our example, the entry is highlighted in Figure 4-4 on page 98. The entry will include the string USB.
If the USB device is attached to a USB hub (for example, a memory key plugged into the USB port of a USB keyboard), then two USB strings will appear in the entry, as shown in Figure 4-5.

```
Shell> map
Device Mapping Table
fs0 : MemMap(0:FF000000-FFFFFFFF)
fs1 : MemMap(0:FF800200-FFBFFFFFFF)
fs2 : Acpi(PNP0A03,0)/Pci(5|1)/Usb(0,0)/Usb(3,0)/HD(Part1,Sig00000000)
blk0 : MemMap(0:FF000000-FFFFFFFF)
blk1: MemMap(0:FF800200-FFBFFFFFFF)
blk2: Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)
blk3: Acpi(PNP0A03,1)/Pci(3|0)/Scsi(Pun1,Lun0)
blk4: Acpi(PNP0A03,0)/Pci(5|2)/Usb(0,0)/Usb(3,0)
blk5: Acpi(PNP0A03,0)/Pci(5|2)/Usb(0,0)/Usb(3,0)/HD(Part1,Sig00000000)
```

Figure 4-5 Map command - USB memory key connected to the USB keyboard

5. Access the device by its name followed by the colon; for example, fs2:.
6. Use the commands **cd**, **ls** (or **dir**), **cp** (or **copy**), **mv**, **rm** to manipulate the files on the device in the same way you would access a device through the operating system.

The commands have a similar syntax to DOS or UNIX. However, where there are duplications only one is listed in Help. For example, only **ls** is listed in Help, but **dir** also works. If **dir /?** is typed the syntax for **ls** is shown.

CD-ROM media supported in the x455 can either be ISO9660 (El-Torito) or Joliet format. This differs from the x455 that only supported ISO9660. Writing to the DVD/CD-RW drive is not supported from the EFI shell.

If Joliet format is used then we recommend that you use 8.3 DOS format for file names, as the EFI will truncate long file names that may then be difficult to decipher. If ISO9660 (El-Torito) is used then EFI displays the full file name.

There are a variety of built-in shell commands. Of note are:

- ▶ **help -b** displays a list of the available commands. The -b parameter pauses the display after each screen full of information and can be used by many EFI commands.
- ▶ To view the contents of a text file, use the **viewer** command.
- ▶ To edit any text file, use the **edit** command.

- ▶ To restart the server (no confirmation), use **reset**. To power off the server, use **reset -s**.

Checking a disk

The third-party utility EFICLK.EFI is useful for checking some basic statistics of a disk. The output will be similar to Figure 4-6.

```
fs5:\efi\tools> efichk fs0
EFI Check Disk Version 1.0
Based on EFI Core Version 1.2.12.38
Checking file system on fs0
The type of the file system is FAT.
Volume Serial Number is 0000BC91-0000512C
EFI is verifying files and folders...
File and folder verification is complete.
EFI has checked the file system and found no problem.
    180723712 bytes total disk space.
    745472 bytes in 1 hidden files.
    20480 bytes in 5 folders.
    495616 bytes in 3 files.
    179462144 bytes available on disk.
    4096 bytes in each allocation unit.
    44122 total allocation units on disk.
    43814 allocation units available on disk.
```

Figure 4-6 EFICLK.EFI output

Deleting the disk content

There are situations when it is necessary to delete the contents of the disk. For example, you want to install an operating system on a disk that was previously used in a non-EFI system (that is, with a disk with a standard Master Boot Record (MBR) or you want to delete a partition that is reserved and cannot be deleted by other means (for example, a Microsoft Reserved Partition—see “Microsoft Reserved Partition” on page 137).

Tip: If you need to perform a low-level format of the disk, use the LSI Configuration utility, which is described in “EFI tips and tricks” on page 102.

To delete the content of the disk, perform the following steps:

1. Download the DISKPART.EFI drive partitioning utility and transfer it to the x455 using a device such as a USB memory key or CD-R. This is available from:

http://developer.intel.com/technology/efi/diskutil_overview.htm

2. Invoke the EFI shell from the EFI Firmware Boot Manager menu. Change to the device where the DiskPart utility is (for example, CD-ROM or USB memory key) and run it. A DiskPart prompt appears, similar to the following:

```
DiskPart Version 1.0
Based on EFI core release Version 1.2.1.0
DiskPart>
```

3. Issue the following commands:

list List all the disks in the system that can be deleted.

select dd Select the disk you will delete.

clean Erase the contents of the selected disk. Confirm the action as prompted.

exit Exit the console and return to the EFI shell.

The overall process is shown in Figure 4-7.

```
fs5:\tools> diskpart
DiskPart Version 1.0
Based on EFI core release Version 1.2.1.0
DiskPart> list
   ##  BlkSize          BlkCount
   ---  -
   0    200             21E6CA0
   1    200             21E6CA0
   2    200             32D9000
DiskPart> select 0
Selected Disk = 0
DiskPart> clean
About to CLEAN (DESTROY) disk 0, are you SURE [y/n]?
CLEAN>> y
If you are REALLY SURE, type '$C'
CLEAN>> $C
DiskPart> exit
Exiting....
```

Figure 4-7 DiskPart commands to erase a disk

The described procedure can be used to delete both the MBR and GPT disks.

Tip: The hard drive partitioning utility DISKPART.EFI may also be used to create the disk partitions, but to do this we recommend using tools available during or after operating system installation.

The numbers for block size (BlkSize) and block count (BlkCount) are displayed in hex. To calculate the overall size open Windows Calculator and change to Hex mode by clicking **View -> Scientific -> Hex**. Multiply both numbers and then click **View -> Scientific -> Dec**. This should give the disk size. Using the examples in Figure 4-7 on page 101, this would be:

```
200 x 21E6CA0 = 43CD94000 (hex)
43CD94000 (hex) = 18200739840 (dec)
18200739840 = 18.2 GB disk
```

```
200 x 32D9000 = 65B200000 (hex)
65B200000 (hex) = 27298627584 (dec)
27298627584 = 27.3 GB disk (this is a RAID array)
```

The actual block count is typically not an exact figure and is based on how EFI views the available space it has access to. However, the number will be accurate enough to determine the physical disk size.

Tip: It is possible to use the LSI (see “EFI tips and tricks” on page 102) or ServeRAID utilities (see “ServeRAID configuration CD” on page 110) to delete partitions by destroying and re-creating the RAID array.

EFI tips and tricks

By default, the EFI has a basic set of tools to get you up and running. However, it is possible to enhance this by installing additional tools and making some basic modifications. The tips below guide you through some of these.

- ▶ A USB memory key is highly recommended for efficient use and maintenance of the EFI environment.
- ▶ When a USB memory key is inserted there is a delay of up to 10 seconds during which the keyboard freezes. The keystrokes are buffered.
- ▶ After an EFI/SAL or diagnostics upgrade, move the update files from the root of the drive from where they were installed. This avoid future problems when upgrading.

- ▶ Boot CDs will have at least two file system entries, one ending with (Entry0) and one ending with (Entry1). The file system with (Entry1) will have the boot loader file(s). It is normally in the following location:

`\efi\boot\bootia64.efi`

This file or a batch file that calls this file may also be in the root directory.

- ▶ Some commands, although not listed, are available. These can be used as well as the official commands. Invoking the help option for these commands shows the help for the official commands.

Table 4-2 Undocumented commands

Undocumented	Official
<code>dir</code>	<code>ls</code>
<code>copy</code>	<code>cp</code>

- ▶ In a multinode partition the EFI will enumerate all keyboards and mice. The net effect of this is that any attached keyboard or mouse can be used.
- ▶ There are a number of EFI executables that are worth copying to the system partition so that they are available for maintenance and emergency purposes without having to resort to a USB drive or CD. These are listed in Table 4-3.

Table 4-3 Useful EFI tools

Tool	Description
DISKPART.EFI	Hard drive partitioning tool
EFICLK.EFI	Hard drive file system checker
EFIFMT.EFI	Hard drive partitioning formater
FLASH2.EFI	RSA configurator (and updater)
RAMDISK.EFI	RAM disk creation driver

4.1.3 Driver Setup

Driver Setup (on the main menu, Figure 4-2 on page 94) allows you to configure different devices in the machine. For example, you can configure the LSI controller to set up a RAID-1 array.

By default there will be five devices. See Figure 4-8 on page 104 for an example.

Driver Configuration Menu (c) copyright IBM Corporation 2003				
Number	Driver Handle	Ctrl Handle	Driver Name	
1	30	45	PCI IDE/ATAPI Bus Driver	
2	3D	59	Broadcom Gigabit Ethernet Driver	
3	3D	5A	Broadcom Gigabit Ethernet Driver	
4	55	57	LSI Logic Ultra320 SCSI Driver	
5	55	58	LSI Logic Ultra320 SCSI Driver	
6	5D	5F	Adaptec RAID Driver	

Options : ↑ or ↓ to select
 <ENTER> to execute
 <ESC> to return to the Boot Manager

Figure 4-8 Driver configuration menu

There are repeated entries for the Ethernet and SCSI drivers, as the machine has two Ethernet ports and two SCSI channels.

The selections in this menu are described below.

Configuring PCI IDE/ATAPI Bus Driver

This option allows you to enable or disable IDE devices. The options are simple yes or no answers. See figure Figure 4-9 for an example.

Enable Primary Master	(Y/N)?	-->Y
Enable Primary Slave	(Y/N)?	-->N
Enable Secondary Master	(Y/N)?	-->N
Enable Secondary Slave	(Y/N)?	-->

Figure 4-9 Configuring PCI IDE/ATAPI Bus Driver

When this is complete you will be prompted to press Enter to restart the controller. This restarts the controller and rescans all the devices in the machine.

Configuring Broadcom Gigabit Ethernet Driver

There were no options available as the time of writing; selecting 02 or 03 returns you immediately to the EFI Firmware Boot Manager. This may change in future updates to the EFI or Broadcom firmware.

Configuring SCSI controller and internal disks

This option allows you to configure the SCSI Controller. Typically the only option that would be changed is the RAID configuration. The internal LSI SCSI/RAID controller can treat the internal disks in two different ways:

- ▶ As two standalone disks
- ▶ As two parts of a RAID-1 array

Both ways are supported for operating system installation, although installing on the single non-redundant disk is not recommended.

RAID-1 is the only possible redundant configuration consisting of a single logical drive spanning the two disks. There is no way to create multiple logical volumes or hot spares.

Using the embedded LSI controller allows the content of the primary disk (the disk marked as Primary in the Mirroring Properties menu) to stay intact. The contents are copied to the secondary disk during the first synchronization. This feature means that the installation of the operating system and the creation of the mirror can normally be done in any order (see warning below).

Warning: Even though it is possible to preserve the contents of the primary disk after the mirror creation, this is only possible if the disk is not full. Full means the whole disk is partitioned; it does not have to have any data. If the disk is full then data cannot be saved and the contents of the primary disk will be lost. We strongly recommend that you configure the RAID options before the installation of the operating system.

The following steps describe the process of creating a RAID-1 array.

1. Select the Driver number for the SCSI Controller. The following screen will be displayed (Figure 4-8 on page 104).

```

LSI Logic MPT SCSI Setup Utility Version v1.02.11

LSI Logic Host Bus Adapters
Adapter PCI PCI PCI RAID FW Revision Pro- LSI RAID IRQ
          Bus Dev Fnc Alert          duct Control Status
<53C1030 01 03 00> None 1.03.08.00 IME Enabled Optimal 0

F1/Shift+1 =Help ArrowKeys/H,J,K,L =Select Item -/+ =Change [Item]
Esc=Exit Menu Home(I)/End(0) =Select Item Enter=Execute <Item>

```

Figure 4-10 Options to choose RAID cards

2. Select the host bus adapter. Normally there is only one option to choose. In this example it is labelled 53C1030. The following screen will be displayed.

```

LSI Logic MPT SCSI Setup Utility Version v1.02.11

Adapter Properties

Adapter PCI PCI PCI
          Bus Dev Fnc
53C1030 01 03 00

    <Device Properties>
    <RAID Properties>

    Boot Support [Enabled BIOS & OS]
    Host SCSI ID [ 7]
    Secondary Cluster Server [No]
    Termination Control [Auto]
    <Restore Defaults>

F1/Shift+1 =Help ArrowKeys/H,J,K,L =Select Item -/+ =Change [Item]
Esc=Exit Menu Home(I)/End(0) = Select Item Enter=Execute <Item>

```

Figure 4-11 Option to choose RAID configuration

3. Select **RAID Properties** and press Enter. If you have already configured RAID there will be an additional screen that allows you to Modify, Sync, or Delete the RAID array.

```

LSI Logic MPT SCSI Setup Utility  Version  v1.02.11

RAID Properties      Array: IM  SCSI ID:  0   Size(MB):  17278

SCSI Device Identifier      Array  Hot   Status   Predict Size
ID                           Disk?  Spare  Failure  Failure (MB)
0  IBM-PSG ST318451LC  !#B833 [Yes] [No]  Primary  No      17278
1  IBM-ESXSST318452LC !#B841 [Yes] [No]  Secondary No      17278
2  -                  [---] [---] ----- --      -----
3  -                  [---] [---] ----- --      -----
4  -                  [---] [---] ----- --      -----
5  -                  [---] [---] ----- --      -----
6  -                  [---] [---] ----- --      -----
7  53C1030           [---] [---] ----- --      -----
8  IBM      25P3495a S320 11 [---] [---] Incompatbl-- -----
9  -                  [---] [---] ----- --      -----
10 -                  [---] [---] ----- --      -----
11 -                  [---] [---] ----- --      -----
12 -                  [---] [---] ----- --      -----
13 -                  [---] [---] ----- --      -----
14 -                  [---] [---] ----- --      -----
15 -                  [---] [---] ----- --      -----
F1/Shift+1 =Help ArrowKeys/H,J,K,L =Select Item  -/+ =Change [Item]
Esc=Exit Menu  Home(I)/End(O) =Select Item      Enter=Execute <Item>

```

Figure 4-12 Configuring RAID 1

4. Change the Array Disk options for SCSI IDs 0 and 1 to Yes. These are the two internal hard disks. Normally SCSI ID 0 is the primary disk, but this may not be the case. All other devices will either have nothing in the Status column, or will say Incompatbl. In both cases these devices cannot be used as part of a RAID configuration.
5. Once all the configuration changes are made press Esc. The following message appears:
Warning: Data on Secondary or Hot Spare! Mirroring will overwrite!
6. Confirm by pressing Enter and continue.
7. Save the changes on the next window and exit the menu. The disks will synchronize in the background.

Help is available by pressing F1, and the instructions to change configuration are listed at the bottom of the screen.

Notes® on the use of RAID:

- ▶ RAID can be configured on either the internal or external SCSI port but not both at the same time. If one is configured and there are disks on the other port, they are only available in a non-redundant configuration. If the external port is used, additional RAID options are available, but will require ServeRAID to configure them. We recommend using the LSI RAID Controller on the 2 internal disk and configure a RAID-1 array. If RAID is required on an external disk then install a ServeRAID card and configure these as required.
- ▶ After confirming the mirror creation, the data on the secondary disk is lost. The secondary disk will be overwritten by the content of the primary disk whose content stays intact.
- ▶ The initial mirror synchronization can take up to several hours. The process itself runs in the background. You can begin the operating system installation at any time, even while the synchronization is running.
- ▶ During the synchronization, the server may be restarted or turned off as needed. If not finished, the procedure will continue after restarting without user intervention.
- ▶ Booting from ServeRAID-attached drives is supported, using one of the supported adapters. See “ServeRAID with external storage enclosures” on page 79 for a list of supported devices.

ServeRAID mini-configuration utility

If a ServeRAID adapter is installed (4Mx or 6M), it will dynamically load a driver within the EFI environment. This can then be selected from the list of drivers. This loads the ServeRAID mini-configuration utility that allows you to view or amend some of the card's configuration parameters. The mini-configuration utility is not as comprehensive (or intuitive) as ServeRAID Manager and does not let you configure a RAID-1 array from scratch. ServeRAID is the recommended program for initial setup and configuration.

The use of the mini-configuration utility is as follows.

From the Driver Configuration menu (Figure 4-8 on page 104) select **Adaptec RAID Driver**. The following screen is displayed.

```

IBM ServeRAID Mini Configuration - Main Menu
Controller 1 : ServeRAID-4MX
EFI Ver:

View Controller Status
View Logical Drive Status
Advanced Functions

Highlight Function using Up/Down arrow keys, to select press Enter.
ESC for Previous Menu

```

Figure 4-13 IBM ServeRAID Mini-Configuration Main Menu

From here you can view the controller status, view the logical drive status, or select more advanced functions.

Figure 4-14 shows an example of controller status.

```

IBM ServeRAID Mini Configuration - Controller Status
Controller 1 : ServeRAID-4MX
EFI Ver:

Registers:ISPR=EF10h BCS= 9h ECS= 0h
Unattend:off ReadAhead:ADP
CompMode:off Clustered:off BBC:64MB
Boot Blk:6.10.32 Code Blk:6.10.32
Rebuild rate:High
Number of Defunct physical drives:0
Number of Offline logical drives:0
Number of Critical logical drives:0

Bad stripe log dvs:None
Locked stripe log drvs:None
Blocked log drvs:None

Ch 1 Ch 2
Array Array
SID SID
0 ONL A 0 ---
1 ONL A 1 ---
2 --- 2 ---
3 --- 3 ---
4 --- 4 ---
5 --- 5 ---
6 --- 6 ---
7 INI 7 INI
8 SBY 8 ---
9 --- 9 ---
10 --- 10 ---
11 --- 11 ---
12 --- 12 ---
13 --- 13 ---
14 --- 14 ---
15 --- 15 ---

Press F1 for Help ESC = Prev Menu

```

Figure 4-14 Controller status

Figure 4-15 on page 110 shows an example of logical drive status.

IBM ServeRAID Mini Configuration - Logical Drive Status								EFI Ver:
Controller 1 : ServeRAID-4MX								
Logical Drive	Size (MB)	RAID Level	Logical State	R5E/R5EE State	Stripe Size	Write Policy	Read Ahead	Date Created mm/dd/yyyy
1	17357	1	okay		8kb	Back	ADP	01/01/1970

ESC= Prev Menu

Figure 4-15 Logical drive status

Figure 4-16 shows the Advanced menu and options available.

IBM ServeRAID Mini Configuration - Advanced Functions Menu		EFI Ver:
Controller 1 : ServeRAID-4MX		
<div style="border: 1px solid black; padding: 5px; margin: 10px auto; width: fit-content;"> Restore to Factory Default Settings Copy the configuration from drives to the controller View controller and PCI information </div>		
Highlight Function using Up/Down arrow keys, to select press Enter. ESC for Previous Menu		

Figure 4-16 Advanced functions

ServeRAID configuration CD

ServeRAID allows you to configure your RAID cards either from a bootable CD or from the operating system. For the x455 it provides the same functions as the LSI Driver Setup utility but presents a graphical front end. ServeRAID has the following advantages:

- ▶ Allows the configurations of all LSI and ServeRAID controllers
- ▶ Will allow you to update the firmware version on ServeRAID controllers (but not on the internal LSI controller)

- ▶ Contains drivers for your operating system
- ▶ Contains the ServeRAID Manager that can be loaded within your operating system
- ▶ Allows data to remain intact on the Primary disk

A ServeRAID CD is shipped with every ServeRAID controller, or the latest version can be downloaded from:

<http://www.ibm.com/pc/support/site.wss/MIGR-495PES.html>

Please refer to the documentation for further information on using ServeRAID.

Booting from ServeRAID

It is possible to boot from a ServeRAID adapter. To do this consult the server's Installation Guide, as this involves the removal and installation of a number of parts. Below are the basic steps.

1. Power off the machine.
2. An internal short SCSI cable will be attached from the back plane of the internal drives to a connector on the planar. An internal long SCSI cable is shipped with the machine. It will be disconnected at both ends.
3. Disconnect both ends of the short SCSI cable and store it in a safe place.
4. Connect one end of the long SCSI cable to the back plane of the internal drives.
5. Connect the other end of the long SCSI cable to the ServeRAID adapter.
6. Ensure that all cable are neat and tidy and power back on the machine.
7. When the machine boots, configure the RAID arrays and reboot.
8. The disks will now be available for the load of the operating system.

Note: Until the arrays are configured and online EFI and EFI tools will not recognize the presence of the disk(s). Only the ServeRAID 4Mx and 6M are supported. Earlier ServeRAID cards are not supported.

Notes on the use of a ServeRAID adapter

Some things to note are:

- ▶ When the ServeRAID CD is in the drive it may not load by choosing the CD option from the EFI Firmware Boot Manager. If it does not, then select **EFI Shell [Built-in]** and this should automatically load the CD. It is not clear why this happens. This behavior only affects the ServeRAID 6.10 CD.

- ▶ When a ServeRAID boot CD is used on a machine with an LSI chipset a warning message is initially displayed implying that the card may be bad. This can be ignored; it is simply warning you that it cannot update the LSI firmware.
- ▶ The lowest version of ServeRAID controller firmware that can reliably be used in the x455 (also x455 and x382) is V6.10. At the time of writing there was an issue with flashing lower versions in these machines. The recommended way to upgrade the version of firmware is to put the controller in an IA-32 machine, flash it, then replace it in the x455. For more information, see RETAIN® tip H18505:
<http://www.ibm.com/pc/support/site.wss/MIGR-53247.html>
- ▶ 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 RETAIN tip H177072:
<http://www.ibm.com/pc/support/site.wss/MIGR-46174.html>

Booting from Fibre Channel

At the time of writing, booting from a Fibre Channel-attached disk was not supported.

4.1.4 Flash update

As part of your installation procedure, we recommend that you check the firmware levels and update to the most current revision. Typical components that may need upgrading are:

- ▶ System EFI/SAL (equivalent to BIOS on IA-32 systems)
- ▶ Onboard diagnostics
- ▶ Remote Supervisor Adapter
- ▶ Additional devices if installed, such as ServeRAID adapters and FAStT host bus adapters

We recommend that you update both the EFI/SAL and Diagnostics at the same time if new versions are available for both. Always check the readme.txt file (the actual name is usually the same as the update executable) to check if there are any prerequisites.

The latest updates and firmware code can be found at:

<http://www.pc.ibm.com/qtechinfo/MIGR-53575.html>

Unlike an IA-32 platform, the normal boot process does not permit booting from a diskette, so a flash update procedure must start differently.

Warning: If a problem is encountered during a firmware upgrade, such as a total loss of power, this can corrupt system components, resulting in loss of data. Always run a backup and verify its integrity prior to any upgrades.

System EFI/SAL and Diagnostics

Once the program loads, it looks exactly like the diskette version but with an important difference. The diskette version searches for the update files on the diskette, whereas the EFI flash program reads the root directories of all accessible file systems until it finds the firmware update files. The implication is that once the utility successfully locates the first flash update file, it will stop searching. This may not be the newest flash file. The routine does not search subdirectories.

You should realize this behavior when you have, for example, the latest firmware updates on a USB memory key (for example, fs5:) and an older version stored in the root of the EFI System Partition (for example, fs4:). In such situations, the flash update utility will keep offering only the file from the first file system it finds that is the EFI System Partition. This is not the one you want. To resolve the issue, you have to delete the firmware update files from the system partition and repeat the procedure. If you want to delete the files from ESP, enter the EFI shell, identify the EFI System Partition, and use the `rm` command (for the exact syntax, use the `help` command).

How to recognize update files used with Flash Update: The firmware EFI/SAL update files are named MZKT*.FLS, for example, MZKT19AUS.FLS.

Diagnostics update files are named MZYT*.US1, where * is a two-digit number indicating the level of the ROM Diagnostics code followed by optional additional characters; for example, MZYT05A.US1.

To see the current version of system components, from the EFI Firmware Boot Manager complete the following steps: Select **Configuration/Setup -> System Information -> System Product Data**. The versions for EFI/SAL (still called BIOS), Service Processor (SP), and Diagnostics are displayed as shown in Figure 4-17 on page 114.

```
IBM xSeries x455 System (c) Copyright IBM Corporation 2002

System Product Data
-----
Machine Type/Model 88552RZ
System Serial      23a0014
System UUID        80659BCDDFF7D711AEBF0000080FFFF
System Asset Tag

BIOS Version       1.01
BIOS Date          09-25-03
BIOS Build Level   OYKT10AUS

SP Version         1.17
SP Date           09-26-03
SP Build Level     MZET28A

Diagnostic Version  1.00
Diagnostic Date     09-17-03
Diagnostic Build Level MZYT15AUS

-----
<F1> Help
<Esc> Quit
```

Figure 4-17 System Product Data window

To update the system firmware and the diagnostics (we always recommend that you do both at once), follow these steps:

1. Select **Flash Update** from the EFI Firmware Boot Manager menu (Figure 4-2 on page 94). The following screen will be displayed.

```
IBM Flash Update Utility
version 4.07

-----
0 - Exit
1 - Update EFI/SAL
2 - Update Diagnostic
3 - Backup Current EFI/SAL
4 - Restore EFI/SAL from the backup version on a drive
   Select an option - 1
```

Figure 4-18 Flash Update Utility

2. Press 1 then Enter to update the EFI/SAL.
3. You are asked if you wish to change the serial number, the machine type, and the asset tag. Typically, you would answer No.

4. You will then be asked if you would like to save the current code to a disk. See the screen below.




Figure 4-19 Prompting to backup

If you select **Y**, you need to specify a file name for the backup. The file name can be any valid file name. If you plan to store the files on CD then we recommend using the 8.3 DOS format. The Flash Update utility saves the backup file to the writable file system with the lowest file system identifier number (FSx). In most cases this is EFI System Partition, but may also be a USB drive such as the example in Figure 4-3 on page 96. If the operation fails, make sure some sort of writable media is attached to the system.

CD-R/W and MemMap mappings are not considered writable devices in EFI.

In our system, as shown in Figure 4-3 on page 96, the backup file would be saved on fs4. Do not try to specify the full path to the file name, because the utility does not accept full path syntax.

5. At this point, the Flash Update utility should detect the presence of the update files on the disk, and you will be asked to choose which language you wish to use during POST and in Setup.
6. Once the flash update completes, you will be prompted to press Enter to restart your system.
7. After rebooting, return to the Flash Update utility.
8. Enter 2 to select **Update Diagnostics** from the menu.
9. Follow the instructions in the window until you successfully update the diagnostics. Press Enter to reboot the server when requested.
10. After rebooting, verify that you have successfully completed the update. From the EFI Firmware Boot Manager menu select **Configuration/Setup -> System Information -> System Product Data** (see Figure 4-17 on page 114).

Note: There is no option to back up the diagnostics.

Remote Supervisor Adapter firmware update

To update the Remote Supervisor Adapter (RSA), use the IBM Remote Supervisor Adapter utility. This allows you to do the following:

- ▶ Update the Remote Supervisor Adapter.
- ▶ Configure the Remote Supervisor Adapter.
- ▶ Update the RXE-100 Expansion Enclosure.

The utility is not a default part of EFI configuration, it may be located on the media shipped with the server but it is always better to check for the latest version that can be downloaded from the x455 driver page:

<http://www.pc.ibm.com/qtechinfo/MIGR-53575.html>

Unlike the Flash Update utility, which requires update files to be in the root directory of an available file system, the Remote Supervisor Adapter update files can reside in any directory on an accessible file system, but the update program (FLASH2.EFI) and data files must reside in the same directory.

The following files are required to be in the same directory to successfully update the Remote Supervisor Adapter and RXE-100:

FLASH2.EFI	Flash update executable
CNETBRUS.PKT	System Management Adapter Boot ROM
CNETMNUS.PKT	System Management Adapter Main Application
CNETRGUS.PKT	Remote Graphics Application
DMCFLASH.PKT	Integrated System Management Processor Application
IIOEFLSH.PKT	RXE-100 Remote Expansion Enclosure Application
IIOEFLPB.PKT	RXE-100 Power Backplane Application

Tip: Attach the RXE-100 before starting the Remote Supervisor Adapter firmware update so it will be updated at the same time.

To apply the Remote Supervisor Adapter (and RXE-100) updates, follow these steps:

1. Enter the EFI shell from the EFI Firmware Boot Manager menu.
2. Navigate to the folder that contains FLASH2.EFI and the update files.
3. At the EFI shell command line, type `flash2` and press Enter. The main window is displayed.

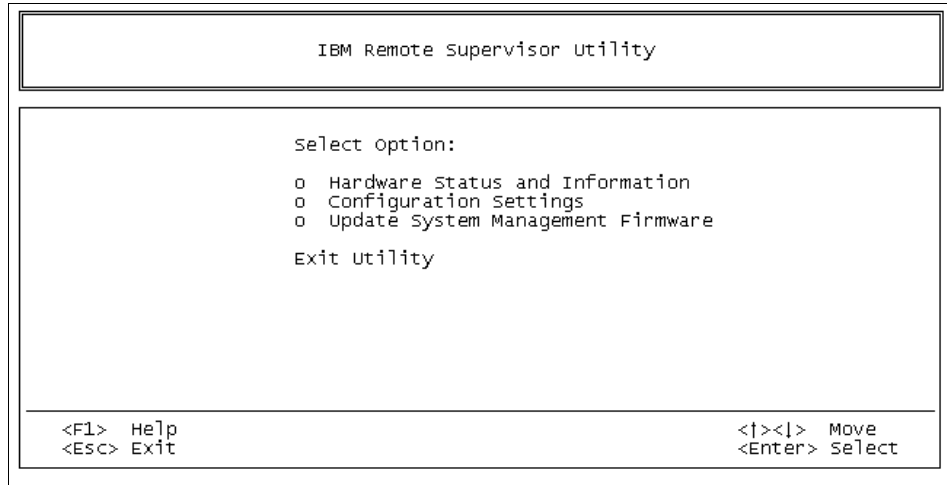


Figure 4-20 Remote Supervisor Adapter update utility

4. Select **Update System Management Firmware** from the main menu and press Enter.
5. Read the warning window and press Enter. You will be presented with the Firmware Update Options menu.

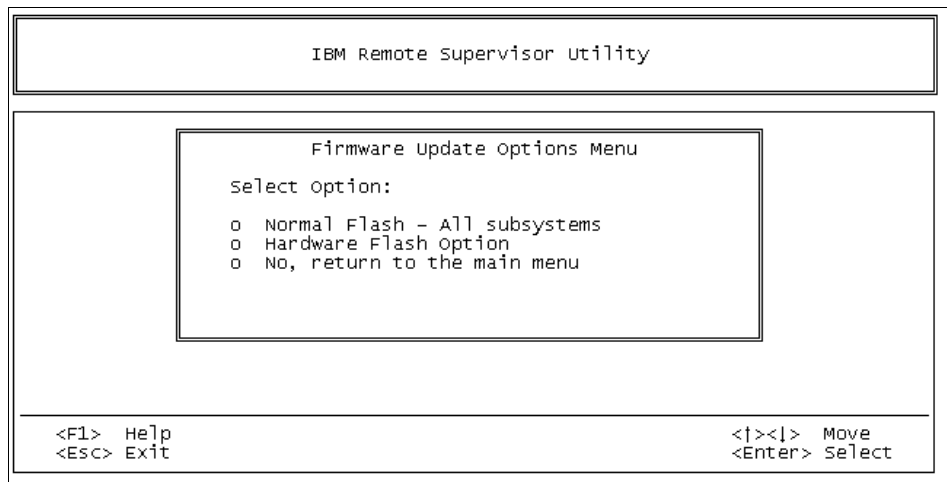


Figure 4-21 Selecting the type of flash

6. Select **Normal Flash - All subsystems** and press Enter. This option will flash all management code in the system. You will see the following components updated:
 - System Management Adapter Boot ROM
 - System Management Adapter Main Application
 - Remote Graphics Application
 - Integrated Systems Management Processor Application
 - RXE-100 Remote Expansion Enclosure Application (if installed)
7. Press Esc to return to the main menu and quit the utility.

4.1.5 Configuration/Setup utility

Select **Configuration and Setup Utility** from the EFI Firmware Boot Manager menu (see Figure 4-2 on page 94) to configure internal options. Figure 4-22 is displayed.

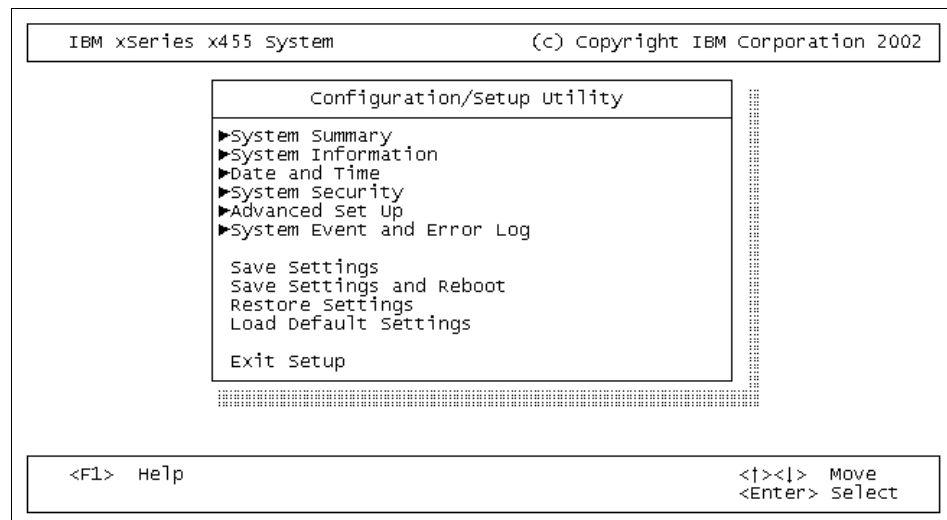


Figure 4-22 Configuration and Setup utility

These options are similar to those available on IA-32 systems. They are well laid out and fairly self explanatory.

Enabling memory mirroring

Memory mirroring is part of IBM's Active Memory technology and provides an additional level of fault tolerance to the memory subsystem. For detailed information and guidelines on memory mirroring, see "System memory" on page 14.

To enable memory mirroring perform the following steps:

1. Allow the system to boot to the EFI Firmware Boot Manager menu.
2. Start the Setup program by selecting the **Configuration/Setup Utility** from the EFI Firmware Boot Manager menu. Figure 4-22 on page 118 appears.
3. Select **Advanced Set Up**. The following screen is displayed

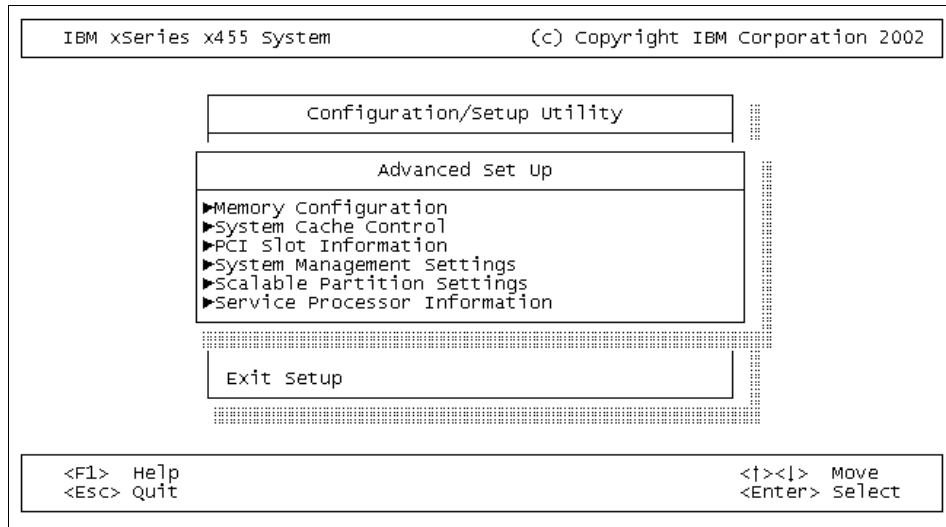


Figure 4-23 Advanced Set Up menu

4. Select **Memory Configuration**. The following screen is displayed.

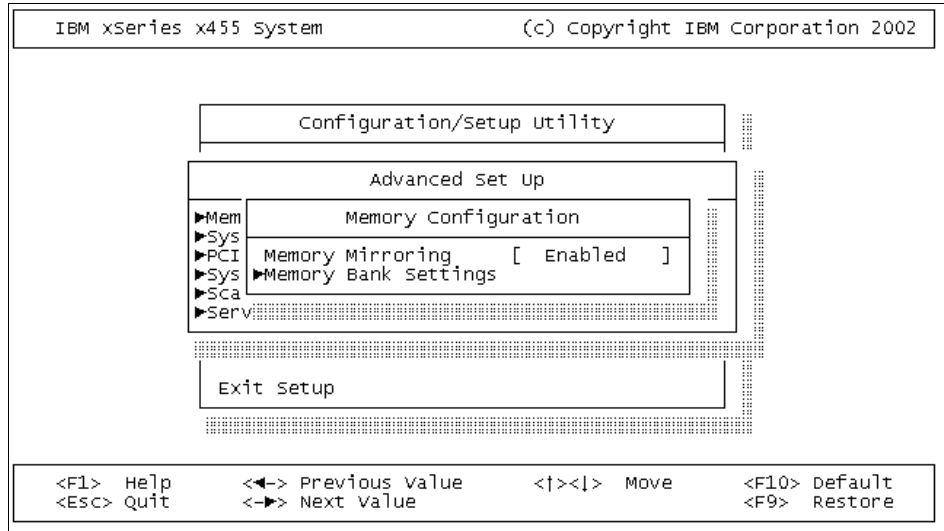


Figure 4-24 Memory Configuration menu

5. Press the right arrow key to change the Memory Mirroring value to **Enabled**.
6. Exit the Setup program and save the changes. If you made any changes while in the menus, the system will reboot. If you made no changes, you will be returned to the EFI Firmware Boot Manager menu.

Tip: Note the following tips:

- ▶ When memory is added, it should be detected automatically but may need two reboots to complete the process.
- ▶ When replacing faulty memory, you will need to re-enable the faulty DIMM(s) by running **Configuration/Setup**, saving the system changes and rebooting again.
- ▶ To ensure the your memory configuration is correct, enter **Configuration/Setup Utility** and verify that the settings in the Memory Configuration option match your configuration.

Error logging

A choice of four different logs can be turned on that can be used for general information or to diagnose problems. These are for:

- ▶ Debug information
- ▶ Debug trace
- ▶ Debug warning
- ▶ Debug error

To turn these on follow the steps below.

1. Select **Configuration/Setup -> System Event and Error Log**. The following screen will be displayed.

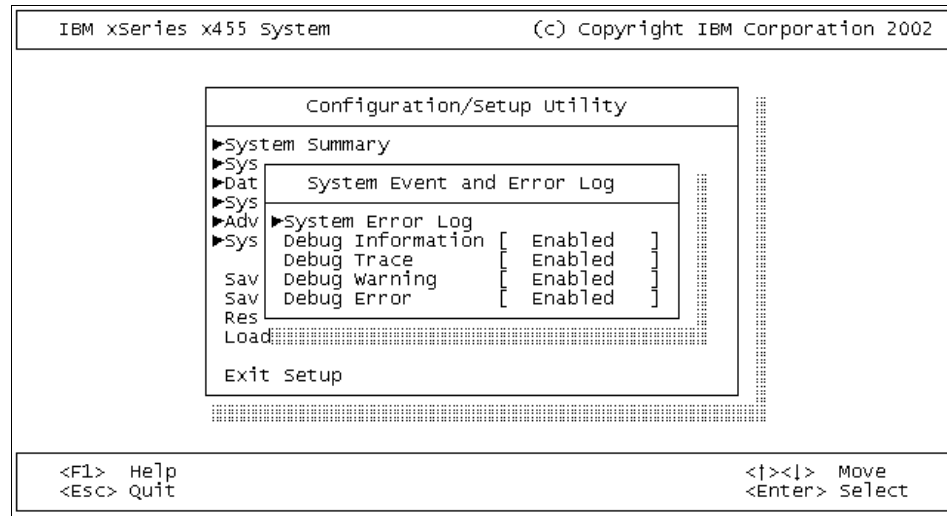


Figure 4-25 System Event and Error Log

2. Set either all or a subset of the logs to be enabled to suit your desired outcome.
3. Press Esc, then save your settings.

4.1.6 Diagnostic utility

The ROM-based diagnostics package can be started from the EFI Firmware Boot Manager menu by selecting the **Diagnostic Utility**. It is also possible to start this from the MemMap file system command prompt by typing `ami diag` (in our example, it was `fs1:`).

The Diagnostic utility looks similar to the IA-32 ROM version.

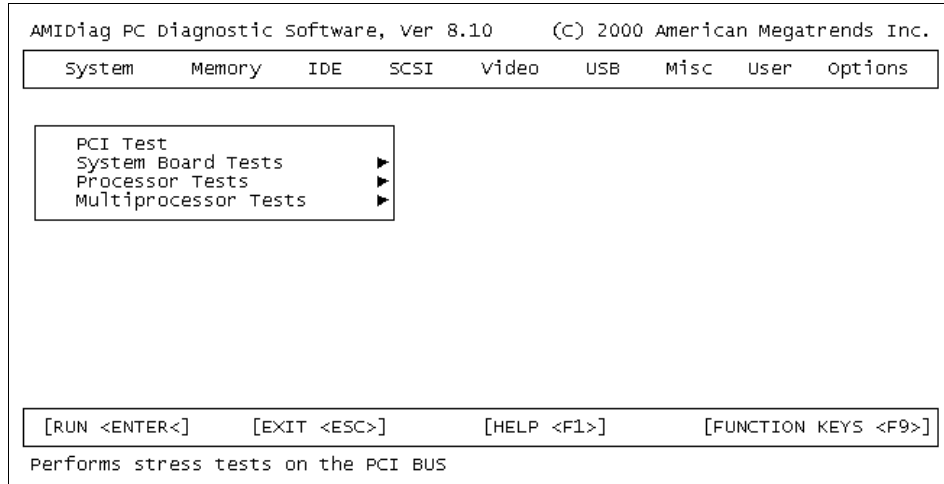


Figure 4-26 Diagnostic utility

In the utility, only non-destructive tests are available by default. To enable all the tests (including destructive SCSI Disk Write/Format Tests), choose **Options** → **Toggle Hidden Test Display**.

4.1.7 Boot Option Maintenance

The Boot Maintenance Manager menu allows you to modify the EFI boot menu options, including the ability to add to, change or delete from the boot menu. The boot order is important because that is the order that the system will try to boot from (similar to the boot sequence option in Setup in IA-32 based xSeries systems).

Figure 4-27 on page 123 shows the Boot Maintenance Manager and the options available.


```
EFI Boot Maintenance Manager ver 1.10 [14.61]

Main Menu. Select an Operation

    Boot from a File
    Add a Boot Option
    Delete Boot Option(s)
    Change Boot Order

    Manage BootNext setting
    Set Auto Boot TimeOut

    Select Active Console Output Devices
    Select Active Console Input Devices
    Select Active Standard Error Devices

    Cold Reset
    Exit

Timeout->[5] sec SystemGuid-->[CD9B6580-F7DF-11D7-AEBF-00000080FFFF]
SerialNumber-->[23a0014]
```

Figure 4-27 Boot Maintenance Manager menu

Backing up and restoring EFI boot options

It is important to back up the EFI options so that in the event that the EFI information is lost it can be restored. The backup can reside on the ESP for easy access but should also be backed up again as part of the normal backup schedule.

To back up the EFI boot options follow the steps below:

1. From the EFI Firmware Boot Manager select **EFI Shell [Built-in]**.
2. Change to the file system where you want to save the file, for example, fs5:.
3. Run the command **backup <filename>**, where <filename> is a valid filename. A screen similar to the following will be displayed.

```

fs5:\> backup x455.bak
backup: Variable NV+RT+BS 'Efi:Lang' DataSize = 3
backup: Variable NV+RT+BS 'Efi:ConIn' DataSize = 18
backup: Variable NV+RT+BS 'Efi:ConOut' DataSize = 16
backup: Variable NV+RT+BS 'Efi:ErrOut' DataSize = 16
backup: Variable NV+RT+BS 'blkio:EDD30' DataSize = 1
backup: Variable NV+RT+BS 'Efi:Driver0000' DataSize = C6
backup: Variable NV+BS 'Efi:DriverOrder' DataSize = 2
backup: Variable NV+BS
'E778C047-2121-4AD1-8236-090740619F05:Configuration' Data
Size = 2
backup: Variable NV+RT+BS 'EB704011-1402-11D3-8E77-00A0C969723B:MTC'
DataSize =4
backup: Variable NV+RT+BS 'Efi:Timeout' DataSize = 2
backup: Variable NV+RT+BS 'Efi:Boot0000' DataSize = 48
backup: Variable NV+RT+BS 'Efi:Boot0001' DataSize = 7E
backup: Variable NV+RT+BS 'Efi:Boot0002' DataSize = 68
backup: Variable NV+RT+BS 'Efi:Boot0003' DataSize = C4
backup: Variable NV+RT+BS 'Efi:Boot0004' DataSize = 10A
backup: Variable NV+RT+BS 'Efi:Boot0005' DataSize = 10A
backup: Variable NV+RT+BS 'Efi:Boot0006' DataSize = 97
backup: Variable NV+RT+BS 'Efi:Boot0007' DataSize = 97
backup: Variable NV+RT+BS 'Efi:BootOrder' DataSize = 10
backup: Completed - Success

```

Figure 4-28 Backing up EFI menu options

4. Exit back to the EFI Firmware Boot Manager.
5. The backup is now complete.

To restore the EFI boot options follow the steps below:

1. From the EFI Firmware Boot Manager select **EFI Shell [Built-in]**.
2. Change to the file system where you saved the backup file, for example, fs5:.
3. Run the command **restore <filename>**, where <filename> is the name of the backup file.
4. You will be presented with the following prompt:

Delete existing non-volatile efi variables? [y/n]

Answer Yes to this question. A screen similar to the following will be displayed.

```
fs5:\> restore x455.bak
Delete existing non-volatile efi variables? [y/n]y
restore: Variable NV+RT+BS 'Efi:Lang' DataSize = 3
restore: Variable NV+RT+BS 'Efi:ConIn' DataSize = 18
restore: Variable NV+RT+BS 'Efi:ConOut' DataSize = 16
restore: Variable NV+RT+BS 'Efi:ErrOut' DataSize = 16
restore: Variable NV+RT+BS 'Efi:Boot0000' DataSize = 48
restore: Variable NV+RT+BS 'Efi:Boot0001' DataSize = 7E
restore: Variable NV+RT+BS 'Efi:Boot0002' DataSize = B2
restore: Variable NV+RT+BS 'Efi:Boot0003' DataSize = B2
restore: Variable NV+RT+BS 'Efi:Boot0004' DataSize = 97
restore: Variable NV+RT+BS 'Efi:Boot0005' DataSize = 97
restore: Variable NV+RT+BS 'Efi:Boot0006' DataSize = 132
restore: Variable NV+RT+BS 'Efi:BootOrder' DataSize = E
restore: Variable NV+RT+BS 'Efi:Driver0000' DataSize = C6
restore: Variable NV+RT+BS 'Efi:DriverOrder' DataSize = 2
restore: Variable NV+RT+BS 'E4329F89-546B-11D4-9A39-0090273FC14D:RTC'
DataSize = 3
restore: Variable NV+RT+BS 'Efi:Timeout' DataSize = 2
restore: Variable NV+RT+BS 'EB704011-1402-11D3-8E77-00A0C969723B:MTC'
DataSize = 4
restore: Completed - Success
```

Figure 4-29 Restoring the EFI menu options

5. Exit back to the EFI Firmware Boot Manager.
6. The restore is now complete.

Adding a boot option

You are most likely to need to add a boot option in the following situations

- ▶ A default option has been deleted and needs re-adding.
- ▶ You have moved disks to another machine and have to update EFI with the required boot options.
- ▶ The system planar has been changed.
- ▶ You have developed a custom option that needs adding.

Below are the steps to add an option for the most common of these situations where the disks have been moved to another machine.

1. Run a backup as described in “Backing up and restoring EFI boot options” on page 123.

2. Go to the EFI Shell and type `reset -s` to shut down the machine.
3. Move the disks to the new machine.
4. Boot the new machine.
5. The EFI Firmware Boot Manager will be displayed with the default options or options previously set on the machine.
6. Go to the EFI Shell and run a restore as described in “Backing up and restoring EFI boot options” on page 123. Note that the file system from where the backup file was saved may be different from the one it is restored from.
7. Exit back to the EFI Shell if you have not already done so.
8. The new option may or may not appear as a boot option at this stage. If it has, this is now complete. If it does not appear follow the steps below.
9. Select **Boot Option Maintenance**.
10. Exit from Boot Option Maintenance. The new option should now appear.

Changing a boot option

You may wish to change the boot options to suit your environment. One of the most common reasons is to enable the server to PXE boot. If your x455 needs to be configured to PXE boot as the first option, follow the steps below:

1. Select the **Boot Maintenance Option** from the EFI Firmware Boot Manager menu. See Figure 4-2 on page 94.
2. Select **Change Boot Order** and press Enter. The following screen is displayed. (You may have slightly different options depending on the operating systems installed.)

```
EFI Boot Maintenance Manager ver 1.10 [14.61]

Change boot order.  Select an Operation

Windows Server 2003, Datacenter
EFI Shell [Built-in]
Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)
Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)/CDROM (Entry0)
Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)/CDROM (Entry1)
Acpi (PNP0A03,1)/Pci (4|0)/Mac (0002551F809E)
Acpi (PNP0A03,1)/Pci (4|1)/Mac (0002559F809E)
Save Settings to NVRAM
Help
Exit

Acpi (PNP0A03,1)/Pci (4|0)/Mac (0002551F809E)
Boot0004
```

Figure 4-30 Change boot order

- 3. Highlight the first PXE option. These end with Mac followed by a 12-digit hexadecimal number. Press U to move up so it is the first option. The number after Boot on the bottom line refers to the option currently selected and is how EFI enumerates that particular option.
- 4. Highlight the second PXE option and press U to move up so it is the second option. If you need to move an option down press D.
- 5. Select **Save Settings to NVRAM** and exit back to the EFI Firmware Boot Manager.

The EFI Firmware Boot Manager will now look like Figure 4-31 on page 128 and will PXE boot by default.

```
IBM                               Firmware Boot Manager                               x455
                                Version 2.00 [2.10] (c) copyright IBM Corporation 2003

Boot options: Enter number or ↑/↓ to select, 'Enter' to execute

1 Acpi (PNP0A03,1)/Pci(4|0)/Mac(0002551F809E)
2 Acpi (PNP0A03,1)/Pci(4|1)/Mac(0002559F809E)
3 Windows Server 2003, Datacenter
4 EFI Shell [Built-in]
5 Acpi (PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)
6 Acpi (PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)/CDROM(Entry0)
7 Acpi (PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)/CDROM(Entry1)

System options: Enter letter or ↑/↓ to select, 'Enter' to execute

    Boot Option Maintenance           Diagnostics           EFI Shell
    Configuration/Setup              Driver Setup          Flash

Messages:
```

Figure 4-31 New default boot order in Firmware Boot Manager

Deleting a boot option

Deleting boot options should be done with caution. The default options can be recreated, but options added after the installation of an operating system are difficult to recreate and require specialist knowledge of the EFI environment and how the option was constructed in the first place. To avoid potential problems you need to follow some general guidelines:

- ▶ Always ensure that you have a current backup of the EFI boot options.
- ▶ Optionally take a Windows Server 2003 “boot option” backup using the `nvrboot.efi` utility.
- ▶ Proceed with caution.

To delete an option follow the steps below:

1. Select **Delete Boot Option(s)** from the EFI Boot Maintenance Manager. The following screen will be displayed.

```
EFI Boot Maintenance Manager ver 1.10 [14.61]

Delete Boot Option(s).  Select an Option

Windows Server 2003, Datacenter
EFI Shell [Built-in]
Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)
Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)/CDROM(Entry0)
Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)/CDROM(Entry1)
Acpi(PNP0A03,1)/Pci(4|0)/Mac(0002551F808F)
Acpi(PNP0A03,1)/Pci(4|1)/Mac(0002559F808F)
Delete All Boot Options
Save Settings to NVRAM
Help
Exit

HD(Part1,SigB8D87F80-9D70-01C3-507B-9E5F8078F531)/\EFI\Microsoft\WINN
T50\ia6
4ldrBoot0006
  WINDOWS
```

Figure 4-32 Delete Boot Option(s) menu

2. Highlight the required option. In this example it is **Windows Server 2003 Datacenter**, and press Enter. The following screen will be displayed.

```

EFI Boot Maintenance Manager ver 1.10 [14.61]

Delete Boot Option(s).  Select an Option

Windows Server 2003, Datacenter
EFI Shell [Built-in]
Acpi(PNPOA03,0)/Pci(5|1)/Ata(Primary,Master)
Acpi(PNPOA03,0)/Pci(5|1)/Ata(Primary,Master)/CDROM(Entry0)
Acpi(PNPOA03,0)/Pci(5|1)/Ata(Primary,Master)/CDROM(Entry1)
Acpi(PNPOA03,1)/Pci(4|0)/Mac(0002551F808F)
Acpi(PNPOA03,1)/Pci(4|1)/Mac(0002559F808F)
Delete All Boot Options
Save Settings to NVRAM
Help
Exit

Delete selected Boot Option [Y-Yes N-No]: 78F531)/
\EFI\Microsoft\WINNT50\ia6 4ldrBoot0006
WINDOWS

```

Figure 4-33 Confirming the delete request

3. Select **Yes** to delete the option.
4. Save and exit.
5. The option has now been removed from the boot options.

Important: In the unlikely event that you have major problems and are not able to get back the default options, then connect an auxiliary console (ANSI terminal connected to the serial port) and press Shift+R when prompted to restore the EFI defaults.

Changing the TimeOut value

To change the TimeOut value follow the steps below:

1. Select **Set Auto Boot TimeOut** from the EFI Boot Maintenance Manager. The following screen will be displayed.
2. Set the timeout value.
3. Save and exit.

Other boot options

The other EFI Boot Maintenance Manager options allow you to invoke alternative boot options and set alternative input and output devices. For more information on how to use these options refer to the user manual.

4.2 Configuring scalable partitions

This section provides information and instructions for creating and deleting scalable partitions through the Configuration/Setup Utility program.

Note: The Scalable Systems Manager program, when available, is an extension to IBM Director that provides a graphical user interface for creating and managing scalable partitions. More information is detailed in “Scalable Systems Manager” on page 177.

Before you create or delete scalable partitions, read the following information:

- ▶ The Remote Supervisor Adapters must be connected to the same network. This is required to create and maintain the partition.
- ▶ PCI-X slot numbering starts with the primary node, including attached remote I/O enclosures, and continues with the secondary nodes, in numeric order of the logical node IDs.
- ▶ To change the primary node, delete the scalable partition from the current primary node (if it exists). Then create a new scalable partition from the new primary node.

4.2.1 Creating a scalable partition

Complete the following steps to create a scalable partition:

1. If you have not already done so, connect the SMP Expansion, RXE Expansion, and RXE Management cables. See “SMP Expansion connectivity” on page 67 and “Remote Expansion Enclosure connectivity” on page 70.
2. Start the Configuration/Setup Utility program from the EFI Firmware Boot Manager.
3. From the main menu, select **Advanced Setup**. See Figure 4-34 on page 132.

Two-chassis configuration:

- a. Select **2 x455 chassis** as the scalable partition configuration.
- b. Type a scalable partition ID.
- c. Type the ASM host name or IP address for the secondary node.
- d. If a remote I/O enclosure is attached (See Figure 4-36), complete the following steps:
 - i. Select the number of PCI-X slots available to each RXE Expansion Port on the remote I/O enclosure.
 - ii. Select the destination of Expansion Port A on the remote I/O enclosure.
 - iii. Select the destination of Expansion Port B on the remote I/O enclosure.

IBM xSeries x455 System		(c) Copyright IBM Corporation 2002	
Create scalable Partition settings			
Scalable Partition Configuration	[2 x455 Chassis]	
Scalable Partition ID	[00000001]	
Second ASM Host Name	[192.168.1.2]	
Third ASM Host Name	[]	
Fourth ASM Host Name	[]	
First RXE Configuration	[Port A 6 Slot, Port B 6 Slot]	
Expansion Port A Destination	[Primary Scalable Node, Port A]	
Expansion Port B Destination	[Secondary Scalable Node, Port A]	
Second RXE Configuration	[]	
Expansion Port A Destination	[]	
Expansion Port B Destination	[]	
Write scalable Partition settings			
.....			
<F1> Help	<<-> Previous Value	< >< > Move	
<ESC> Quit	<-> Next Value		

Figure 4-36 Two-node configuration with RXE-100

Four-chassis configuration:

- a. Select **4 x455 chassis** as the scalable partition configuration.
- b. Type a scalable partition ID.
- c. Type the ASM host names or IP addresses for the secondary nodes.

- d. If a remote I/O enclosure is attached in the first RXE configuration, complete the following steps:
 - i. Select the number of PCI-X slots available to each RXE Expansion Port on the remote I/O enclosure.
 - ii. Select the destination of Expansion Port A on the remote I/O enclosure.
 - iii. Select the destination of Expansion Port B on the remote I/O enclosure.
 - e. If a remote I/O enclosure is attached in the second RXE configuration, complete the following steps:
 - i. Select the number of PCI-X slots available to each RXE Expansion Port on the remote I/O enclosure.
 - ii. Select the destination of Expansion Port A on the remote I/O enclosure.
 - iii. Select the destination of Expansion Port B on the remote I/O enclosure.
7. Select **Write Scalable Partition Settings** and follow the prompts to exit the menu item.
 8. Save the changes and exit the Configuration/Setup Utility program.

4.2.2 Booting a scalable partition

Once a scalable partition has been created the boot procedure will reflect the new configuration. In a 2-node partition the initial boot screens will be displayed as shown in Figure 4-37 and Figure 4-38.

```

IBM                               Eserver                               x455
Build..0YKT12A 2003/10/18 08:52:22      (c) Copyright, IBM Corporation, 2003
Partition.....1 Processors....8[.4|.4|..|..] Checkpoint....42-53
Node.....1 of 2 Memory(GB)....8[.4|.4|..|..] Node:Cpu.....1:1
=====

```

Figure 4-37 Node 1 boot up screen

4.2.4 Deleting a scalable partition

Complete the following steps to delete the scalable partition:

1. Start the Configuration/Setup Utility on the primary server.
2. From the main menu, select **Advanced Setup**.
3. Select **Scalable Partition Settings**.
4. Select either **Delete ALL Scalable Partition Settings** or **Delete Local Scalable Partition Settings**.

Important: Do not select **Write Scalable Partition Settings** when you delete the scalable partition.

5. Exit the Configuration/Setup Utility program.
6. Turn off the servers.
7. If necessary, remove the SMP Expansion cables.
8. Restart the servers.

Error messages

When one of the scalability cables is not connected or faulty you will see the following error message:

```
9912 - Scalability cable fault on port 3, check cable, performance degraded
```

4.3 Installing Windows Server 2003

In this section, we guide you through the installation procedure for Windows Server 2003, Enterprise Edition and Datacenter Edition on the x455. The installation for each version is the same. It is also the same on a single node or multinode partition. We will simply refer to Windows Server 2003, which will apply to both editions.

We assume that you are familiar with the installation of Windows Server in the IA-32 platform, so we will highlight the important information and installation procedures that are specific to the 64 bit platform and x455.

4.3.1 Important information

There are a number of things you should consider before starting the installation. These are detailed below.

Microsoft Reserved Partition

The *Microsoft Reserved Partition* (MSR partition) reserves space on each disk drive for subsequent use by the operating system software. For example, dynamic disk configuration is stored in the MSR partition.

During the installation, the MSR partition is automatically created. When initially created, the size of the MSR partition depends on the size of the disk drive:

- ▶ On drives less than 16 GB in size, the MSR partition is 32 MB.
- ▶ On drives greater than or equal to 16 GB, the MSR partition is 128 MB.

Every GPT disk used by Windows must contain an MSR partition. The order of the partitions on the disk should be ESP (if the disk is a boot disk), MSR partition, primary data partition(s). It is particularly important that the MSR partition is created before other primary data partitions.

Note: Each Windows *boot drive* must contain an ESP, an MSR partition, and at least one basic data partition that contains the operating system. Each GPT *data drive* must contain at least an MSR partition and one basic data partition. The MSR partition is not accessible to the operating system users by any means.

The order of partitions on the Windows boot disk is shown in Figure 4-39.




Figure 4-39 Windows boot GPT disk structure

For information on the partitioning required by a Windows 64-bit installation, see:

http://www.microsoft.com/technet/prodtechnol/winxppro/reskit/prkb_cnc_helm.asp

GPT disks

The dynamic disk concept introduced in Windows 2000 is not affected by introducing GUID Partition Table (GPT) disks in Windows Server 2003. Both Master Boot Record (MBR) and GPT disks can be present in a single dynamic disk group. Volume sets can span both MBR and GPT disks.

The OS loader and the boot partition must reside on a GPT disk. Other hard disks in the system may be either MBR or GPT.

4.3.2 Preparing to install

Before you install Windows 2003 Server it is important that all hardware and software is prepared properly to ensure a smooth and trouble-free installation. Perform the following tasks before you begin the installation:

1. Before powering on the server, verify that the hardware is configured properly, including:
 - All hardware options are installed.
 - All cables are attached and secured to the correct ports.
 - Optional RXE-100 Expansion Enclosure is connected.
 - All EXP storage expansion units are connected.
 - All FASt storage expansion units are connected.
2. Ensure that firmware levels of the following devices have been updated to the latest levels:
 - System EFI/SAL
 - Diagnostics
 - Remote Supervisor Adapter
 - RXE-100; refer to “Remote Supervisor Adapter firmware update” on page 116
 - ServeRAID adapter; refer to “ServeRAID configuration CD” on page 110
 - Fibre Channel host bus adapters
3. Ensure that you have all the required OS drivers, including those for:
 - Broadcom Ethernet controllers
 - Active PCI
 - Remote Supervisor Adapter
 - ServeRAID adapter
 - Fibre host bus adapter

These can be downloaded from the x455 driver page:

<http://www.pc.ibm.com/qtechinfo/MIGR-53575.html>

4. Optionally configure the two internal disks in a RAID-1 array. See “Configuring SCSI controller and internal disks” on page 105.
5. Optionally, delete all existing partitions (including the system partition and Microsoft Reserved partition) using the `diskpart` command, as described in “Checking a disk” on page 100.

Note: Earlier problems that affected the installation of Windows Server 2003 on the x455 with an RXE-100 attached have now been resolved. It is no longer necessary to disconnect the RXE-100 before the installation.

4.3.3 Installation

In this section we discuss installation.

Tip: Official installation instructions are typically available on the Web. However, at the time of writing none were published for the x455. Once they are available, they will be accessible at:

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

Select the x455 (8855) and select **OS Installation** document type.

The installation here assumes that you will be installing Windows Server 2003 on a brand new pair of disks, disks that have no data on them, or disks that have data that can be destroyed.

If you are installing on disks that have data that you want to keep, you will see additional messages. Please read this information carefully and refer to earlier sections of this chapter for additional help.

Tip: If the disk you install to is still an MBR disk, the installation procedure will prompt you to convert the disk to a GPT structure.

Starting the installation CD automatically

To start the installation automatically, proceed as follows:

1. Insert the installation CD in the drive.
2. From the EFI Firmware Boot Manager menu, select the CD-ROM entry in the list (as shown in Figure 4-40 on page 140) and press Enter.

```
IBM                               Firmware Boot Manager                               x455
Version 2.00 [2.10] (c) copyright IBM Corporation 2003

Boot options: Enter number or ↑/↓ to select, 'Enter' to execute

1 EFI Shell [Built-in]
2 Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)
3 Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)/CDROM (Entry0)
4 Acpi (PNP0A03,0)/Pci (5|1)/Ata (Primary,Master)/CDROM (Entry1)
5 Acpi (PNP0A03,1)/Pci (4|0)/Mac (0002551F809E)
6 Acpi (PNP0A03,1)/Pci (4|1)/Mac (0002559F809E)

System options: Enter letter ↑/↓ or to select, 'Enter' to execute

Boot Option Maintenance           Diagnostics           EFI Shell
Configuration/Setup              Driver Setup          Flash

Messages:
```

Figure 4-40 Boot Manager menu - Selecting the CD-ROM

For three seconds, there will be a message on the screen that says Press any key to boot from CD-ROM. If you do not press any key, the installation will not start and you will return to the Firmware Boot Manager.

Starting the installation CD manually

To start the installation manually, proceed as follows:

1. Insert the installation CD-ROM into the drive.
2. Enter the EFI shell from the EFI Firmware Boot Manager menu. Access the CD-ROM drive as described in “File operations in the EFI shell” on page 97.
3. Run the SETUPLDR.EFI program. If the CD-ROM drive is mapped as fs2: and fs3:, the command line looks similar to:

```
fs3:\> setupldr
```

Tip: With a bootable CD, there will be two file system entries, typically fs2: and fs3:. The first of the two entries will contain all the program files, such as I386, and the second will contain the boot loader files.

Text-mode setup

The installation process will now begin. The screens in this section will not necessarily follow consecutively. Only those important to 64-bit and the x455 are listed. For all other screens please read carefully and choose options appropriate to your installation.

If you are using the serial port as an auxiliary console, when you see the Setup is starting Windows message on the video display, the auxiliary console (a terminal session via the serial port) will stop functioning. This is normal. If you are using the RSA Remote Console feature, the session will continue.

If you install the operating system on a disk without the EFI System Partition (that is, a first-time installation), you will be asked to create one, as shown in Figure 4-41.

```
Windows Server 2003, Datacenter Edition Setup
=====
Setup could not locate an existing system partition.

System partitions contains diagnostic or hardware configuration
programs, programs to start operating systems (such as Windows), or
other manufacturer-supplied programs.

Setup will try to create a system partition for you automatically.

To allow Setup to create a system partition, press ENTER.
If you want to create a system partition on your own, press ESC.

Enter=Continue    Esc=Cancel
```

Figure 4-41 ESP creation during Windows installation

Allow Setup to create the EFI System Partition by pressing Enter.

Tip: At this point, in addition to the EFI System Partition, Setup also creates the Microsoft Reserved Partition (see “Microsoft Reserved Partition” on page 137).

If the disk you are installing on is still the old Master Boot Record (MBR) partitioning structure, you should delete all partitions, and press S to change to a GUID Partition Table (GPT) partitioning structure. The OS loader and the boot partition must reside on a GPT disk (see figure 4-42).

```
Windows Server 2003, Datacenter Edition Setup
=====

Setup could not create a system partition automatically.

System partitions contains diagnostic or hardware configuration
programs, programs to start operating systems (such as Windows),
or other manufacturer-supplied programs.

You need to create a partition manually and mark it as a system
partition.

If all the disks are MBR disks then convert one of disks where
you want the system partition to reside to GPT disk.

To convert a disk to GPT delete all the partitions on the disk,
and select "S=Change Disk Style" option. When setup prompts for
auto creation of system partition, press ENTER.

    · To continue, press ENTER.
```

Figure 4-42 Convert to GPT disk

Tip: Once created, the MSR partition is a reserved partition and cannot be deleted from the Disk partitioning menu. To delete the MSR partition, the entire disk can be deleted as described in “Checking a disk” on page 100.

After both partitions are automatically created, the Disk partitioning screen is displayed, as in Figure 4-43 on page 143.

```
Windows Server 2003, Datacenter Edition Setup

The following list shows the existing partitions and
unpartitioned space on this computer.

Use the UP and DOWN ARROW keys to select an item in the list.

• To set up windows on the selected item, press ENTER.
• To create a partition in the unpartitioned space, press C.
• To delete the selected partition, press D.

17278 MB Disk 0 at Id 0 on bus 0 on symmpi [GPT]
  -: Partition1 [FAT]                173 MB ( 172 MB free)
  -: Partition2 [Reserved]          126 MB ( 0 MB free)
  Unpartitioned space                16980 MB

8679 MB Disk 0 at Id 0 on bus 0 on symmpi [GPT]
  -: Partition2 [Reserved]           31 MB ( 0 MB free)
  Unpartitioned space                8647 MB

ENTER=Install C=Create Partition D=Delete Partition F3=Quit
```

Figure 4-43 Disk partitioning screen

These partitions are as follows:

- ▶ Partition1 [FAT]: EFI System Partition
- ▶ Partition2 [Reserved]: Microsoft Reserved Partition (both disks)

Do not attempt to install the operating system into the first two partitions. Instead, use the unpartitioned space on the disk to create the data partition. The internal disk is recommended if you have both internal and external disks.

The sizes of the ESP and the MSR partition are calculated according to the rules described in “EFI System Partition” on page 33 and “Microsoft Reserved Partition” on page 137, respectively.

Tip: Do not remove or resize the ESP and MSR partitions after they are created. The ESP partition appears as Partition1 [FAT] and MSR partition as Partition2 [Unknown] in the Disk partitioning menu during text-mode setup.

Continue to the formatting screen. There is now an option to quick format (see Figure 4-44 on page 144).

```
Windows Server 2003, Datacenter Edition Setup
=====

The partition you selected is not formatted. Setup will now
format the partition.

Use the UP and DOWN ARROW keys to select the file system
you want, and then press ENTER.

If you want to select a different partition for Windows,
press ESC.

Format the partition using the NTFS file system (Quick)
Format the partition using the FAT file system (Quick)
Format the partition using the NTFS file system
Format the partition using the FAT file system
```

Figure 4-44 Formatting screen

Choose an option then follow the instructions to complete the text-mode portion of the installation.

At the end of the text-mode portion, a new Windows entry is added to the EFI Firmware Boot Manager menu. After you restart the system at the end of text-mode setup, the entry is launched automatically. (If Windows does not boot automatically, select it from the EFI Firmware Boot Manager menu.)

GUI setup

Complete the standard Windows Server 2003 installation as per the instructions on the screen and those provided in the product documentation. This is similar to the installation of Windows 2000/2003 on IA-32 platforms.

4.3.4 Post-setup phase

Once Windows has been installed, you should perform the following steps.

Installing additional drivers

Specific device drivers are available for the x455 that are not included with the base operating system. Drivers that you should obtain separately for the operating system include:

- ▶ Broadcom Ethernet controller
- ▶ Remote Supervisor Adapter service processor

► Active PCI-X controller

The CDs that ship with the server may contain these drivers, but there are likely to be later versions available from the IBM Web site. The latest versions can be downloaded from the x455 device driver page:

<http://www.pc.ibm.com/qtechinfo/MIGR-53575.html>

Important: Drivers that were developed for any 32-bit versions of Windows will not work in the 64-bit version of Windows. You must use 64-bit drivers.

After Windows Server 2003 is installed, Device Manager will report a number of unknown devices. These unknown devices correspond to specific components, as shown in Figure 4-45.

Some device drivers released for use with Windows Server 2003 may not be digitally signed. Therefore, when you install the drivers, you may have to accept the warning that you are about to install an unsigned driver.

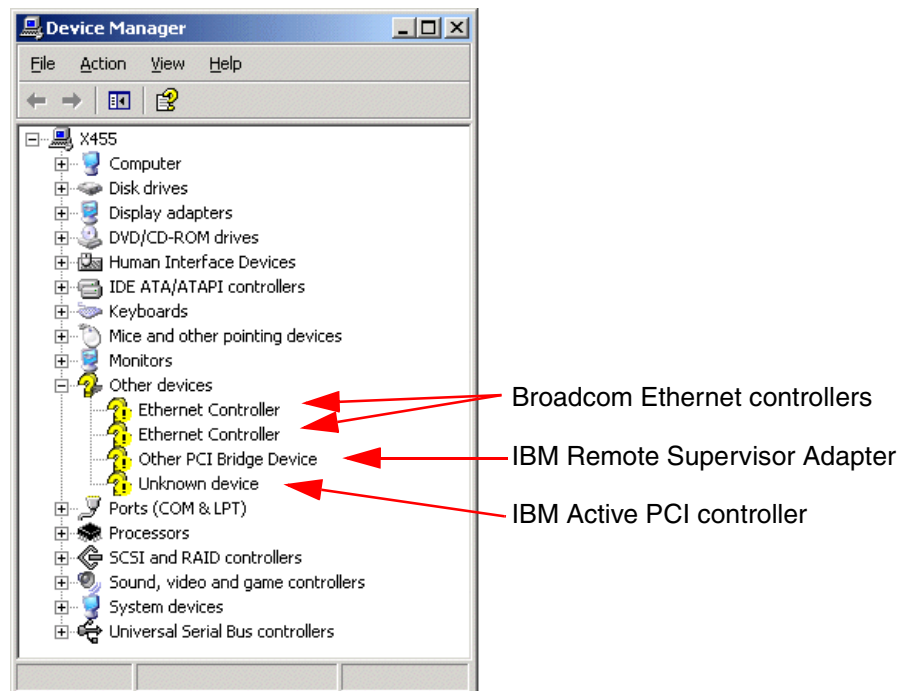


Figure 4-45 Device Manager after Windows installation

When the Windows Server 2003 installation has completed, you should install drivers for each of the unknown devices. For each one, right-click the entry in

Device Manager and click **Update Driver**. Select the options that let you specify a location for the drivers and point the installer to the directory where you downloaded the drivers.

If you have other devices installed, such as FASTT Fibre Channel or ServeRAID adapters, update or add the drivers for these. ServeRAID device drivers should be at the same level as the installed ServeRAID firmware.

Figure 4-46 shows Device Manager after the Active PCI, RSA, and Broadcom Gigabit Ethernet controllers drivers have been installed.

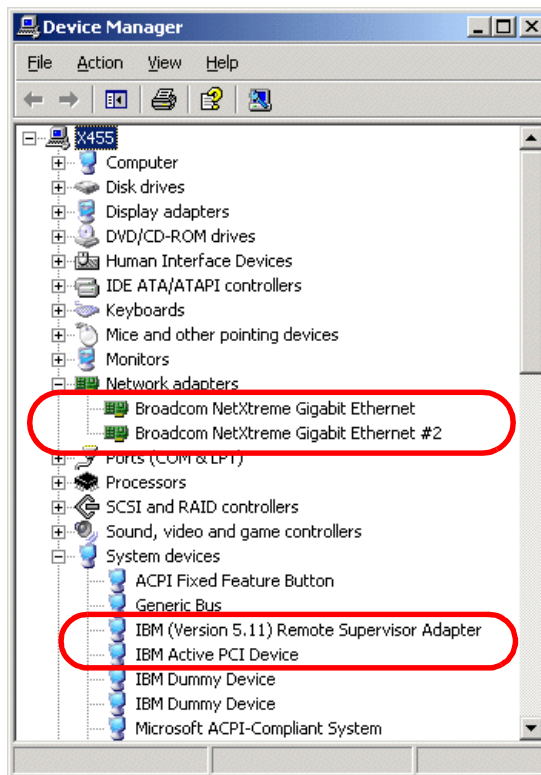


Figure 4-46 Device Manager after device driver updates

Installing additional drivers - Multi-node

The installation of a driver in a multi-node partition is identical to a single node partition. The only difference is that some of the hardware will be multiplied by a factor of 2 or 4, depending on the partition size. Below are some screen shots of what you can expect to see in a 2-node partition.

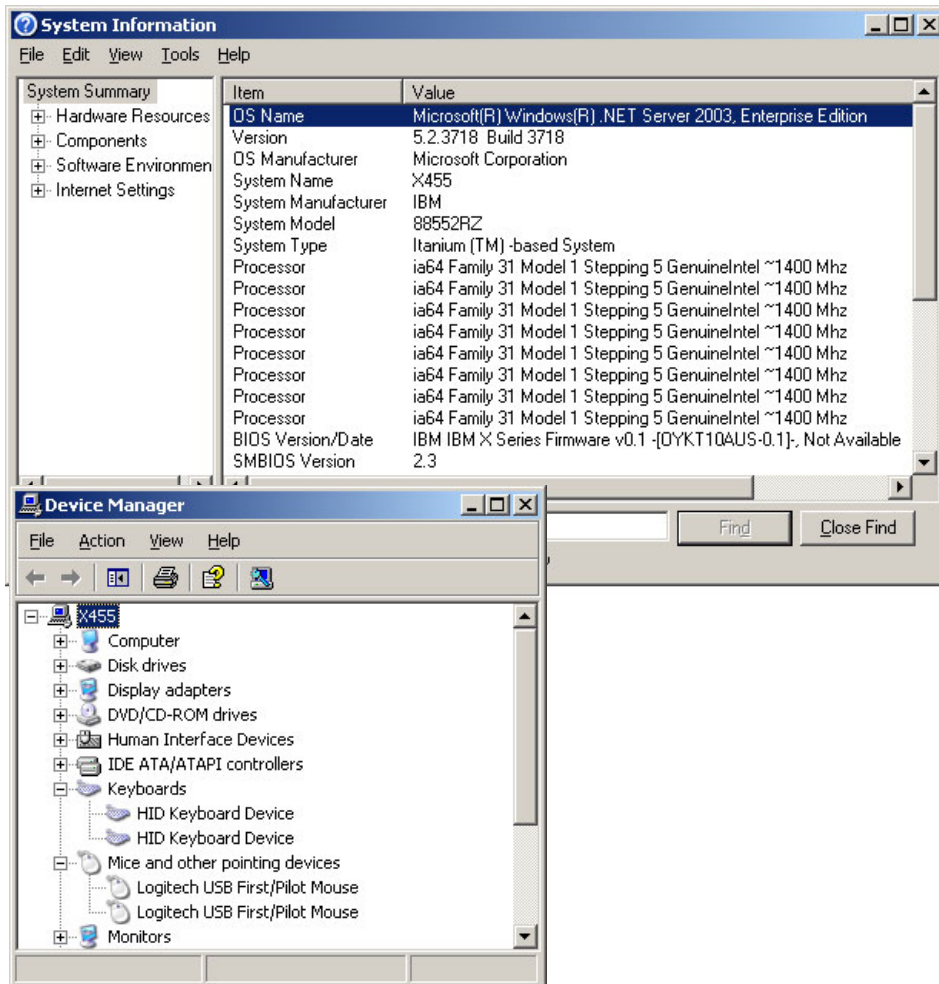


Figure 4-47 Multinode configurations “see” the devices from both nodes

Accessing EFI System Partition from Windows

To access the EFI System Partition from within the operating system, use the following utility:

```
mountvol DriveLetter: /s
```

Where DriveLetter is any spare drive letter in the system.

The ESP will be accessible for both read and write operations. Altering the contents of the ESP should be carefully considered. It should not be used as a general storage area. Recommended uses of the ESP are:

- ▶ BIOS update files
- ▶ Diagnostic update files
- ▶ RSA firmware update
- ▶ Firmware update files for other devices
- ▶ Backup of the EFI/SAL
- ▶ Backup of the EFI Firmware Boot Manager boot options
- ▶ Backup of the Windows boot record
- ▶ Useful EFI utilities

Important: Do not alter the content of the EFI directory on the EFI System Partition. As described in “EFI System Partition” on page 33, this directory contains the files necessary to boot installed operating systems.

Always make a backup of important files that can be taken off site.

Exporting and importing the Windows EFI boot record

During the installation process, the Setup program adds a new entry to the EFI Firmware Boot Manager menu.

It will be added to the top of the menu so that if no keys are pressed once the EFI Firmware Boot Manager menu appears, Windows will automatically start. The default delay is 30 seconds but may be changed using the Set Auto Boot Timeout option in the Boot Option Maintenance menu.

The EFI Firmware Boot Manager records are stored in NVRAM. If the boot record is deleted from EFI NVRAM, there is no other way to start Windows except by restoring the boot record from backups. To back up and restore Windows boot records use the NVRBOOT.EFI utility.

In this section we show you how to save the Windows options in the EFI Firmware Boot Manager menu and how to recover it in case of an NVRAM failure.

Tip: This procedure actually saves all the boot options in the EFI Firmware Boot Manager, but if you require all the options then always use the EFI **backup** command. The **backup** command is designed for this, nvrboot.efi features may change as Microsoft make modifications to it.

In “GUID Partition Table disk” on page 31, we describe how GPT disks are identified by the global unique identifier (GUID). The entries in the EFI Firmware Boot Manager menu use the GUID to identify which device to boot from.

There may be several scenarios when the EFI Firmware Boot Manager menu will need to be restored:

- ▶ Disks rearrangement - If the disks are moved from one server to another, the NVRAM contents and the disk contents are out of sync, meaning that there will be no record in the new server's EFI Firmware Boot Manager menu to point at the recently added disks.
- ▶ Planar board replacement - The planar board contains the system's NVRAM, so when it is replaced, the operating system boot entry in the EFI Firmware Boot Manager menu will need to be manually restored.
- ▶ If the boot record is deleted or becomes corrupt.

To export the Windows EFI boot menu record, perform the following steps:

1. Invoke the EFI shell from the EFI Firmware Boot Manager menu. Change to `fsn:`, where `n` is the system partition (see "The EFI shell" on page 95).
2. Change the directory to `MSUTIL` (`cd msutil`).
3. Run the `NVRBOOT.EFI` utility. See Figure 4-48 on page 150.
4. Enter `X` (for Export) from the OS Boot Options Maintenance Tool menu.
5. Enter the number of Windows boot records you want to export, or `*` to export all.
6. Type a file name to store the configuration in and press Enter. A confirmation message, as shown in Figure 4-48 on page 150, will be displayed. The file will be saved in the root of the EFI System Partition. Back up the file during your regular backup.
7. Press `Q` and exit the utility.

Note: The utility can export/import all boot menu records, but can only display and modify the Windows boot record.

```

NVRBOOT: OS Boot Options Maintenance Tool [Version 5.2.3683]

* 1. Windows Server 2003, Datacenter
  2. EFI Shell [Built-in]
  3. Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)
  4 Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)/CDROM(Entry0)
  5 Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)/CDROM(Entry1)
  6. Acpi(PNP0A03,1)/Pci(4|0)/Mac(0002551F809E)
  7. Acpi(PNP0A03,1)/Pci(4|1)/Mac(0002559F809E)

* = Windows OS boot option

(D)isplay (M)odify (C)opy E(x)port (I)mport (E)rase (P)ush (H)elp
(Q)uit

Select>
Select> x
Enter OS boot option to export (* = All - Maximum of 30): 1
Enter EXPORT file path: win2003
Saving boot option 1...
Saved Boot Option 1 to file: win2003.cfg
Use Import command to retrieve saved boot option

Press enter to continue
Select>q

```

Figure 4-48 Exporting Windows boot record with NVRBOOT.EFI

To recover the boot menu record, use the **(I)mport** command from the NVRBOOT utility.

To illustrate the export/import process, here is an example. If you were to take a set of disks from one x455 to another, you would do the following:

1. Export the Windows boot entry from the first server and save it on the EFI System Partition (which is a partition on the disk).
2. Insert the disks into the second server.
3. Import the file.
4. Verify that the import was successful by booting from the newly created boot entry.

4.4 Installing Linux

The goal of the project to port the Linux kernel to the Itanium platform was to have a single optimized port of the Linux kernel for every Itanium-based machine. The kernel has been available at <http://www.kernel.org> since early 2000.

The development model used to create this new porting of the kernel was the same as used for other Linux development work: Open source developers working together around the world on a cooperative effort to make it available before any other operating system and even before the official release of the processor itself. This is one of the key points of Linux' success, and also of this Itanium 2 port of the kernel.

This effort was also possible thanks to many companies and institutions that contributed on this development effort, such as Intel, IBM, HP, Caldera (now SCO Group), the European Organization for Nuclear Research (CERN), Red Hat, SGI, SUSE LINUX, Turbolinux, and VA Linux Systems (now VA Software).

Each company distributed the work among their developers to achieve the best results. The main tasks were:

- ▶ Port the kernel itself
- ▶ Initial port of GNU Compiler Collection (GCC) to IA-64 architecture
- ▶ Port of GAS, emacs, ld, GNUPro toolkit
- ▶ Port of performance, measurement, and analysis tools

Intel provided the IA-32 and IA-64 platform specifications and the EFI, and helped on the Apache port. The rest of the work was distributed among the companies involved, and also by private individuals around the world.

Linux for IA-64 provides a pure 64-bit kernel but maintains compatibility at the API level with Linux for x86 wherever this is possible. The error messages, the system signal, and the ioctl codes will remain as compatible as possible between the two platforms. Also in this version, as started in the Linux kernel from Version 2.1.126 onwards, developers tried to minimize changes to make the code as platform-independent as possible, while optimizing the platform-specific features to improve performance on Itanium processors.

The IA-64 Linux kernel follows the standards defined by Intel and other companies. These include:

- ▶ EFI (Extensible Firmware Interface); see:
<http://developer.intel.com/technology/efi/>
- ▶ DIG64 (Developers Interface Guide for 64-bit Intel Architecture Servers); see:
<http://www.dig64.org>

- ▶ System V UNIX ABI (Application Binary Interface); see:
<http://www.linuxbase.org/spec/refspecs/elf/gabi4+/contents.html>

4.4.1 Linux IA-64 kernel overview

The following information is helpful if you plan to recompile or port some software to IA-64 kernel. The main objective of this section is to introduce some basic concepts about the Linux kernel and the most important differences between the IA-32 version and the new 64-bit IPF (Itanium Processor Family, or IA-64) architecture. Your preferred Linux distribution may have done this recompilation for you, but it is very common to recompile software under Linux, for example, if you get only the source code and not an installable package.

Data types

In 1996, The Open Group decided to extend the *Single UNIX Specification for 64-bit Systems* to create a new standard on 64-bit for any UNIX operating system. Linux kernel developers attempt to follow the Open Group standards as much as possible, so the Linux IA-64 kernel uses the LP64 Data Model, as defined by The Open Group. This data model is LP64 because it allows long and integer types up to 64 bits. 32-bit Intel processors use the LP32 Data Model. This model is an industry standard and it is used in all other 64-bit Linux and UNIX implementations.

Some of the usual data types on C change between the LP64 and the LP32 Data Models, as shown in Table 4-4.

Table 4-4 Comparison table of C types between LP32 and LP64

Data type in C	LP32	LP64
char	8	8
short	16	16
int	32	32
long	32	64
long long	64	64
pointer	32	64
enum	32	32
float	32	32
double	64	64

Data type in C	LP32	LP64
long double	128	128

Keep these comparisons in mind while programming or porting applications to any 64-bit operating system, not only Linux. It is easy to assume that every data type is the same size, and that is why the LP Data Model standard was created.

Note: The LP64 Data Model example discussed in this publication is not the only 64-bit Data Model defined by The Open Group. There are other models, such as LP64, ILP64, or LLP64 for 64-bit; and ILP32 or LP32 for 32-bit. For more information about the data models defined by The Open Group and also about their standardization activities, go to:

<http://www.opengroup.org>

Byte order

The IA-64 kernel of Linux also uses standard byte order at a byte level. The native byte order is little-endian, while big-endian processes are still possible.

All Intel processors are little-endian based. That means that the bytes at lower addresses have lower significance given a 16-bit or 32-bit word (the word is stored “little-end-first”). In big-endian architectures, the bytes addresses have higher significance (the word is stored “big-end-first”). The IA-64 Linux kernel uses little-endian by default, but allows for the possibility of using big-endian byte order.

4.4.2 Choosing a Linux distribution

Unlike other operating systems, when we talk about Linux, we are not talking about one company, one version, or one distributor. Linux is the kernel, the base of the system, and there are many companies that develop Linux distributions built around the kernel.

In this section we discuss the major Linux distributors that work with IBM, and their different products available for the Itanium 2 platform: Red Hat and UnitedLinux.

Note that the IBM Linux support page is:

<http://www.pc.ibm.com/qtechinfo/MIGR-48NT8D.html>

Here you can find information about the Linux distributions supported by IBM, and documents, information, and drivers for IBM products.

Note: Debian Linux is also available for Itanium 2 platforms, but we will not discuss this Linux distribution because it is not supported by IBM.

Red Hat Linux for Itanium

Red Hat is a major Linux distributor and has two products for the Itanium processors:

- ▶ Red Hat Advanced Server 2.1
- ▶ Red Hat Enterprise Linux AS 3.0, the replacement for Red Hat AS 2.1

Both products support Itanium 2 platforms, and we made our tests with both also. In Table 4-5 you have a comparison between these products.

Table 4-5 RH AS 2.1 and RH EL AS 3.0

Feature	Red Hat AS 2.1	Red Hat EL AS 3.0
Maximum CPUs supported	8	16
Maximum memory	16 GB	64 GB
Maximum SCSI devices	128	256
Kernel version	2.4.18	2.4.21
Desktop GUI	GNOME	Red Hat Bluecurve
Compiler	GCC 2.9.6	GCC 3.2
Logical Volume Manager	No	Yes

For the latest information on Red Hat's plans and IBM's support plans, go to:

<http://www.redhat.com/software/rhel/as/>

<http://www.pc.ibm.com/us/compat/nos/redchat.html>

The UnitedLinux initiative

In 2002, four of the companies that were developing different Linux distributions—Conectiva, The SCO Group, SUSE LINUX and Turbolinux—announced that they will work together to build a unique base for a Linux distribution. The primary idea is to create a base system from which each company will develop their products, solutions, and services.

The official definition of UnitedLinux, as found on their Web site is:

"UnitedLinux is a standards-based, worldwide Linux solution targeted at the business user and developed by Conectiva, The SCO Group, SUSE LINUX, and Turbolinux. Designed to be an enterprise-class, industry-standard Linux

operating system, UL provides a single stable, uniform platform for application development, certification, and deployment and allows Linux vendors, Independent Software Vendors (ISVs), and Independent Hardware Vendors (IHV) to support a single high-value Linux offering rather than many different versions.”

UnitedLinux attempts to offer enterprise-level solutions based on Linux, starting from a standard base that each distributor can freely adapt to the needs of its customers. UnitedLinux will also be based on all the Linux standards defined as of this date, such as the Linux Standard Base standard. The main advantage for customers is that they can work with any of UnitedLinux companies, with any of their products, and the base system, the configuration files, and the file locations will remain the same, at the same place, and with the same properties. This is a main advantage for customers that use Linux, but do not have a unified distribution across their systems.

For more information about the UnitedLinux initiative, go to:

<http://www.unitedlinux.com/en/info/faqs.html>

The IPF (IA-64) Linux plans for each of the four UnitedLinux partners are as follows:

► **Conectiva Linux**

Conectiva has not officially released any version of their Conectiva Linux product for any Itanium processor. To find out about Conectiva’s products and services, go to:

<http://www.conectiva.com>

► **Caldera SCO Linux 64**

The SCO Group, previously called Caldera, has announced a version of their SCO Linux product for the Itanium family. This product is called Caldera SCO Linux 64.

To find out about SCO’s products on Linux for 64-bit platforms or IBM support for SCO products, go to:

<http://www.sco.com/products/>

<http://www.pc.ibm.com/us/compat/nos/thescocgroup.html>

► **Turbolinux 7.0 for Itanium**

Turbolinux also has a product for Itanium processors, Turbolinux 7.0 for Itanium, but it is not their latest version, which is 8.0.

To find out about Turbolinux products on Linux for 64-bit platforms or IBM support for SCO products, go to:

<http://www.turbolinux.com/products>

<http://www.pc.ibm.com/us/compat/nos/turbo linux.html>

- ▶ SUSE LINUX Enterprise Server 8 for IA-64

We also installed SUSE LINUX Enterprise Server 8 with a 2.4.19 kernel optimized for Itanium 2.

For information on IBM Support of SUSE LINUX products, go to:

<http://www.pc.ibm.com/us/compat/nos/suse.html>

4.4.3 Installing SUSE LINUX Enterprise Server 8.0

In this section we describe the installation procedure of the SUSE LINUX Enterprise Server (SLES) 8 and the issues you may find.

Tip: Official installation instructions are typically available on the Web. However, at the time of writing none were published for the x455. Once these are available, they will be accessible at:

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

Select the x455 (8855) and select **OS Installation** document type.

IBM employees can access draft installation instructions from:

<http://nosinfo.raleigh.ibm.com>

Requirements and hardware used

According to the *Installation and Administration Addendum* found in the doc directory on the first CD of SLES 8, the requirements to use this Linux distribution are:

- ▶ An Itanium or Itanium 2 processor
- ▶ At least 256 MB of main memory
- ▶ A hard disk larger than 10 GB
- ▶ Firmware version EFI 1.0.2 or 1.1.0 (preferred)

Our test system was a multi-node configuration with two pre-production x455 units with:

- ▶ Four Itanium 2 processors running at 1.4 GHz per server
- ▶ Four GB of DDR RAM per server
- ▶ EFI Version 1.10

We also connected an RXE-100 expansion enclosure during the installation.

Before starting the installation

Before starting the installation, we recommend that you get and apply the latest updates and firmware. You can download them from the x455 driver matrix:

<http://www.pc.ibm.com/qtechinfo/MIGR-53575.html>

SLES 8 supports LVM volumes during the installation, so you can install it on LVM partitions. We used only internal disks for our tests. Notice that the default file system used on SUSE LINUX is the journaling file system ReiserFS. For more information about the ReiserFS file system, go to:

<http://www.namesys.com/>

Installation procedure

The installation procedure of SLES Version 8 is like any other SUSE LINUX installation, except that you work on a different architecture, so the bootable CD-ROM will not boot directly, as it would do on any IA-32 architecture.

These are the steps to follow for a SLES 8 installation:

1. Power on the server and insert the first SUSE LINUX CD-ROM in the drive.
Although it is not necessary to insert the CD at this point, it makes the installation process easier because the CD-ROM file system is automatically detected and mapped.
2. Wait for the EFI Firmware Boot Manager menu to appear.
3. Select the first entry for the CD-ROM drive. In the x455, this is the ACPI entry with the strings "PCI" and "ATA(Primary/Master)". Press Enter (Figure 4-49 on page 158).

```
IBM                               Firmware Boot Manager                               x455
                                Version 2.00 [2.10] (c) copyright IBM Corporation 2003

Boot options: Enter number or ↑/↓ to select, 'Enter' to execute

1 EFI Shell [Built-in]
2 Acpi (PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)
3 Acpi (PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)/CDROM(Entry0)
4 Acpi (PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)/CDROM(Entry1)
5 Acpi (PNP0A03,1)/Pci(4|0)/Mac(0002551F809E)
6 Acpi (PNP0A03,1)/Pci(4|1)/Mac(0002559F809E)

System options: Enter letter ↑/↓ or to select, 'Enter' to execute

    Boot Option Maintenance           Diagnostics           EFI Shell
    Configuration/Setup              Driver Setup          Flash
Messages:
```

Select this option to boot the CD-ROM

Figure 4-49 EFI Boot Manager - Select this option to boot from the CD-ROM

Manually starting the CD-ROM: If you did not have the CD-ROM inserted before the EFI Boot Manager menu loaded, do the following:

1. Insert SUSE LINUX Disk 1 into the drive.
 2. Select **EFI Shell** and press Enter.
 3. Enter the command **map -r** to remap all file systems.
 4. Enter **exit** to go back to the Boot Manager menu.
 5. Select the CD-ROM entry **Ata(Primary,Master)/CD**.
4. The SUSE LINUX welcome screen appears.
- The installation procedure is like any other version of SUSE LINUX. This will let you select your installation method or boot a rescue environment.
5. Select **linux** and press Enter to begin the installation. This will boot the YaST2 software that will help you to select the packages that you may want to install. It will also help you with the hardware configuration and then will install all the packages.

Installation

To configure your system and select your packages, here are the steps you will encounter while installing SLES 8:

1. License Agreement.
2. Language Selection. English is the default.
3. Installation Mode Selection. Select **New installation**.
4. Installation Settings, where you can select what you want to install and where.
5. Package Installation, which is the part of the installation that copies the files from the installation CDs and installs the packages.
6. Prepare for the initial boot. This is the part of the installation that runs the configuration scripts and installs the boot manager.

Once all your packages are copied to the hard disk, reboot the server.

Once the system has rebooted, on the EFI menu select **SuSE SLES** to start the Linux boot process. This entry will also be configured to be the default entry to boot, so if you do not press any key, the system will automatically boot SLES.

You can modify your boot settings EFI boot maintenance menu and the `elilo.conf` file as we discuss later in Example 4-1 on page 162.

The remainder of the installation is completed. The first time SLES 8 boots, it will finish configuring your system. This will boot the graphical interface and then YaST will be launched. If you encounter problems during this first boot, refer to “Linux boot process” on page 162.

7. Root password: The installation program will ask you to enter a password for the root user.
8. Add user: The installation program will ask you to create a user. You need to create this user to continue the installation. This is done for security reasons. You should log on as a normal user, and use the root account as little as possible.
9. Desktop settings: You will be prompted to configure your desktop settings and choose which mode (text or graphical) you want to use the next time the machine is booted.
10. Hardware configuration: If you have some additional hardware, such as network interfaces, you can configure them here.

Installation notes: If you install on a mirrored disk configuration while the mirror is being created, the SCSI driver detection on the installation program may take some time. The driver works fine, but it takes longer to load than in a non-mirrored configuration. You may also notice that the I/O throughput slows down.

4.4.4 Installing Red Hat Enterprise Linux AS

In this section we focus on the installation of the Red Hat Enterprise Linux AS 3.0, although we also did the installation of the RH AS 2.1 on the x455.

Tip: Official installation instructions are typically available on the Web. However, at the time of writing none were published for the x455. Once they are available, they will be accessible at:

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

Select the x455 (8855) and select **OS Installation** document type.

IBM employees can access draft installation instructions from:

<http://nosinfo.raleigh.ibm.com>

Before starting the installation

Before starting the installation, we recommend that you get and apply the latest updates and firmware. You can download them from:

<http://www.pc.ibm.com/qtechinfo/MIGR-53575.html>

Installation procedure

The installation of Red Hat Enterprise Linux AS 3.0 is like any other Red Hat Linux installation, except that you work on a different architecture, so the bootable CD-ROM will not boot directly, as it would do on any IA-32 architecture. These are the steps to follow for a Red Hat Enterprise Linux AS 3.0 installation:

1. Power on the server and insert the first Red Hat CD-ROM in the drive.
Although it is not necessary to insert the CD at this point, it makes the installation process easier because the CD-ROM file system is automatically detected and mapped.
2. Wait for the EFI Firmware Boot Manager menu to appear.
3. Select the first entry for the CD-ROM drive. In the x455, this is the ACPI entry with the strings “PCI” and “ATA(Primary/Master)”. Press Enter (Figure 4-49 on page 158).

Manually starting the CD-ROM: If you did not have the CD-ROM inserted before the EFI Boot Manager menu loaded, do the following:

1. Insert Red Hat Disk 1 into the drive.
2. Select **EFI Shell** and press Enter.
3. Enter the command `map -r` to remap all file systems.
4. Enter `exit` to go back to the Boot Manager menu.
5. Select the CD-ROM entry **Ata(Primary,Master)/CD**.

Installation

To configure your system and select your packages, here are the steps you will encounter while installing Red Hat EL AS 3.0:

1. A message asking you to test your installation CDs appears. You can test the disk or just skip to the next step. We recommend that you test your CDs to avoid problems during the installation.
2. The Red Hat welcome screen appears. The installation procedure is like any other version of Red Hat Linux.
3. Select the following:
 - Language Selection.
 - Keyboard Configuration.
 - Mouse Configuration.
 - Disk Partitioning Setup. (We recommend Automatically, which creates a boot partition that is the EFI System Partition, then the Swap and Root as default.)
 - Network Configuration.
 - Firewall Configuration.
 - Additional Language Support.
 - Time Zone Selection.
 - Set Root Password.
 - Select packages for the installation.
4. Follow the instructions for video card setup and monitor selection.
5. Once the installation is complete, reboot the server.
6. Once the system has rebooted, on the EFI menu select **Red Hat Enterprise Linux AS** to start the Linux boot process. This entry will also be configured to be the default entry to boot, so if you do not press any key, the system will automatically boot Red Hat Enterprise Linux AS.

7. For the first boot you will be prompted to complete some post-installation settings, like the License Agreement, Date and Time, and User Account.

4.4.5 Linux boot process

The Linux boot process on IA-64 systems uses elilo instead of the classic lilo or grub boot loaders. elilo.conf is a file stored in the root director of the EFI System Partition.

There is no longer a need for boot loaders because the EFI system can manage the boot procedure for the different operating systems. The boot loader is helpful, however, if we need to provide boot parameters to the kernel; for example, if we want to recover from some special situations by putting parameters to the boot command line, or we want to boot on initlevel 3 only once and our system is configured to boot on initlevel 5 by default, etc.

If you do not press any key before the delay estimated on the elilo.conf file, Linux will start loading automatically. To reduce the delay or remove it completely, change or remove the corresponding line in elilo.conf.

Example 4-1 shows a typical elilo.conf configuration file. As you can see, the elilo configuration is very similar to the old lilo configuration. One of the main advantages of elilo is that you do not need to execute any command to make your changes active: elilo will automatically take any changes you make on the elilo.conf file.

Example 4-1 An example of elilo.conf

```
# ELILO config file
prompt
timeout = 80
read-only

image = /vmlinuz
  label = linux
  root = 806
  initrd = /initrd
  append = "hda=ide-scsi"
```

Once elilo is loaded, select which Linux kernel image you want to load, and the boot procedure is the same as other Linux startup processes:

1. The Linux kernel image loads on memory.
2. The Setup function is called, enabling basic devices, buses and video.

3. The startup function is called, enabling stacks, decompressing the kernel, and identifying the processor type.
4. The `start_kernel` function is called, starting all memory procedures, and kernel cache, and executes `/sbin/init`.
5. The boot procedure continues.

IA-32 applications

SLES 8 provided the IA-32 runtime environment installed by default. This environment contains `glibc`, `ncurses`, `C++`, and `X11` libraries to guarantee compatibility with IA-32 applications.

If you want to compile any application from code source to IA-32 platform, it is still possible, and you may also compile it to IA-64. This will offer you the possibility to still use older IA-32 applications, if you don't have the source code or you do not want to recompile it to IA-64.

4.4.6 Information about the installed system

Our recently installed Linux system can offer us some information about the x455, such as CPU details, memory, and some other useful information. The following sections outline what the Linux kernel found about the CPU used during our tests. We used the `/proc` virtual file system to obtain this information.

CPU information

Linux recognizes the x455's processors as an Itanium 2, from the IA-64 architecture. During our tests we had two-node configuration, each server with four processors and we will find all the information regarding every CPU, the Example 4-2 shows the two last processors.

Example 4-2 /proc/cpuinfo output

```
[root@x455 root]# cat /proc/cpuinfo
...
processor      : 6
vendor        : GenuineIntel
arch          : IA-64
family        : Itanium 2
model         : 1
revision      : 5
archrev       : 0
features      : branchlong
cpu number    : 0
cpu regs      : 4
cpu MHz       : 1400.000000
```

```
itc MHz      : 1400.000000
BogoMIPS     : 2080.37
```

```
processor    : 7
vendor      : GenuineIntel
arch        : IA-64
family      : Itanium 2
model       : 1
revision    : 5
archrev     : 0
features    : branchlong
cpu number  : 0
cpu regs    : 4
cpu MHz     : 1400.000000
itc MHz     : 1400.000000
BogoMIPS    : 2080.37
```

Memory information

We used 8 GB of DDR memory (each server had 4 GB), and it was fully recognized by Linux. We also set up some swap memory, but it was never used during our tests. Example 4-3 shows the output of the meminfo file on the /proc virtual file system, and more details about the memory.

Example 4-3 /proc/meminfo output

```
[root@x455 root]# cat /proc/meminfo
          total:    used:    free:    shared: buffers:  cached:
Mem:  8469430272 673955840 7795474432      0 15974400 232194048
Swap: 2089172992      0 2089172992
MemTotal:      8270928 kB
MemFree:       7612768 kB
MemShared:           0 kB
Buffers:       15600 kB
Cached:        226752 kB
SwapCached:           0 kB
Active:        276608 kB
ActiveAnon:    113360 kB
ActiveCache:   163248 kB
Inact_dirty:   79104 kB
Inact_laundry:      0 kB
Inact_clean:     0 kB
Inact_target:   71136 kB
HighTotal:     6291456 kB
HighFree:      5894384 kB
```

```
LowTotal:      1979472 kB
LowFree:       1718384 kB
SwapTotal:    2040208 kB
SwapFree:     2040208 kB
HugePages_Total:  0
HugePages_Free:  0
Hugepagesize:  262144 kB
```

Modules used by the default kernel

Linux provides a default kernel that loads almost all the drivers we will need for our x455, and also some drivers we will never use.

As we can see in Example 4-4, it also loads some modules that are not used, identified with `unused`. You can unload any unused module with the `rmmod` command. Enter `man rmmod` for more information about `rmmod` options and parameters.

Example 4-4 /proc/modules output

```
[root@x455 root]# cat /proc/modules
usb-storage      155624  1
soundcore        13632   0 (autoclean)
autofs           30864   0 (autoclean) (unused)
tg3              147432  4
sg               78056   0 (autoclean)
sr_mod           39648   0 (autoclean)
ide-scsi         28480   0
ide-cd           72016   0
cdrom            70244   0 [sr_mod ide-cd]
nls_iso8859-1    5984    2 (autoclean)
nls_cp437        7680    2 (autoclean)
vfat             30336   2 (autoclean)
fat              85708   0 (autoclean) [vfat]
keybdev          5976    0 (unused)
mousedev         14688   1
hid              49544   0 (unused)
input            13536   0 [keybdev mousedev hid]
usb-uhci         67400   0 (unused)
usbcore          180928  1 [usb-storage hid usb-uhci]
ext3             214024  1
jbd              127680  1 [ext3]
mptscsih         92168   3
mptbase          96992   3 [mptscsih]
sd_mod           32352   8
```

```
scsi_mod          229248  6 [usb-storage sg sr_mod ide-scsi
mptscsih sd_mod]
```

Partitions on IA-64 Linux

The EFI system implements a new way to work with partitions. That means that the IA-32 tools to manage partitions on Linux may not work properly. For example, if you run **fdisk** to see the content of a partition or disk, you do not see the real content of the disks. Modifying this partition can destroy the Linux installation.

Example 4-5 shows the **fdisk** output for our test partitions. We had two arrays RAID level 1, each per server, configured on the system: One to run Linux, and the other to run Windows 2003 Server.

Tip: In our case the third partition (sdc) is a USB 256 MB memory key.

Example 4-5 fdisk output of two EFI partitions

```
Disk /dev/sda: 18.1 GB, 18117296128 bytes
255 heads, 63 sectors/track, 2202 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1		1	2203	17692671+	ee	EFI GPT

```
Disk /dev/sdb: 18.1 GB, 18117296128 bytes
255 heads, 63 sectors/track, 2202 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		1	2203	17692671+	ee	EFI GPT

```
Disk /dev/sdc: 262 MB, 262144000 bytes
16 heads, 32 sectors/track, 1000 cylinders
Units = cylinders of 512 * 512 = 262144 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdc1	*	1	999	255728	6	FAT16

The two disks included more than one partition. In fact, there were at least three partitions on each: One for EFI, one for the operating system boot images, and another for the data. As you can see in the example, **fdisk** does not see any of these partitions, but only a big partition of type EFI GPT.

Instead of using **fdisk**, use **parted**.

The parted utility is included on the Linux (SUSE LINUX and Red Hat) base system. As shown in Example 4-6, the parted output does not show the same information in the partition table, but recognizes the real partitions on the system. The parted utility can identify every partition on the partition table.

Example 4-6 GNU parted output (some lines removed)

```
[root@x455 root]# parted
GNU Parted 1.6.3
...
Using /dev/sda
(parted) print
Disk geometry for /dev/sda: 0.000-17278.000 megabytes
Disk label type: gpt
Minor    Start      End        Filesystem  Name                Flags
1         0.031      172.573    fat16       EFI system partition boot
2         172.573    298.081                    Microsoft reserved partition
3         298.081    10299.484 ntfs        Basic data partition
4         10299.485  10353.750 fat32                               boot
5         10353.750  11381.281 linux-swap
6         11381.281  17271.938 reiserfs
(parted) select /dev/sdb
Using /dev/sdb
(parted) print
Disk geometry for /dev/sdb: 0.000-17278.000 megabytes
Disk label type: gpt
Minor    Start      End        Filesystem  Name                Flags
1         0.031      125.507                    Microsoft reserved partition
2         125.508    3125.507 ntfs        Basic data partition
3         3125.508    3225.507 fat16                               boot
4         3225.508    5217.944 linux-swap
5         5217.944    17277.983 ext3
(parted) select /dev/sdc
Using /dev/sdc
(parted) print
Disk geometry for /dev/sdc: 0.000-250.000 megabytes
Disk label type: msdos
Minor    Start      End        Type        Filesystem  Flags
1         0.016      249.750    primary     fat16       boot
```

In our case, we had a two-node configuration, and per server we made an array RAID level 1, and installed the following operating systems:

- ▶ Windows Server 2003, Datacenter Edition in disk sda and sdb, and it has the following partitions:
 - EFI System Partition (FAT16) assigned as sda1
 - Two Microsoft Reserved Partitions; one assigned as sda2 and the another as sdb1
See “Microsoft Reserved Partition” on page 137 for more information.
 - Two Basic data partitions (NTFS); one assigned as sda3 and the another as sdb2
- ▶ SUSE LINUX SLES 8 in disk sda
 - ESP (FAT32) named boot assigned as sda4
 - Swap partition assigned as sda5
 - Root or data partition (ReiserFS) assigned as sda6
See Figure 4-50 on page 169.

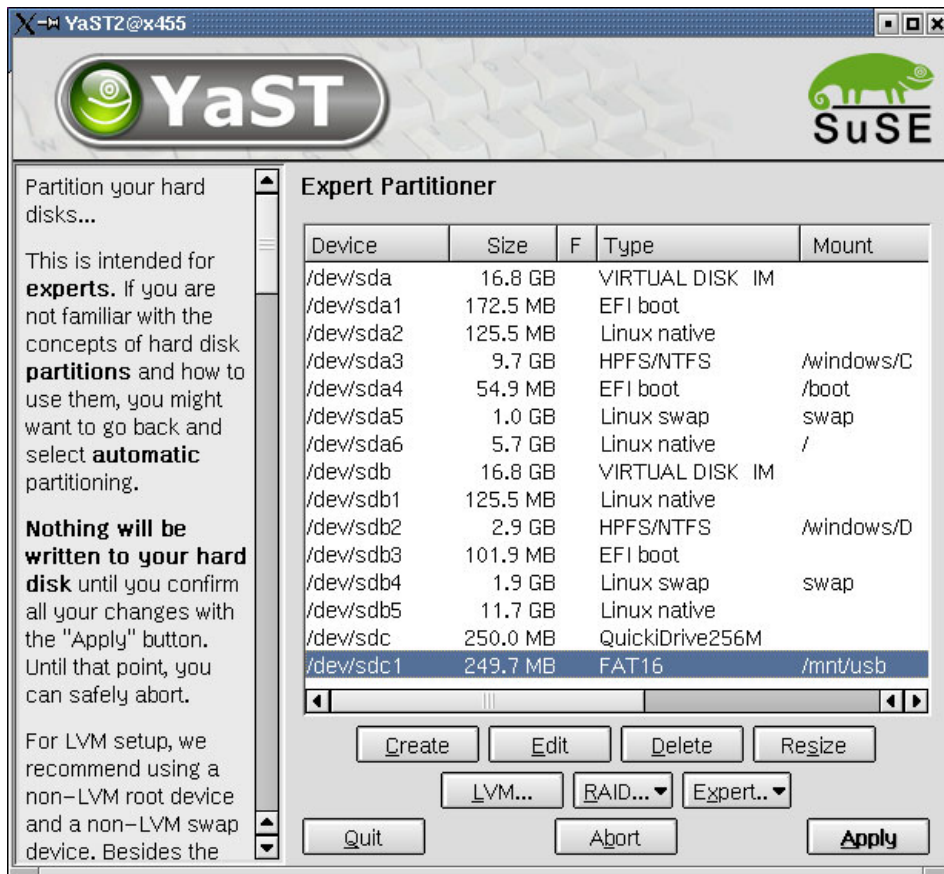


Figure 4-50 Expert Partitioner, SLES 8

- ▶ Red Hat EL AS 3.0 in disk sdb
 - ESP (FAT16) named boot assigned as sdb3
 - Swap partition assigned as sdb4
 - Root or data partition (ext3) assigned as sdb5
- See Figure 4-51 on page 170.

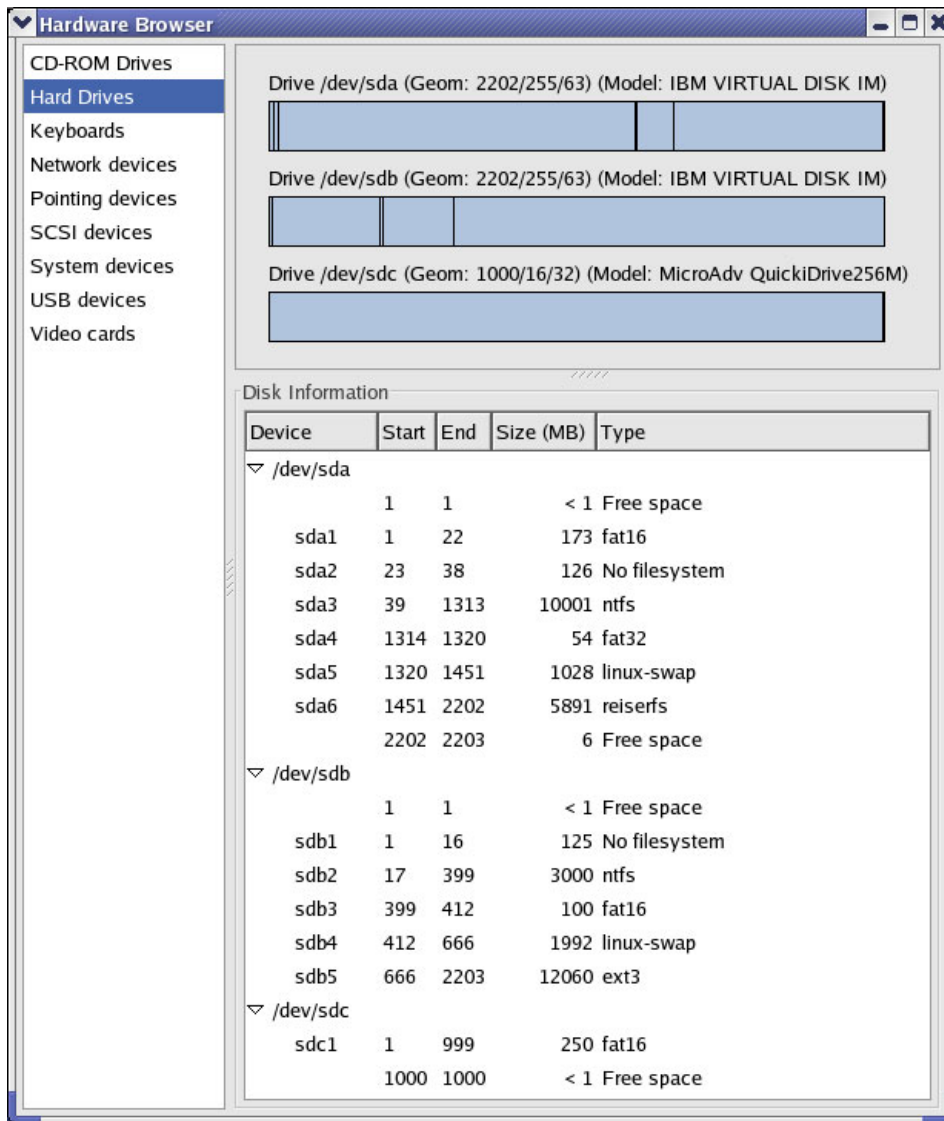


Figure 4-51 Partitions in Hardware Browser, Red Hat EL AS 3.0

Tip: In our case the third partition (sdc) is a USB 256 MB memory key.

The FAT16 and FAT32 partitions are the EFI System Partitions (one per operating system) where the kernel images and other files needed to boot Linux

are stored. The EFI System Partition is also available as the /boot partition mounted directory once you log into Linux.

For more information about the parted utility, use `man parted` or go to:

<http://www.gnu.org/software/parted>

4.4.7 Using the serial port for the Linux console

You can redirect all kernel console messages to the serial port. This can be done by entering the following command at the elilo prompt or adding it to the elilo.conf file:

```
append="console=ttyS0,115200"
```

You can also redirect not only kernel messages, but any other system message from the default messages console (which is virtual terminal number 10 on SUSE LINUX, and /var/log/messages on other Linux distributions) to the serial port. To do so, you must use the syslog daemon. This daemon controls the flow of system messages and redirects them wherever we want them to be redirected. To redirect all syslog messages to the serial port, edit the /etc/syslog.conf file and add the following string to the line that you want to be redirected to the serial port:

```
/dev/ttyS0
```

Example 4-7 shows an example of this.

Example 4-7 Redirection of kernel messages to the serial port with syslogd daemon

```
# This is only a part of the /etc/syslog.conf file
# Your file may be different. Here we will send all the kernel
# warning messages to the serial port
```

```
kern.warn; /dev/ttyS0
```

To activate the changes you made in the syslog daemon, enter the command:

```
/etc/init.d/syslog restart
```

Refer to the syslog daemon man page for more information about the configuration.

4.4.8 RXE-100 Expansion Enclosure

The RXE-100 Expansion Enclosure is now fully supported under Linux. The x455 will recognize any device attached to any of the 12 RXE-100 slots. You do not need to disconnect the RXE-100 before installing Linux.

Tip: In the case of SUSE LINUX Enterprise Server 8 you can easily view all your recognized PCI devices with the YaST tool. To invoke YaST, simply enter **yast** at the command prompt, then go to the Hardware and PCI information.

As you can see in Figure 4-52 on page 173, all the PCI slots are recognized, including the 12 external PCI slots.

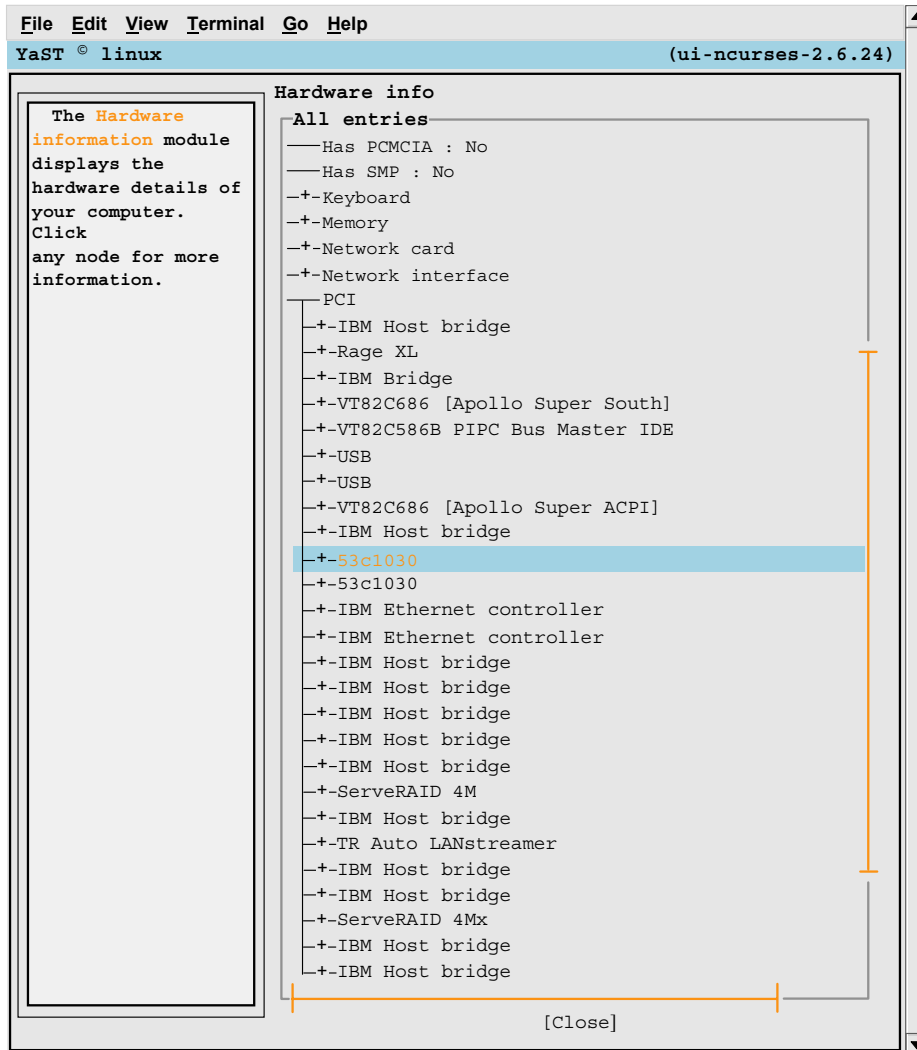


Figure 4-52 PCI information on YaST

4.4.9 Upgrading drivers

Any update to the Linux installation should be done via the official support service provided from IBM and SUSE LINUX. To visit the xSeries support page, go to:

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

The most significant drivers and modules that you should maintain on the x455 are:

- ▶ **Broadcom Gigabit Ethernet**

The Broadcom Gigabit Ethernet adapter is supported on Linux with the “bcm5700” module. This module supports all the 5700 series of Gigabit adapters from Broadcom. This module comes with SUSE LINUX SLES 8 and Red Hat EL AS 3.0, and both Broadcom ports are detected during the installation.

- ▶ **LSI SCSI chipset**

The default kernel that comes with SUSE LINUX SLES 8 and Red Hat EL AS 3.0 has the LSI chipset support precompiled. This will automatically support any SCSI devices connected to your LSI chipset.

- ▶ **ATI Rage XL video controller**

The video controller used in the x455 is an ATI Rage XL device. This device is supported by XFree86, with the “ati” driver. This driver supports almost all the ATI graphics controllers. ATI controllers are supported on both major versions of XFree86, the 3.3.6 and the 4.2.1. To follow the status of any supported ATI graphic device supported by the XFree86 project, see:

<http://www.xfree86.org/current/Status6.html#6>

For generic information about the status of any driver, see:

<http://www.xfree86.org/current/Status.html>

ATI also has a support page for Linux, where you may find all the information you may need for finding any new drivers or simply to know more about ATI's support of Linux:

<http://mirror.ati.com/support/faq/linux.html>

- ▶ **XA-64 chipset**

At the time of this writing, there was no official kernel update to support the new 64-bit chipset. The x455 will work in 64-bit native mode without any kernel patches. If IBM, SUSE LINUX, or Red Hat release any specific patch for the XA-64 chipset, you will find more information about how to upgrade your current kernel on the xSeries support page:

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



Management

As systems grow in power and complexity, systems management becomes a critical component of a complete system solution. This chapter discusses several techniques and tools that can be used to manage the x455. The topics covered are:

- ▶ “IBM Director” on page 176
- ▶ “The Remote Supervisor Adapter” on page 177
- ▶ “Management using the Remote Supervisor Adapter” on page 187
- ▶ “Windows System Resource Manager” on page 196

Important: The available date for the IBM Director agent 64-bit (Windows and Linux) and SSM plug-in is first quarter 2004. Consequently, this chapter concentrates on out-of-band alerting and management through the Remote Supervisor Adapter.

5.1 IBM Director

IBM Director is designed to manage a complex environment that contains numerous servers, desktop computers, workstations, notebook computers, and various types of SNMP-based devices.

Today it is common to find a mixture of system types, wiring structures, and network protocols used within a single corporate IT infrastructure. IBM Director supports Intel-based hardware from many manufacturers, as well as multiple network connection types and protocols, enabling you to manage a heterogeneous environment. With IBM Director, IT administrators can view and track the hardware configuration of remote systems in detail and monitor the usage and performance of critical components, such as processors, disks, and memory.

The hardware in an IBM Director environment can be divided into the following groups:

- ▶ Management servers: One or more servers on which IBM Director Server is installed
- ▶ Managed systems: Servers, workstations, desktop computers, and notebook computers that are managed by IBM Director
- ▶ Management consoles: Servers, workstations, desktop computers, and notebook computers from which you communicate with one or more IBM Director Servers
- ▶ SNMP devices: Network devices, printers, or computers that have SNMP agents installed or embedded

IBM Director software has three components:

- ▶ IBM Director Server
- ▶ IBM Director Agent
- ▶ IBM Director Console

The x455 needs only the IBM Director agent 64-bit to be managed and the features available for this agent are listed in “IBM Director support” on page 89.

Note: This publication assumes that you are familiar with IBM Director. For information on how to configure and use IBM Director, see the publication *Implementing Systems Management Solutions using IBM Director*, SG24-6188.

5.1.1 Scalable Systems Manager

Scalable Systems Manager is an extension of IBM Director that communicates solely out-of-band with the service processor in supported xSeries servers, like x455, to create and manage static hardware partitions. These static partitions are referred to as scalable partitions throughout the Scalable Systems Manager interface and documentation.

The console component of Scalable Systems Manager is an extension to IBM Director that provides the following console tasks:

- ▶ Scalable Systems Configuration

With this task you can view and manage scalable objects in a graphical user interface. It is possible to create and delete scalable objects, add and delete unassigned scalable nodes to and from IBM Director, insert and remove scalable nodes to and from a scalable partition, view information about and configure the remote I/O enclosures attached to a scalable node, and perform other operations on scalable partitions, such as power on, shutdown, and power off.

- ▶ Scalable Systems Discovery

Create a manageable scalable system of the unassigned scalable and also update it to add unassigned scalable nodes.

- ▶ Scalable Systems Validation

This task is used when you need to validate the cabling of the scalable nodes in a manageable scalable system.

Note: The available date for the IBM Director agent 64-bit (Windows and Linux) and SSM plug-in is first quarter 2004.

5.2 The Remote Supervisor Adapter

The x455 is equipped with a Remote Supervisor Adapter integrated into the system. See Figure 5-1 on page 178 for the location of the Remote Supervisor Adapter.




Figure 5-1 Location of Remote Supervisor Adapter in the x455

The features of the Remote Supervisor Adapter include the following:

- ▶ Continuous health monitoring and control
- ▶ Advanced Predictive Failure Analysis (PFA)
- ▶ Configurable notification and alerts
- ▶ Event logs that are time stamped, saved in nonvolatile memory, and can be attached to e-mail alerts
- ▶ Remote graphical console redirection
- ▶ Remote access either via Ethernet, serial, or Advanced System Management (ASM) interconnect connectivity
- ▶ Point-to-point protocol (PPP) support
- ▶ Simple Network Management Protocol (SNMP) support
- ▶ Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP) support
- ▶ Remote power control
- ▶ Remote firmware update and access to critical server settings

The functions and information made available by the Remote Supervisor Adapter can typically be accessed either in-band via IBM Director (via IBM Director Agent installed on the server), or via one of the out-of-band interfaces.

In-band means by means of the operating system and device drivers installed on the server. Out-of-band means bypassing the operating system and connecting to the adapter directly.

There are three physical connections to the adapter for out-of-band management (that is, management without going through the operating system on the server), as shown in Figure 5-1 on page 178:

- ▶ Ethernet port, supporting TCP/IP connections
- ▶ Serial (COM) port, supporting null modem or PPP connections
- ▶ RS-485 ASM interconnect port, for connection to other service processors

The adapter supports three user interfaces for out-of-band management:

- ▶ Web browser
- ▶ Telnet
- ▶ ANSI terminal

The recommended and usually the easiest ways to perform out-of-band management are:

- ▶ Using a Web browser to connect to the adapter via its Ethernet port
- ▶ Connecting to the adapter from a Remote Supervisor Adapter installed in another server, using the RS-485 ASM interconnect port

5.2.1 Connecting via a Web browser

To use a Web browser to access the Remote Supervisor Adapter remotely, ensure that your browser meets the following requirements:

- ▶ Microsoft Internet Explorer 4.0 (with Service Pack 1) or later
- ▶ Netscape Navigator 4.72 or later (Version 6.x is not supported)
- ▶ Java-enabled Web browser
- ▶ Support for JavaScript 1.2 or later
- ▶ Minimum display resolution of 800 x 600 pixels and 256 colors

Note: Neither the Web interface nor the text-based interface support double-byte character set (DBCS) languages.

To initially configure the adapter, you will need to know either the host name or the IP address of the adapter. Since both of these items are initially unknown, you can use a Dynamic configuration, or see “Configuring a static IP address” on page 181.

Dynamic configuration

By default, TCP/IP on the Remote Supervisor Adapter is as follows:

- ▶ DHCP enabled
- ▶ Host name set to ASMA followed by the adapter’s burned-in MAC address

Follow these steps to determine the host name or the IP address of the adapter from the System Log:

Tip: The following assumes you have a DHCP server configured.

1. Connect an active Ethernet connection to the Ethernet port of the adapter.
2. Power on the server.
3. From the EFI Firmware Boot Manager menu (Figure 4-2 on page 94), select **Configuration/Setup**.
4. Select **System Event and Error Log** and press Enter.
5. Scroll through the event logs until you see an entry similar to the following:

```
ENET[1] DHCP-HstNme=455rsa, DN=itso.ra1.ibm.com, IP@=192.168.1.151,  
GW@=192.168.1.1, NMsK=255.255.255.0, DNS1@=192.168.1.201,  
DNS2@=192.168.1.202
```
6. From this information, you can connect to the adapter via a Web browser.
7. When prompted, enter a valid user ID and password. By default, only one user ID is configured:
 - User ID: USERID (all caps)
 - Password: PASSWORD (all caps and zero, not the letter O)

Important: We strongly recommend that you disable this default user ID and create new more secure user IDs before you put your server into production. Access to the Remote Supervisor Adapter should be carefully restricted.

8. Once you are connected, the main menu should appear.
9. You should change the adapter name (via System Settings) and host name (via Network Interfaces). We recommend that you make these two names the same to eliminate confusion.
10. We also recommend that you configure a static IP address for your Remote Supervisor Adapter. This will ensure that the adapter is still accessible from the network even when the DNS and DHCP servers are unavailable. This is described in the next section.

Configuring a static IP address

If you do not have a DHCP server on the network, you will need to configure the adapter to use a static IP address.

The easiest way is to connect to the x455's adapter using the ASM interconnect bus. You can then change the TCP/IP setting via the Network Interfaces menu using a Web browser. See "Connecting via the ASM interconnect" on page 184.

However, if another Remote Supervisor Adapter is not available, or close enough to be connected via the ASM interconnect, then the next easiest way is to use the EFI-based flash/configuration utility FLASH2.EFI.

Download this utility as part of the latest x455 Remote Supervisor Adapter firmware update from:

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

Once you have downloaded it, you will need to copy it to a USB memory key (see "EFI tips and tricks" on page 102) or a CD-R, or copy it directly to the EFI system partition. Only the FLASH2.EFI file is needed to configure the adapter.

To configure a static IP address using FLASH2, perform the following steps:

1. Boot the x455 into the EFI shell.
2. Place the flash2 application on media that can be read by the x455. This could be a CD, memory key, or the application can be copied to a partition on a hard disk drive (assuming that a partition has been created).
3. From the EFI shell prompt, navigate to where you have placed the flash2 application (enter `fs2:` or `fs3:`, for example, to select the file system that has been mapped to the system partition).
4. Enter the command `flash2` and press Enter.
5. Select **Configuration Settings**.
6. Select **Ethernet Settings**. The Ethernet Settings window opens, as shown in Figure 5-2 on page 182.

```

IBM Remote Supervisor utility

Ethernet Settings

Network Interface:      ENABLED_
IP Address:            192.168.  1.  1
Subnet Mask:          255.255.255.  0
Gateway:              0.  0.  0.  0
Line Type:            Ethernet__
Host Name:            X455.itso.ibm.com_____
Data Rate:            AUTO
Duplex:              AUTO
MTU size:            1500____
User Defined MAC address: 00-00-00-00-00-00
DHCP Setting:        DISABLED
To allow ASM to use static IP address, subnet mask, and gateway values,
you must set the DHCP setting to "DISABLED".

<F1> Help      <F2> Refresh   <-> Next Value   <↑><↓> Move
<Esc> Exit     <F6> Apply     <←> Previous Value <F9> Restart

```

Figure 5-2 RSA Ethernet settings

7. Enter the values specific to your organization for the IP address, subnet mask, and gateway.
8. Enter your host name. To eliminate confusion, consider making this field the same value as the system name under General.
9. Change the DHCP Setting field to Disabled. The default value is Enabled.
10. Press F6 to apply the changes.
11. Press Esc. After the Warning window opens, press Enter to restart the Remote Supervisor Adapter.
12. Exit the utility and return to the EFI Boot Manager menu.
13. You should now be able to connect to the adapter using the static IP address you supplied.
14. From another machine, that is also connected to the same network, connect to the adapter via the browser. In our example, the URL would be the IP address:

```
http://192.168.1.1
```
15. When prompted, enter a valid user ID and password. By default, only one user ID is configured:
 - User ID: USERID (all caps)
 - Password: PASSWORD (all caps and zero, not the letter O)

Important: We strongly recommend that you disable this default user ID and create new more secure user IDs before you put your server into production. Access to the Remote Supervisor Adapter should be carefully restricted.

16. Once you are connected, the main menu should appear. See Figure 5-3.

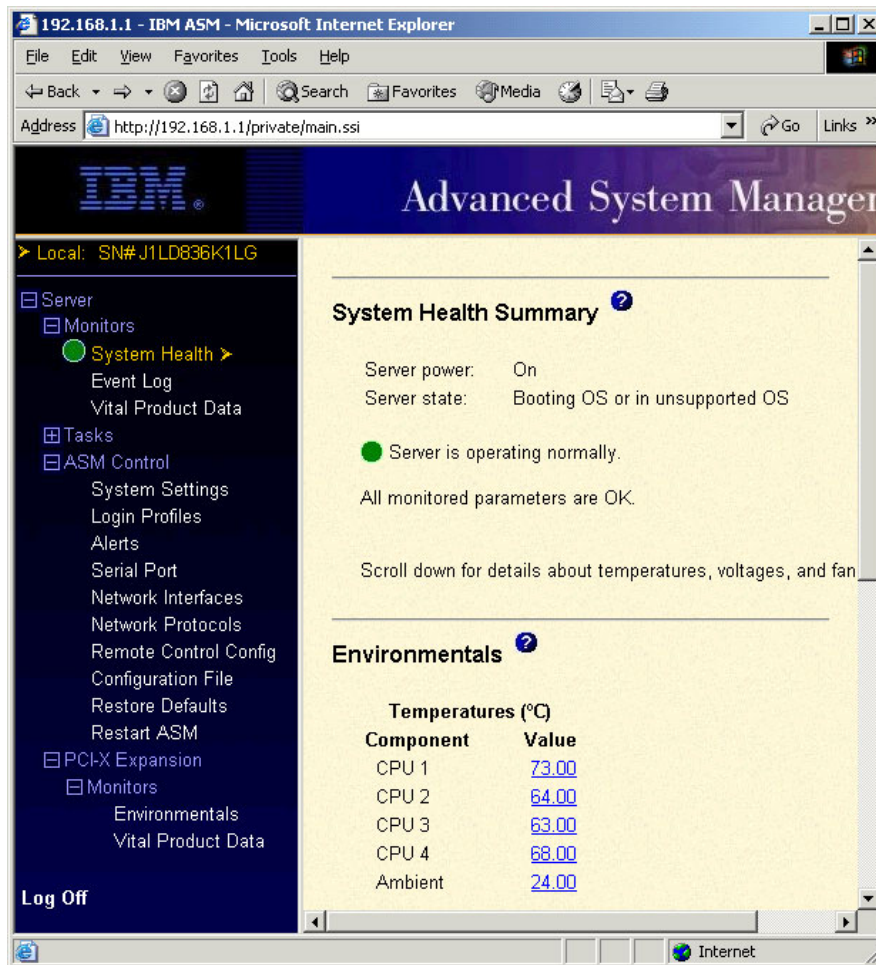


Figure 5-3 Remote Supervisor Adapter Web interface

Tip: For detailed information about the Ethernet settings or about the other connectivity options using the serial port (null modem or PPP), see the *Remote Supervisor Adapter User's Guide*.

5.2.2 Connecting via the ASM interconnect

In addition to the Ethernet interface, the ASM interconnect is often used to manage the Remote Supervisor Adapter when you have other service processors in other xSeries servers.

To connect the Remote Supervisor Adapter in the x455 to an ASM interconnect bus, you will need:


- ▶ Advanced System Management Interconnect Cable Kit, part 03L9309
- ▶ Standard Cat5 Ethernet cable (not crossover)






Follow the instructions that ship with the cable kit, or review the publication *Implementing Systems Management Solutions using IBM Director*, SG24-6188.

Once you have connected the x455 Remote Supervisor Adapter to other service processors including another Remote Supervisor Adapter, do the following:

1. From your Web browser, connect to the other Remote Supervisor Adapter (for example, one in an x445).
2. Select **Access Remote ASM** from the main menu (Figure 5-3 on page 183).
3. A list of the other service processors connected on the ASM interconnect bus will be displayed.

Tip: It can take up to about a minute (the online help says 45 seconds) before a new service processor connected to the interconnect bus appears on this page.

Remote ASM Access 

System Health	ASM Name	ASM Interconnect Connection	Direct LAN Connection
	x440top	login	9.24.104.124
	Oslo-RSA	login	(no LAN connection)
	Lenzburg	login	9.24.104.105
	SN# J1LD836K1LG	login	192.168.1.1
	chicago-rsa	login	9.24.104.254

Click on "login" to establish a session with a specified ASM.
Click on the IP address to start a direct LAN session in a new browser window.

Figure 5-4 Connecting remotely to the Remote Supervisor Adapter in the x455

4. Next to the x455 (IP address 192.168.1.1 in our example), click **Login**. If you click the IP address in that row, it connects to the adapter via its Ethernet port.
5. When prompted, enter a valid user ID and password. By default, only one user ID is configured:
 - User ID: USERID (all caps)
 - Password: PASSWORD (all caps and zero, not the letter O)

Tip: We strongly recommend that you disable this default user ID and create new more secure user IDs before you put your server into production. Access to the Remote Supervisor Adapter should be carefully restricted.

6. Once you are connected, the main menu of the x455 Remote Supervisor Adapter should appear.

5.2.3 Installing the device driver

You should install the ASM device driver once the operating system is installed. The device driver enables communication between the adapter and the installed operating system (Windows or Linux). In particular, the driver provides:

- ▶ The ability to gracefully shut down the server
- ▶ Indications as to whether the operating system is booted
- ▶ The ability to determine if the operating system is hung and automatically restart the system

The device driver is on the Resource CD-ROM that is shipped with the server. The latest version of the driver is also available from:

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

Installation procedures are documented in the following sections:

- ▶ Windows: “Installing additional drivers” on page 144
- ▶ Linux: “Upgrading drivers” on page 173

5.2.4 Configuring the remote control password

If you want to use the remote control features of the Remote Supervisor Adapter Web interface, you must configure the remote control password using the Configuration/Setup utility program.

To set up a remote control password, complete the following steps:

1. Start the server.
2. From the EFI Firmware Boot Manager startup options, select **Configuration/Setup** and press Enter.
3. Select **System Security**.
4. Select **Remote Control Security**.
5. In the Enter New Remote Control Password field, enter the Remote Control password that you wish to use each time you log into the remote control features of the Remote Supervisor Adapter Web Interface. See Figure 5-5 on page 187.

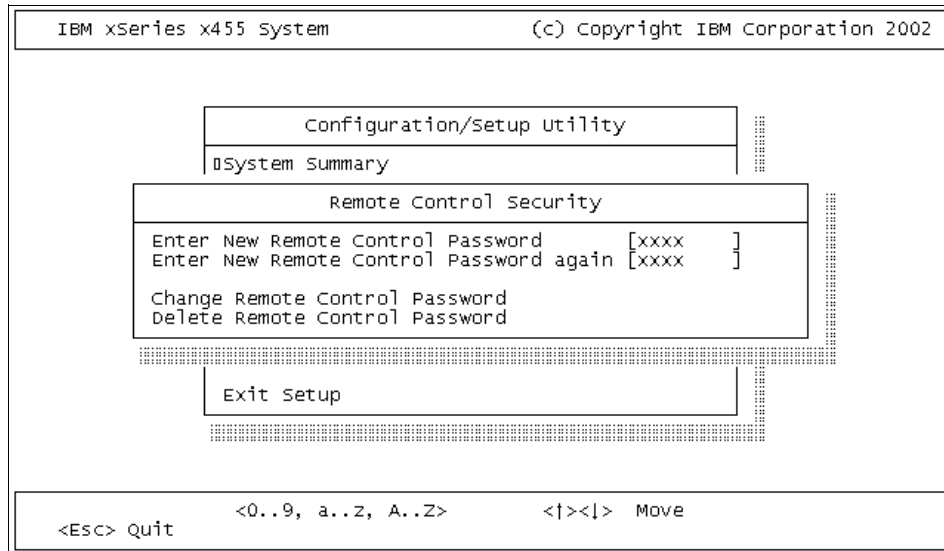


Figure 5-5 Configuring the remote control password

6. Select **Change Control Password** and then press Enter.

5.3 Management using the Remote Supervisor Adapter

Once you have configured the Remote Supervisor Adapter as described in “The Remote Supervisor Adapter” on page 177, you can then use the adapters to manage x455 by forwarding events to IBM Director. IBM Director can then be configured to act upon those events in a variety of ways (for example, by sending e-mail or a pager alert).

The Remote Supervisor Adapter supports the following methods of forwarding alerts to IBM Director:

- ▶ Simple Network Management Protocol (SNMP), as described in “Configuring SNMP” on page 190
- ▶ Communication directly with IBM Director, either via Ethernet or via modem, as described in “Sending alerts directly to IBM Director” on page 191

Tip: While all methods are supported, we recommended that only one method be utilized. Otherwise, you may have the situation where a single event on the x455 triggers multiple events in IBM Director.

5.3.1 Configuring which alerts to monitor

Before configuring the method of sending the alert to IBM Director, you should first define which events the Remote Supervisor Adapter is to send.

To enable alerting on the Remote Supervisor Adapter, do the following:

1. From the Web browser main menu, click **Alerts** in the navigation bar.
2. Scroll to the Monitored Alerts section. There are three types of events listed, as follows:
 - Critical Alerts: These are hard disk drive, multiple fan failure, power failure, tamper, temperature, voltage, and VRM failure
 - Warning Alerts: These are single fan failure, temperature, voltage, and redundant power supply
 - System Alerts: These are boot failure, loader timeout, O/S timeout, PFA, POST timeout, power off, and power on
3. Select the events you wish to monitor. We recommend you enable all events. See Figure 5-6 on page 189.

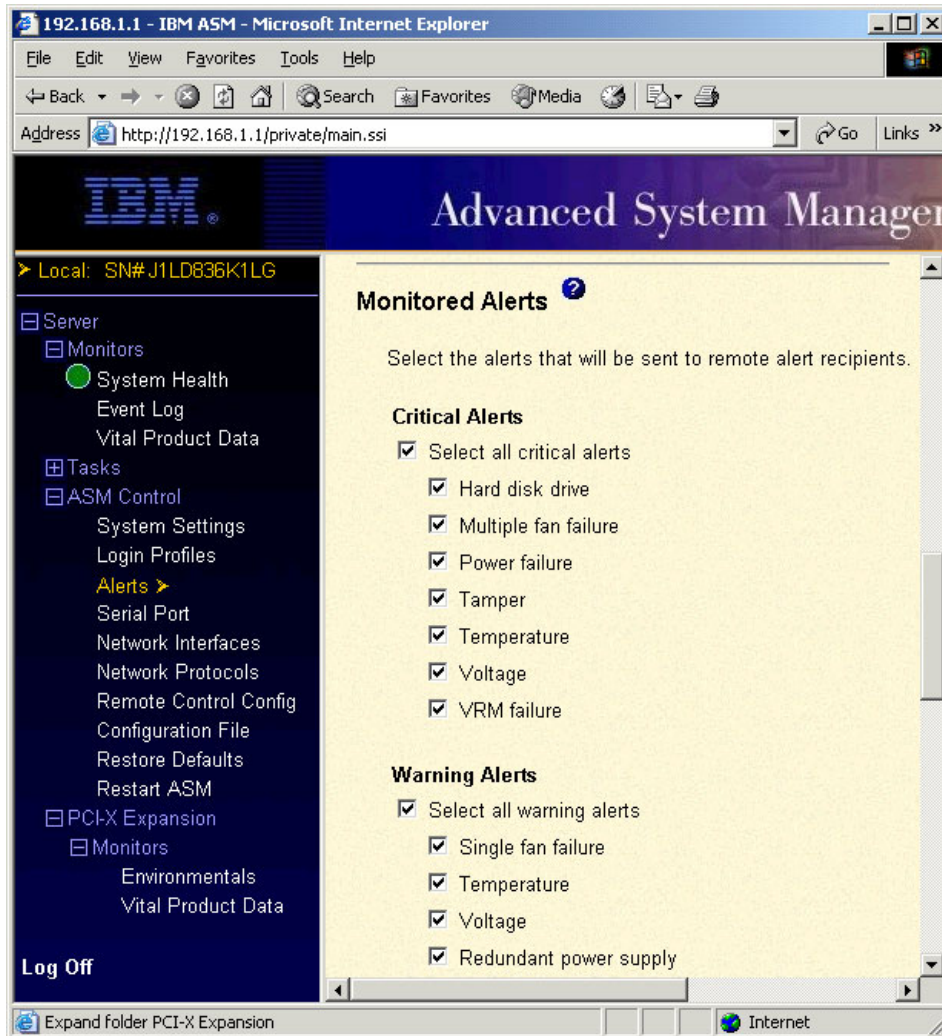


Figure 5-6 Configuring monitored alerts

4. Click **Save**.

Tip: Below the Monitored Alerts section is the Monitored Local Alerts section. These options are applicable only when the IBM Director Agent is installed on the x455. When the agent is installed and those options are selected, the events are sent to IBM Director agent and appear as sourced from the x455.

5.3.2 Configuring SNMP

The Remote Supervisor Adapter supports the sending of alerts to IBM Director using the Simple Network Management Protocol (SNMP).


Before alerts can be sent, you need to configure the adapter for SNMP so that IBM Director Server can discover and manage it.

Note: We assume that you have already set up a system running IBM Director Server. If you need help configuring IBM Director, please see the documentation that comes with IBM Director or the IBM Redbook *Implementing Systems Management Solutions using IBM Director*, SG24-6188.

Configure SNMP as follows:

1. If you have not already done so, configure the adapter so that it can be accessed from a Web browser, either through a direct Ethernet connection or via an ASM interconnect. See “Connecting via a Web browser” on page 179 or “Connecting via the ASM interconnect” on page 184.
2. Connect to the adapter using the Web browser. From the home page, click **Network Protocols**.
3. In the SNMP section (see Figure 5-7 on page 191) enter the following information:

SNMP agent	Enabled (default is Disabled)
SNMP traps	Enabled (default)
Community name	Set your SNMP community name (default is public for IBM Director)
Host Name or IP address	Host name or IP address of the system running IBM Director Server

Simple Network Management Protocol (SNMP) 

SNMP agent

SNMP traps

Community Name	Host Name or IP Address
<input type="text" value="public"/>	1. <input type="text" value="192.168.1.20"/>
	2. <input type="text"/>
	3. <input type="text"/>
<input type="text"/>	1. <input type="text"/>
	2. <input type="text"/>
	3. <input type="text"/>
<input type="text"/>	1. <input type="text"/>
	2. <input type="text"/>
	3. <input type="text"/>

Figure 5-7 Remote Supervisor Adapter - SNMP configuration

4. Once all values are entered, click **Save**. You will be prompted that you need to restart the Remote Supervisor Adapter to bring the changes into effect.
5. In the left-hand navigation pane, click **Restart ASM**.
6. Perform a discovery at the IBM Director Console. IBM Director should detect the Remote Supervisor Adapter in the x455 as an SNMP device.
By default, the name listed in IBM Director is Remote Supervisor Adapter <IP address>, where <IP> Address is the address of the Remote Supervisor Adapter you configured. You can choose to make this name more descriptive for your environment by right-clicking and choosing **Rename**.
7. Proceed to “Creating a test event action plan in IBM Director” on page 193 for instructions on how to verify the functionality of the Remote Supervisor Adapter.

5.3.3 Sending alerts directly to IBM Director

The Remote Supervisor Adapter can be configured to send alerts directly to IBM Director Server via one of these ports on the adapter:

- ▶ Ethernet port
- ▶ COM port and an attached modem

- ▶ ASM interconnect connection to another service processor that has an Ethernet connection or modem attached

The ASM interconnect bus transparently handles the transfer of all outgoing alerts to Ethernet connections or modems on any of the service processors connected together. In this way, for example, as long as one of the connected service processors has an Ethernet connection or modem, then all alerts from the Remote Supervisor Adapter in the x455 will be sent without any extra configuration required.

Important: Sending alerts directly to IBM Director, while convenient, may cause difficulties processing incoming events at the IBM Director Server. This is because the events that are sent directly from the x455 Remote Supervisor Adapter are reported as being sourced from the IBM Director Server itself, instead of being sourced from the x455.

Once the Director Agent is available, alerts can be sent through it and event processing will be easier.

To configure the Remote Supervisor Adapter to send alerts to IBM Director Server via Ethernet, perform the following steps:

1. If you have not already done so, configure the Ethernet adapter settings on the adapter. See “Connecting via a Web browser” on page 179 for details on how to complete this configuration.
2. Connect to the Remote Supervisor Adapter using a Web browser. From the home page, click **Alerts**.
3. The upper portion of this Web page lists those recipients configured for remote alerts. By default, this list is empty. Click an empty link (labeled ~not used~) to configure a remote recipient.
4. To configure IBM Director as a remote recipient, you will need to enter the following values. All fields not listed can be left in their default (blank) state:

Status	Enabled (default is Disabled)
Name	Enter a descriptive name
Notification Method	IBM Director over LAN or IBM Director over modem
Number	IP Address of the system running IBM Director Server

5. Click **Save** to save your configuration. You should see your recipient listed in the list.

Tip: At time of writing, clicking the **Generate Test Alert** button did not properly transmit an alert to IBM Director when the **IBM Director over LAN** method was selected.

5.3.4 Creating a test event action plan in IBM Director

Once you have configured the Remote Supervisor Adapter to communicate with IBM Director, we recommend that you test the functionality to ensure proper communication.

For more information on alerts and event management in IBM Director, see the IBM Redbook *Implementing Systems Management Solutions using IBM Director*, SG24-6188.

To verify the functionality of the Remote Supervisor Adapter in the x455, perform the following steps:

1. Launch the IBM Director Console, either directly from the IBM Director Server system or from a workstation with the IBM Director Console installed.
2. Click the **Event Action Plan Builder** icon located in the IBM Director Console.
3. The Event Action Plan Builder window will appear. Right-click **Event Action Plan** in the left-hand pane and select **New** to create a new plan.
4. Provide a descriptive name for your plan and click **OK**. For our testing, we entered RSA Test.

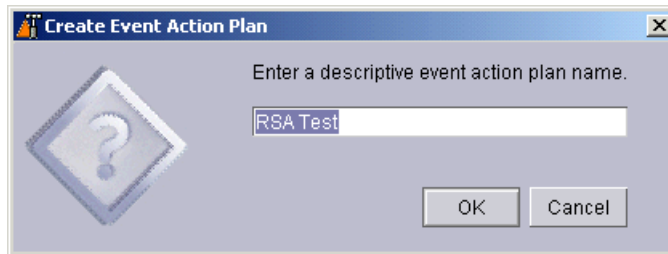


Figure 5-8 New event action plan

5. Once you click **OK**, your new event action plan will appear in the left-hand pane under Event Action Plans.
6. In the middle pane, click the blue bullet next to **Simple Event Filter**. A menu tree will drop down showing all pre-defined event filters created during installation of IBM Director. Click and drag **All Events** to the event action plan you just created. This will add the filter to the action plan. The All Events filter will cause any event generated by the system to be reported to IBM Director, and is good for testing. For production, you may wish to limit the events sent to IBM Director.

7. In the right-hand pane, right-click **Send an Event Message to a Console User** and choose **Customize**. This action will pop up the Message Console on the console currently being used when the alert is triggered.
8. A new window will appear. For the message, type &Text. This will show the error message in the pop-up window. For User(s), type *. This will cause the message to appear on all current Console users. See Figure 5-9 for more details.

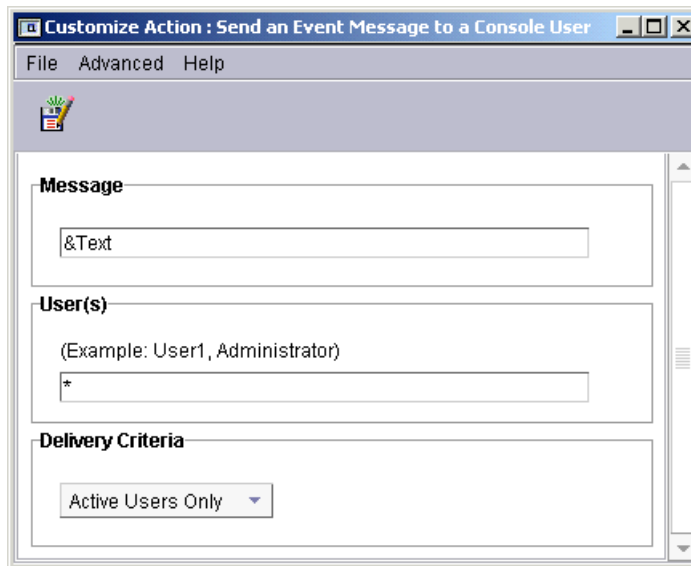


Figure 5-9 Customize action configuration window

9. Click the **Save** icon in the upper left-hand corner of the window (it looks like a diskette) to save the changes.
10. Your new customized action should appear in the right-hand pane. Click and hold the action and drag the icon to the filter you created.
11. Close the Event Action Plan Builder.
12. If you chose to configure the Remote Supervisor Adapter using Alert Forwarding, continue with step 13. If you chose to configure your adapter using SNMP, jump to step 14.
13. In the IBM Director Console, click and drag your test event to your IBM Director Server system in the middle pane. See Figure 5-10 on page 195 for how your configuration should look.

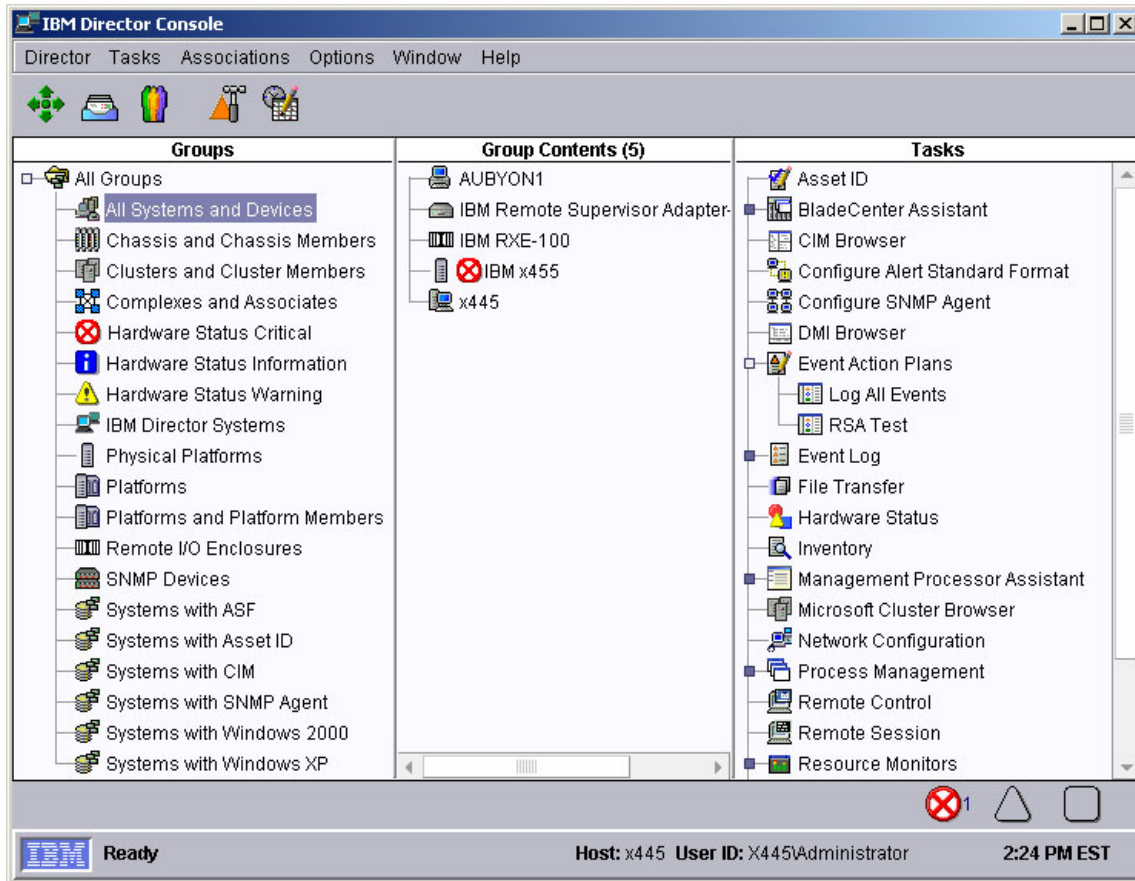


Figure 5-10 Completed configuration

14. In the IBM Director Console, click and drag your completed test configuration to the Remote Supervisor Adapter for the x455 located in the middle pane.
15. In either case, a pop-up window will appear notifying you that the event action plan has been added (Figure 5-11 on page 196).




Figure 5-11 Event action plan added

16. To complete testing, we recommend removing one of the fans from the top of the x455. Within approximately 10 seconds, an alert should appear on the Director Console, similar to the one in Figure 5-12.




Figure 5-12 Message alert

Once functionality has been confirmed, you can now build event action plans to suit your environment.

5.4 Windows System Resource Manager

Microsoft Windows System Resource Manager (WSRM) is a utility used for managing CPU and memory resource allocation among multiple applications that run on the same Windows Server 2003 operating system image. It is a key factor in Microsoft's server consolidation strategy. The server component is available for use on:

- ▶ Microsoft Windows Server 2003, Enterprise Edition
- ▶ Microsoft Windows Server 2003, Datacenter Edition

WSRM is supplied on a separate CD as part of the operating system package and can be installed by inserting the CD and following the instructions.

Note: The WSRM utility also can be downloaded, in ISO (CD image) format, from:

<http://www.microsoft.com/windowsserver2003/downloads/wsrmspx>

5.4.1 WSRM description

Is very common that customers run multiple production applications on one server. These applications often compete for system resources, such as processor and memory. When competition is highest among applications (during peak periods, for example), the applications that need, for business reasons, to have the most resources available to them often do not get this allocation, and performance suffers. Also older applications might have the tendency not to release the resources they have been using after they have finished completing their tasks. Applications causing memory leaks also present significant problems to the function of the system, and to other applications that are installed.

All of these circumstances can cause the application or the operating system to become unstable. WSRM helps avoid these situations by specifically allocating and managing resources to applications, processes, and services, thereby restricting the applications from consuming more resources than they have otherwise been allocated. These are some environments of use:

- ▶ Consolidate applications from several servers.
- ▶ Improve server hardware utilization.
- ▶ Fair use of resources.
- ▶ Limit resource-intensive applications.

Figure 5-13 on page 198 shows the main window of WSRM.

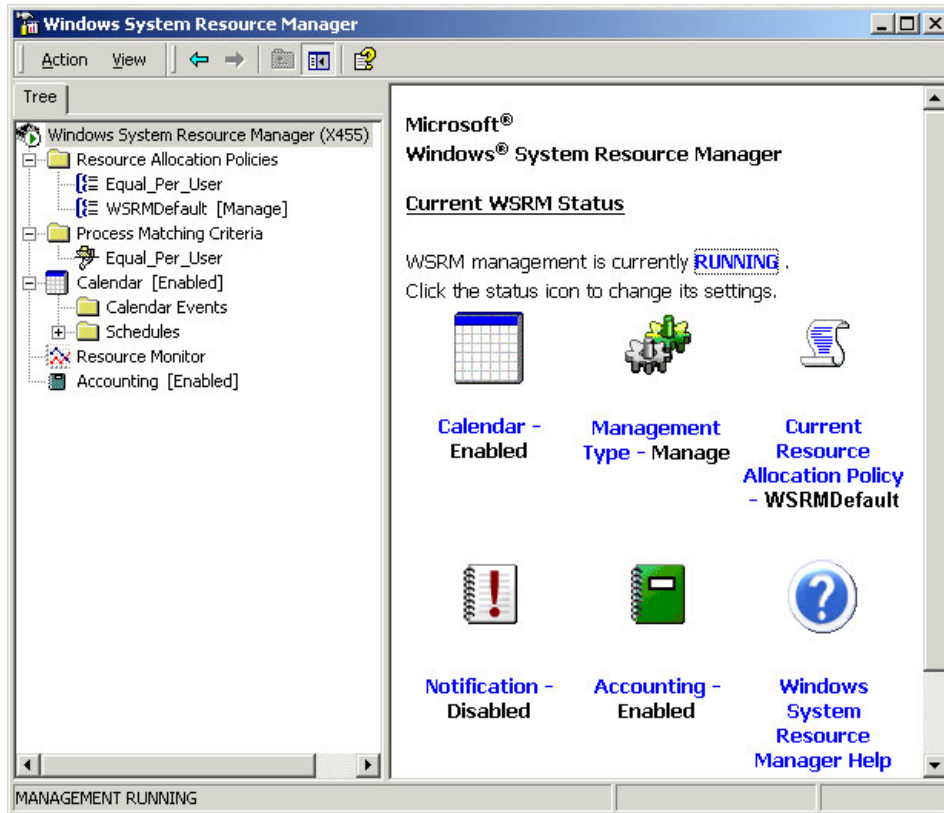


Figure 5-13 WSRM main window

5.4.2 WSRM features

In this section we discuss WSRM features.

Process matching criteria

Services, processes, or applications that are to be managed by WSRM must be defined using process matching criteria. This criteria specifies the logic used to match and group processes. When a process is started, WSRM tries to match that process with process matching criteria. If it finds a match to the process, the process will then be managed according to the managing resource allocation policy with which it is associated, and rules specified in the resource allocation policy are applied.

Resource allocation policies

Resource allocation policies enables WSRM to specify the resource usage for a process. Before you use this feature you have to create a process matching criteria. The process matching criteria is then added to a resource allocation policy that then defines one or more resource allocations.

Calendar

The calendar provides the ability to schedule the automatic activation of a resource allocation policy. Resource allocation policies can be entered in the calendar to run once, or on a regular basis by defining a recurring calendar event.

Resource Monitor

WSRM provides performance objects for System Monitor. This provides the ability to measure the performance of the WSRM system and other systems on the network. This real-time view of system performance can be saved and viewed at a later date and can also be saved as an HTML file.

Accounting

Accounting logs information based on the behavior of managed processes. The information logged is stored in the accounting database. By default, information is logged into the database every 10 minutes, and this setting can be modified.

Note: If you want to learn more about WSRM, you should read *Implementing Windows System Resource Manager*, REDP-3701.

5.4.3 WSRM in the x455

The WSRM can be installed in the x455 only as a server. This is because the WSRM installer does not have a client for 64-bit. In this case you have to install a client in other Windows 32-bit in order to manage the WSRM of the x455.

Note: The client component runs on Windows XP, the Windows Server 2003 family (32-bit), and Windows 2000 with Service Pack 2 or later.

Once you have installed both (client and server), it is started on the client machine by clicking **Start -> Programs -> Administrative Tools -> Windows System Resource Manager**. Once WSRM has started, click **Action -> Connect to Computer** to connect with x455. See Figure 5-14 on page 200.

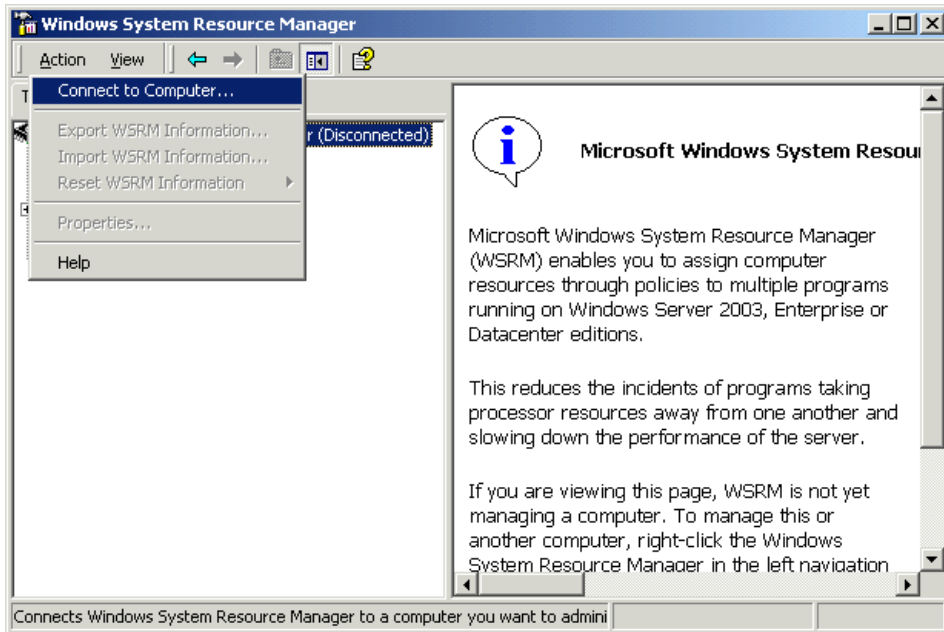


Figure 5-14 Connect to the x455

Abbreviations and acronyms

ABI	Application Binary Interface	EB	Exabytes
ACPI	Advanced Configuration and Power Interface	ECC	error checking and correcting
ACT	Advanced Connectivity Technology	EFI	Extensible Firmware Interface
ANSI	American National Standards Institute	EL	Enterprise Linux
API	application programming interface	EMC	electromagnetic compatibility
AS	advanced server	EPIC	Explicitly Parallel Instruction Computing
ASM	advanced system management	ER	Enterprise Rack
ASMA	Advanced System Management Adapter	ERP	Enterprise Resource Planning
ASR	automatic server restart	ESCON	enterprise systems connection
ATS	Advanced Technical Support	ESP	EFI System Partition
AWE	Advanced Windowing Extensions	ESS	Enterprise Storage Server
BI	business intelligence	FAT	file allocation table
BSP	bootstrap processor	FC	Fibre Channel
CAD	computer-aided drawing	FICON™	Fiber Connection
CAE	computer-aided engineering	GAS	GNU Assembler
CD-ROM	compact disc-read only memory	GB	gigabyte
CERN	Conseil Européen pour la Recherche Nucleaire	GCC	GNU Compiler Collection
CIM	Common Information Model	GNU	GNU's Not Unix
CNA	Certified NetWare Administrator	GPT	GUID Partition Table
CPU	central processing unit	GUI	graphical user interface
CRM	customer relationship management	GUID	global unique identifier
DDR	double data rate	HA	high availability
DDS	digital data storage	HAL	hardware abstraction layer
DHCP	Dynamic Host Configuration Protocol	HBA	host bus adapter
DIMM	dual inline memory module	HCL	hardware compatibility list
DLT	Digital Linear Tape	HP	Hewlett Packard
DNS	Domain Name System	IEC	International Electrotechnical Commission
DOS	disk operating system	ILP	instruction-level parallelism
DRDA	Distributed Relational Database Architecture	IPC	interprocess communications
		IPF	Itanium Processor Family
		ISA	industry standard architecture

IT	information technology	RISC	reduced instruction set computing
KB	kilobyte	RMB	rectifier monophas battery
KVM	keyboard video mouse	ROM	read only memory
LAN	local area network	RPM	revolutions per minute
LCCM	LANClient Control Manager	RSA	Remote Supervisor Adapter
LCM	Local Console Manager	SAL	system abstraction layer
LED	light emitting diode	SAN	storage area network
LTO	Linear Tape-Open	SAR	solution assurance review
LVM	Logical Volume Manager	SCM	supply chain management
MAC	medium access control	SCSI	small computer system interface
MB	megabyte	SDK	software development kit
MBR	master boot record	SDRAM	synchronous dynamic random access memory
MCSE	Microsoft Certified Systems Engineer	SLES	SUSE LINUX Enterprise Server
MSCS	Microsoft Cluster Services	SMP	symmetric multiprocessor
MSR	Microsoft Reserved Partition	SNMP	Simple Network Management Protocol
NAS	Network Attached Storage	SP	service processor
NEMA	National Electrical Manufacturers Association	SQL	Structured Query Language
NUMA	Non-uniform Memory Access	SR	standard rack
NVRAM	non-volatile random access memory	SSL	Secure Sockets Layer
OEM	original equipment manufacturer	SVGA	super video graphics array
OLAP	Online Analytical Processing	TB	terabyte
PCD	Personal Computing Division	TCP/IP	Transmission Control Protocol/Internet Protocol
PCI	Peripheral Component Interconnect	TSC	time stamp counter
PFA	Predictive Failure Analysis	UCO	USB Conversion Option
POST	power on self test	UDB	Universal Data Base
PPP	point-to-point protocol	UL	UnitedLinux
PXE	Preboot eXecution Environment	UPS	uninterruptible power supply
RAID	redundant array of independent disk	URL	Uniform Resource Locator
RAM	random access memory	USB	universal serial bus
RCM	remote console manager	VRM	voltage regulator module
RDIMM	registered dual inline memory module	WSRM	Windows System Resource Manager
RDM	Remote Deployment Manager	ZIF	zero insertion force
RETAIN	remote technical assistance information network		
RH	Red Hat		

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 208. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *Implementing Systems Management Solutions using IBM Director*, SG24-6188
- ▶ *IBM @server xSeries 445 Planning and Installation Guide*, SG24-8870
- ▶ *IBM Tape Solutions for Storage Area Networks and FICON*, SG24-5474
- ▶ *Netfinity Tape Solutions*, SG24-5218
- ▶ *The IBM LTO Ultrium Tape Libraries Guide*, SG24-5946
- ▶ *Windows Server 2003, Datacenter Edition on the IBM @server xSeries 445*, REDP-3700
- ▶ *Implementing Windows System Resource Manager*, REDP-3701

Other publications

These product publications are also relevant as further information sources:

- ▶ IBM @server xSeries 455 Installation Guide
<http://www.ibm.com/pc/support/site.wss/MIGR-53736.html>
- ▶ IBM @server xSeries 455 Option Installation Guide
<http://www.ibm.com/pc/support/site.wss/MIGR-53735.html>
- ▶ IBM @server xSeries 455 Rack Installation Instructions
<http://www.ibm.com/pc/support/site.wss/MIGR-53737.html>
- ▶ IBM @server xSeries 455 User's Guide
<http://www.ibm.com/pc/support/site.wss/MIGR-53738.html>

- ▶ IBM @server xSeries 455 Hardware Maintenance Manual
<http://www.ibm.com/pc/support/site.wss/MIGR-53800.html>
- ▶ IBM @server xSeries 455 FAQ Hints and Tips
<http://www.ibm.com/pc/support/site.wss/MIGR-53820.html>

Online resources

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

IBM product information

- ▶ IBM @server xSeries 455
<http://www.pc.ibm.com/us/eserver/xseries/x455.html>
- ▶ Enterprise X-Architecture
<http://www.pc.ibm.com/us/eserver/xseries/xarchitecture/enterprise>
- ▶ Windows Datacenter from IBM
<http://www.pc.ibm.com/ww/eserver/xseries/windows/datacenter.html>
- ▶ xSeries configuration tools
<http://www.pc.ibm.com/europe/configurators>
<http://www.pc.ibm.com/us/eserver/xseries/library/configtools.html>
- ▶ IBM tape products
<http://ibm.com/pc/ww/eserver/xseries/tape.html>
- ▶ IBM SCSI products
http://ibm.com/pc/ww/eserver/xseries/scsi_raid.html
- ▶ IBM Fibre Channel products
<http://www.storage.ibm.com/hardsoft/disk/fastt/index.html>
- ▶ Clustering information
<http://www.pc.ibm.com/ww/eserver/xseries/clustering/info.html>

IBM support information

Also see “Linux” on page 206.

- ▶ xSeries support site
<http://www.pc.ibm.com/support>

- ▶ IBM x455 device driver matrix
<http://www.ibm.com/pc/support/site.wss/MIGR-53575.html>
- ▶ IBM ServerProven
<http://www.pc.ibm.com/ww/eserver/xseries/serverproven>
- ▶ ServerProven option compatibility page for the x455
<http://www.pc.ibm.com/us/compat/machines/x455.html>
- ▶ ServerProven operating system support matrix
<http://www.pc.ibm.com/us/compat/nos/matrix.shtml>
- ▶ ServerProven partners
<http://www.pc.ibm.com/us/compat/serverproven>
- ▶ ServerProven LAN adapter support matrix
<http://www.pc.ibm.com/us/compat/lan/matrix.html>
- ▶ IBM solution assurance
<http://w3.ibm.com/support/assure>
<http://www.ibm.com/support/assure/esar>
- ▶ RETAIN tip H18505, use of ServeRAID on IA-64 systems
<http://www.ibm.com/pc/support/site.wss/MIGR-53247.html>
- ▶ Linux operating system installation instructions (IBM employees only)
<http://nosinfo.raleigh.ibm.com>
- ▶ RETAIN tip H177072, differences between ServeRAID and LSI RAID-1
<http://www.ibm.com/pc/support/site.wss/MIGR-46174.html>
- ▶ ServeRAID driver downloads
<http://www.ibm.com/pc/support/site.wss/MIGR-495PES.html>

Intel

- ▶ Extensible Firmware Interface information
<http://www.intel.com/technology/efi>
- ▶ EFI disk utilities (diskpart, efichk, efifmt)
http://developer.intel.com/technology/efi/diskutil_overview.htm
- ▶ EFI specification in PDF format
<http://developer.intel.com/technology/efi/download.htm>
- ▶ Itanium 2 product information
<http://www.intel.com/design/itanium2>

- ▶ Itanium product information
<http://www.intel.com/products/server/processors/server/Itanium>

Microsoft

- ▶ EFI information from Microsoft
<http://www.microsoft.com/hwdev/platform/firmware/EFI/default.asp>
- ▶ IA-32 Execution Layer 4.3 Software Driver
<http://www.microsoft.com/windowsserver2003/64bit/ipf/ia32el.msp>
- ▶ Windows Hardware Compatibility List
<http://www.microsoft.com/whdc/hcl/search.msp>
- ▶ Windows Server 2003 64-bit versions
<http://www.microsoft.com/windowsserver2003/64bit>
- ▶ Partitions required on GPT disks for Windows 64-bit operating systems
http://www.microsoft.com/technet/prodtechnol/winxppro/reskit/prkb_cn_c_helm.asp
- ▶ Windows System Resource Manager
<http://www.microsoft.com/windowsserver2003/downloads/wsrp.msp>

Linux

- ▶ ServerProven support of Red Hat Linux products
<http://www.pc.ibm.com/us/compat/nos/redchat.html>
- ▶ ServerProven support of Turbolinux products
<http://www.pc.ibm.com/us/compat/nos/turbolinux.html>
- ▶ ServerProven support of SUSE LINUX products
<http://www.pc.ibm.com/us/compat/nos/suse.html>
- ▶ Linux on IBM desktop PCs and notebooks
<http://www.pc.ibm.com/qtechinfo/MIGR-48NT8D.html>
- ▶ IBM Storage Systems for Linux
<http://www.storage.ibm.com/linux>
- ▶ The Linux Kernel Archives
<http://www.kernel.org>
- ▶ System V UNIX ABI (Application Binary Interface)
<http://www.linuxbase.org/spec/refspecs/elf/gabi4+/contents.html>

- ▶ The Open Group
<http://www.opengroup.org>
- ▶ Red Hat Enterprise Linux AS
<http://www.redhat.com/software/rhel/as>
- ▶ UnitedLinux FAQ
<http://www.unitedlinux.com/en/info/faqs.html>
- ▶ Conectiva Linux
<http://www.conectiva.com>
- ▶ SCO Linux
<http://www.sco.com/products>
<http://www.pc.ibm.com/us/compat/nos/thescocgroup.html>
- ▶ Turbolinux products
<http://www.turbolinux.com/products>
- ▶ GNU Parted — modifying disk partitions and file systems
<http://www.gnu.org/software/parted>
- ▶ ATI video driver support from The XFree86 Project
<http://www.xfree86.org/current/Status6.html#6>
- ▶ ATI video driver information
<http://mirror.ati.com/support/faq/linux.html>
- ▶ ReiserFS information
<http://www.namesys.com>

Other

- ▶ Developers Interface Guide for IA-64 processors
<http://www.dig64.org>
- ▶ SAP R/3 benchmarks
<http://www.sap.com/benchmark>
- ▶ Standard Performance Evaluation Corporation
<http://www.spec.org>

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

Numerics

03K9309, ASM Interconnect Cable Kit 77
22P6801, NetXtreme 1000 T Ethernet adapter 64
22P7801, NetXtreme 1000 SX Fiber Ethernet 64
31P5998, RXE-100 6-slot Expansion Kit 70
31P6087, 3.5 m management cable kit 72
31P6088, 8 m management cable kit 72
31P6102, 3.5 m Remote I/O cable kit 71
31P6103, 8 m Remote I/O cable kit 71
31P6301, NetXtreme 1000T Ethernet adapter 64
32-bit compatibility driver 40
32-bit PCI adapters 61
32-bit to 64-bit migration 40
33L5038, 512 MB DDR DIMM 57
33L5039, 1 GB DDR DIMM 57
33L5040, 2 GB DDR DIMM 57
73P5832, USB Conversion Option 19, 66
73P7076, 1.3 GHz processor option 56
73P7077, 1.4 GHz processor option 56
73P7078, 1.5 GHz processor option 56
73P9715, 8-way to 16-way Supplemental Kit 67
73P9911, 4-way to 8-way Scalability Kit 67
8855-1RX 4, 56
8855-2RX 4, 56
8855-3RX 4, 56

A

Active Memory 14
 mirroring 15
Active PCI driver
 Windows 2000 145
Adaptec RAID Driver 108
adapter speeds 62
adapter teaming 64
alerts 188
application consolidation 44
ASM
 device driver 185
 interconnect 184
AWE in 32-bit environments 47

B

backup 84
bandwidth calculations 8
BIOS 29–31
block diagram 8
boot maintenance menu 122
Boot Manager Menu 94
bootstrap processor 9
Broadcom Gigabit Ethernet 18, 64
 adapter teaming 64
 Windows 2000 145
business intelligence 49

C

cabling
 Ethernet 19, 64
 RXE-100 70–71
 scalability 67
cache 11–12
Cache and Scalability Controller 7
Caldera SCO Linux 155
CD-ROM file formats 99
centralization 44
checkpoint 93
Chipkill memory 16
clocking 67
clustering 88
comparison, x455 and x450 2
Conectiva Linux 155
Configuration/Setup Utility 118
CPUs 56
CRM 50
cyclone chip 8

D

data consolidation 44
database applications 46
Datacenter Solution Program 45
DB2 Universal Database 48
DDR memory 12–13
diagnostics 121
 updates 113
DIMMs

- installation order 57
- maximum 57
- memory-board assembly 13
- standard memory 13
- disk drive bays
 - multinode configuration 26
 - redundant with RAID-1 27
- DISKPART.EFI 101, 103
- drive bays 2, 27
- Driver Setup 103

E

- e-business 50
- EFI 29, 92
 - accessing the system partition 147
 - adding a boot option 125
 - backup and restore 123
 - boot maintenance menu 122
 - Boot Manager Menu 94
 - CD-ROM file formats 99
 - changing a boot option 126
 - check disk command 100
 - Configuration/Setup Utility 118
 - deleting a boot option 128
 - deleting a disk 100
 - DHCP address of RSA 179
 - diagnostics 121
 - Driver Setup 103
 - file operations 97
 - file systems 99
 - flash updates 112
 - hints & tips 102
 - IDE devices 104
 - map command 98
 - mapping table 96
 - MemMap entries 96
 - PXE boot 94
 - RAID-1, creating 105
 - remote control password 186
 - RSA firmware 116
 - SCSI controller 105
 - ServeRAID 108
 - shell 95
 - system partition 97, 147
 - tasks 92
 - TimeOut value 130
 - Windows boot record 148
- EFI System Partition 33

- EFICLK.EFI 100, 103
- EFIFMT.EFI 103
- EI-Torito format 99
- Enterprise Storage Server 84
- Enterprise X-Architecture 1, 35, 41
- ERP 49
- error logs 120
- eSAR 90
- ESS 84
- Ethernet
 - adapter teaming 64
 - features 18
 - multinode configuration 26
- event action plans 193
- events 188
- Execution Layer software driver 40
- EXP300 80
- EXP400 80
- EXP500 81
- EXP700 83
- Explicitly Parallel Instruction Computing 11
- Extensible Firmware Interface 29, 92

F

- fans 27
- FAST 80
- FAST200 81
- FAST600 81
- FAST700 82
- FAST900 83
- features 2
 - not supported 3
- Fibre Channel
 - booting from 112
 - storage options 80
- file operations, EFI 97
- flash updates 112
- FLASH2.EFI 103, 116, 181
- four-node configuration 68
- front view 4, 27
- frontside bus 9, 11
- FS entries 96

G

- Gigabit Ethernet 18, 64
- GPT disks 31, 137
- GUID Partition Table 31

H

hints & tips for EFI 102
hot-add of memory (not supported) 3
hot-swap fans 27
hot-swap memory 59

I

IA-32 Execution Layer software driver 40
IBM Datacenter Solution Program 45
IBM Director 89, 176
 event action plans 193
IBM ServerProven 44
IBM XA-64 chipset 7
IDE devices, configuring 104
in-house applications 51
in-memory databases 47
installing
 in a rack 85
 Red Hat Linux 160
 SUSE LINUX 156
 Windows Server 2003 136
Intel PRO/1000 XT 64
interleaving, memory 13–14
Itanium 2 10

J

Joliet format 99

K

keyboard port 3, 19
KVM USB conversion option 19

L

LCCM 3
LEDs, light path diagnostics 28
legacy 34
Level 3 cache 11
Level 4 cache 12
light path diagnostics 27
Linux
 boot process 162
 byte order 153
 Caldera SCO Linux 155
 Conectiva Linux 155
 CPU information 163
 data types 152
 distributions 153

 drivers 173
 IA-32 applications 163
 kernel 151
 memory information 164
 modules used 165
 partitions 166
 Red Hat Linux 160
 RXE-100 172
 serial port use 171
 SUSE LINUX 156
 Turbolinux 155
 UnitedLinux 154
logging 120
LSI controller, configuring 105

M

Madison 4, 10
management 175
map command 98
mapping table 96
MBR 32
McKinley 10
media bays, multinode configuration 26
MemMap entries 96
memory 14
 Active Memory 14
 Chipkill memory 16
 disabling during POST 16
 hot-add (not supported) 3
 hot-swap 59
 installation order 57
 interleaving 14
 matched pairs 14, 57
 memory mirroring 15, 58, 118
 memory ports 58
 Memory ProteXion 14
 memory scrubbing 15
memory-board assembly 6, 13, 57
 LEDs 13
Microsoft Reserved Partition 137
Microsoft SQL Server 47
migrating to 64-bit 40
models 4
mountvol command 147
mouse port 3, 19
MSCS 88
multinode
 cabling 67

- configuring 131
- four-node 68
- two-node 68
- multiple chassis 22
- multiple monitors 135

N

- NetBurst architecture 9
- NetXtreme 64
- nodes 22
- non-redundant LED 87
- NUMA 36

O

- OLAP 49
- OnForever 41
- operating system support 87
- operating systems supported 35
- OS loader 29

P

- parallel port 3
- partitions 22
- PCI-X
 - See also* RXE-100
 - board assembly 6
 - bridges 17
 - coexistence with PCI 62
 - scan order 63
 - slots in the RXE-100 70
 - slots in the x455 17, 61
- performance
 - memory 14, 60
 - PCI-X slots 18
- physical consolidation 44
- planning 55–90
- POST 93
- power
 - considerations 87
 - modules (processor board) 9–10
 - supplies 27
- power supplies 27
- primary node 23, 26
- processor-board assembly
 - LEDs 10
- processors
 - installation order 9, 56

- option part numbers 56
- processor-board assembly 6, 9
- upgrades 56
- protective MBR 32
- PS/2 keyboard and mouse 3
- PXE boot 94

R

- rack installation 85
- RAID-1 array, creating 105
- RAMDISK.EFI 103
- rear view 5, 17
- Red Hat Linux 154
 - installing 160
- Redbooks Web site 208
 - Contact us xii
- reduced-legacy 34
- redundancy 27
- remind button 10
- remote control password 186
- Remote Deployment Manager 3
- Remote Expansion I/O port 17
- Remote Supervisor Adapter 177
 - alerts 188
 - ASM interconnect 184
 - connectivity 76
 - device driver 185
 - DHCP address 179
 - driver 145
 - features 19
 - firmware updates 116
 - in-band 178
 - out-of-band 179
 - remote control password 186
 - static address 181
- RS-232 serial port 19
- RS-485 19
- rules
 - PCI slots 61
- RXE-100 21, 70
 - cabling 71
 - connectivity options 21, 24, 72
 - Linux 172
 - slots 71

S

- SAL updates 113
- scalability

- cables 67
- ports 23
- scalable partitions, configuring 131
- Scalable Systems Manager 89, 177
- scan order, PCI 63
- scrubbing, memory 15
- SCSI
 - configuring the controller 105
 - multinode configuration 26
 - ports 18
- SDRAM 13
- secondary nodes 26
- serial port 19, 78
 - Linux, use with 171
 - multinode configuration 26
- server consolidation 43
- ServeRAID
 - booting from 111
 - configuration CD-ROM 110
 - configuring 108
 - support 79
- ServerGuide 3
- ServerProven 44
- shark 84
- shell, EFI 95
- SNMP 187, 190
- Solution Assurance Review 90
- SQL Server 47
- storage 79
- Summit chipset 7
- supply chain management 49
- SUSE LINUX, installing 156
- system partition, EFI 97

T

- tape 84
- teaming 64
- time stamp counter 67
- tornado chip 7
- TotalStorage solutions 81
- Trigger Tool 90
- Turbolinux 155
- two-node configuration 68

U

- Ultra 320 SCSI ports 18
- UnitedLinux 154
- UpdateExpress 3

- USB 2.0 ports 19
 - multinode configuration 26
- USB Conversion Option 19

V

- video
 - controller 19
 - multinode configuration 26
- video adapters, not supported 63

W

- Wake on LAN, multinode configuration 26
- Web browser 179
- why x455 52
- Windows Server 2003
 - IA-32 Execution Layer software driver 40
 - installing 136
 - unknown devices 145
- Windows System Resource Manager 196
- winnipeg chip 8

X

- x450 comparison 2
- XA-64 chipset 7
- X-Architecture 1
- XceL4 Server Accelerator Cache 12
- XpandOnDemand 36, 41

Z

- ZIF socket 10



IBM SERV_ER xSeries 455 Planning and Installation Guide



IBM @server xSeries 455 Planning and Installation Guide



Describes the technical details of the new 64-bit server

Covers supported Windows and Linux 64-bit operating systems

Helps you prepare for and perform an installation

The IBM @server xSeries 455 is the second-generation Enterprise X-Architecture server using the 64-bit IBM XA-64 chipset and the Intel Itanium 2 processor. Unlike the x450, its predecessor, the x455 supports the merging of four server chassis to form a single 16-way image, providing even greater expandability and investment protection.

This IBM Redbook is a comprehensive resource on the technical aspects of the server, and is divided into five 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 would be used on a server such as the x455.
- ▶ Chapter 3, “Planning”, describes the considerations when planning to purchase and planning to install the x455.
- ▶ Chapter 4, “Installation”, covers the process of installing Windows Server 2003, SuSE Linux Enterprise Server, and Red Hat Enterprise Linux AS on the x455.
- ▶ Chapter 5, “Management”, describes how to use the Remote Supervisor Adapter to send alerts to an IBM Director management environment.

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