


Planning and Installing the IBM @server X3 Architecture Servers

Covers the xSeries 460, xSeries 366,
and xSeries 260

Learn the technical details of
these high-performance servers

How to configure, install,
and manage the servers



David Watts
Sergey Gorbas
Christian Schroeder
Joshua Young



International Technical Support Organization

**Planning and Installing the IBM @server X3
Architecture Servers**

January 2006

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

First Edition (January 2006)

This edition applies to the following servers:

- ▶ IBM @server xSeries 460, machine type 8872
- ▶ IBM @server xSeries MXE-460, machine type 8874
- ▶ IBM @server xSeries 366, machine type 8863
- ▶ IBM @server xSeries 260, machine type 8865

© Copyright International Business Machines Corporation 2006. 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	xi
Chapter 1. Technical overview	1
1.1 The X3 Architecture servers	2
1.2 Key features	3
1.3 Multi-node capabilities and partitioning	4
1.4 IBM XA-64e third generation chipset	8
1.5 Scalable implementation	9
1.6 Current models	12
1.6.1 x460 and MXE-460 models	12
1.6.2 x366 models	13
1.6.3 x260 models	14
1.7 Front and rear layouts	15
1.7.1 x460 and MXE-460	15
1.7.2 x366	16
1.7.3 x260	17
1.8 Processors	18
1.8.1 Models with dual-core processors	20
1.8.2 667 MHz front-side bus	21
1.8.3 Hyper-Threading	21
1.8.4 Intel 64-bit Extension Memory Technology (EM64T)	23
1.8.5 The benefit of 64-bit computing	26
1.9 XceL4v cache	27
1.10 Memory	28
1.11 Serial Attached SCSI	31
1.12 PCI-X slots	34
1.13 Remote Supervisor Adapter II SlimLine	36
1.14 Baseboard Management Controller	37
1.15 Light path diagnostics and PFA	38
Chapter 2. Positioning	41
2.1 Positioning the servers	42
2.1.1 x366 versus x360	46

2.1.2	x460 versus x445	48
2.1.3	x260 versus x255	50
2.2	Focus market segments and target applications	52
2.2.1	x260 target applications	54
2.2.2	x366 target applications	55
2.2.3	x460 target applications	55
2.3	Server consolidation	56
2.3.1	Why consolidate servers?	56
2.3.2	Types of server consolidation	58
2.3.3	Benefits of server consolidation	60
2.4	XpandOnDemand	64
2.5	Positioning the x460 and the MXE-460	66
2.6	Application scalability	68
2.6.1	Enterprise application scaling	70
2.6.2	SAP product offering	71
2.6.3	Oracle/J.D. Edwards product offering	73
2.6.4	Siebel product offering	75
2.6.5	Oracle/PeopleSoft product offering	77
2.6.6	Microsoft SQL product offering	79
2.6.7	DB2 Universal Database	80
2.7	Scale-up versus scale-out	82
Chapter 3. Hardware planning		89
3.1	Processor subsystem	90
3.1.1	Dual-core upgrades	91
3.1.2	Processor configuration options	92
3.2	Memory subsystem	95
3.2.1	Memory mirroring	98
3.2.2	Hot-swap memory	99
3.2.3	Hot-add memory	100
3.2.4	Memory ProteXion: redundant bit steering	101
3.2.5	Memory configuration in BIOS	102
3.3	Multi-node configurations	104
3.3.1	Cabling	107
3.3.2	Scalable system setup	110
3.3.3	Partitioning	116
3.3.4	Export regulations for 32-core complexes	118
3.4	RSA II SL and BMC initial setup	119
3.4.1	RSA II SlimLine setup	120
3.4.2	BMC IP address setup	122
3.5	Storage options	122
3.5.1	Serial Attached SCSI (SAS) subsystem	122
3.5.2	ServerRAID-8i	123

3.5.3	Internal storage	124
3.5.4	External storage options	126
3.6	Power considerations	126
3.6.1	Working in 110V environments	127
3.7	Performance tuning and optimization	129
3.7.1	Optimal memory module installation	129
3.7.2	Memory settings in BIOS	129
3.7.3	CPU settings in BIOS	130
3.7.4	PCI adapter placement	130
3.8	ServerProven	130
3.9	Solution Assurance	131
3.9.1	Solutions Assurance Reviews	133
Chapter 4. Operating system installation		135
4.1	Supported operating systems	136
4.1.1	Operating systems scalability	137
4.1.2	Hyper-Threading support	137
4.2	Updating BIOS and firmware	140
4.3	Available memory in multi-node configurations	141
4.4	Microsoft Windows Server 2003	142
4.4.1	Installing Windows Server 2003	145
4.5	Datacenter offerings	147
4.6	Microsoft Windows 2000 Server	152
4.7	Red Hat and SUSE LINUX	153
4.7.1	Additional information for SLES 9	155
4.8	VMware ESX Server 2.5.1	155
4.8.1	Support for applications running on ESX Server	157
4.8.2	Pre-install information	158
4.8.3	Installing ESX Server	159
Chapter 5. Management		161
5.1	IBM Director	162
5.2	Remote Supervisor Adapter II SlimLine	164
5.2.1	Installing the RSA II SlimLine	165
5.2.2	Connectivity	166
5.2.3	Network settings	166
5.2.4	Web interface	168
5.2.5	Updating firmware	170
5.2.6	Installing the device driver	171
5.2.7	Remote console and media	174
5.2.8	TCP/UDP ports used by the RSA II SlimLine	176
5.2.9	MIB files	176
5.3	Baseboard Management Controller	177

5.3.1	Connectivity	177
5.3.2	Updating the BMC firmware	178
5.3.3	Configuring the BMC in BIOS	179
5.3.4	Event Log	180
5.3.5	Remote control	182
5.3.6	Installing the BMC device drivers	182
5.3.7	Ports used by the BMC	186
5.4	Integrating the service processors with IBM Director	187
5.4.1	Adding the service processor to IBM Director	187
5.4.2	Adding users	188
5.5	OSA SMBridge utility	194
5.5.1	Configuring BIOS	196
5.5.2	Installation	197
5.5.3	Connecting via the telnet server	201
5.5.4	Configuring Windows Server 2003 to support SOL	206
5.5.5	Configuring Red Hat Linux to support SOL	212
5.5.6	Configuring SUSE LINUX to support SOL	215
5.5.7	Connecting via the command-line interface	217
5.6	Predictive Failure Analysis (PFA)	220
5.7	IBM Dynamic System Analysis	220
5.7.1	Introduction	220
5.7.2	Usage	222
5.8	Partition management	223
Abbreviations and acronyms		225
Related publications		229
IBM Redbooks		229
Other publications		230
xSeries 460 and MXE-460 publications		230
xSeries 366 publications		231
xSeries 260 publications		231
Online resources		231
IBM Web pages		231
Intel Web pages		234
Microsoft Web pages		234
Others		234
Help from IBM		235
Index		237

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:

AIX®	iSeries™	ServerProven®
BladeCenter®	Lotus®	ServeRAID™
Chipkill™	Netfinity®	ServicePac®
DB2 Universal Database™	NetView®	Tivoli Enterprise™
DB2®	PartnerWorld®	Tivoli®
DFS™	PowerPC®	TotalStorage®
@server®	Predictive Failure Analysis®	WebSphere®
@server®	pSeries®	X-Architecture™
i5/OS®	Redbooks™	xSeries®
Informix®	Redbooks (logo)  ™	zSeries®
IBM®	ServerGuide™	

The following terms are trademarks of other companies:

Opteron, AMD Opteron, AMD, the AMD Arrow logo, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

IPX, Java, J2EE, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

i386, Intel, Itanium, Pentium, Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Preface

The IBM @server® X3 Architecture servers are the new third-generation Enterprise X-Architecture™ servers from IBM®. The xSeries® 366 was announced in February 2005, the xSeries 460 was announced in May, and the xSeries 260 was announced in August. These X3 Architecture servers are ideal for random commercial workloads with high processor and memory bandwidth requirements.

Although aimed at different client sets, they share common components and options. The x460 is targeted at clients who need to implement a large single image “scale-up” configuration of up to 32 processors and 512 GB of RAM. The x366 is targeted at the high-performance environment where rack space is a premium. The x260 is for clients who need high performance but also require large amounts of internal disk storage.

This IBM Redbook provides a detailed technical description of the three servers and explains how to plan, install, configure, and manage these high-performance servers running 32-bit and 64-bit versions of Windows® Server 2003, Red Hat Enterprise Linux®, SUSE Linux Enterprise Server, and VMware ESX Server.

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 IBM ITSO Center in Raleigh. He manages residencies and produces Redbooks™ on hardware and software topics related to IBM @server xSeries systems and associated client platforms. He has authored over 30 Redbooks and Redpapers. He holds a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM for over 15 years. He is an IBM @server Certified Specialist for xSeries and an IBM Certified IT Specialist.

Sergey Gorbis is an IBM @server® xSeries Technical Sales Team leader in Russia. He has been working for IBM for 11 years. He has more than nine years of experience in xSeries and Netfinity® servers, including pre-sales and post-sales technical support. He holds a degree in Computer Engineering from the Moscow Engineering Physics Institute in Moscow. He is an IBM @server Certified Specialist for xSeries and his areas of expertise include xSeries systems management hardware and software, including IBM Director,

IBM @server Cluster 1350, Microsoft® clustering, Microsoft Datacenter solutions, and IBM DS Storage Servers.

Christian Schroeder is a Support Specialist in Germany. He has been working in this function and with IBM for six years. His areas of expertise include IBM @server xSeries, BladeCenter®, systems management hardware, and Windows operating systems.

Joshua Young is a university intern at the International Technical Support Organization. He has four years of experience in Information Technology related fields. He will earn a degree in Information Sciences and Technology from Pennsylvania State University. His skill areas include server environments, operating systems, team management, and problem-based learning. He also co-authored the *IBM @server xSeries 260 Solution Assurance Product Review Guide*, REDP-4007.



Figure 1 The team (l-r): Sergey, David, Josh, and Christian

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

- ▶ Daniel Ghidali, IBM Center for Microsoft Technologies, Kirkland
- ▶ Michael Lee, IBM Center for Microsoft Technologies, Kirkland
- ▶ Silvio Erdenberger, xSeries Pre-Sales Support Specialist, Germany
- ▶ Ralph Begun, xSeries Engineering, Raleigh
- ▶ Gregg McKnight, xSeries Performance Lab, Raleigh

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 clients.

Your efforts will help increase product acceptance and client 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 e-mail 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 overview

The IBM @server X3 Architecture is the culmination of many years of research and development and has resulted in what is currently the fastest processor and memory controller in the Intel® processor marketplace. With support for up to 32 Xeon® MP processors and over 20 GBps of memory bandwidth per 64 GB of RAM (up to a maximum of 512 GB), the xSeries servers that are based on the X3 Architecture offer maximum performance and broad scale-up capabilities.

Topics in this chapter are:

- ▶ 1.1, “The X3 Architecture servers” on page 2
- ▶ 1.2, “Key features” on page 3
- ▶ 1.3, “Multi-node capabilities and partitioning” on page 4
- ▶ 1.4, “IBM XA-64e third generation chipset” on page 8
- ▶ 1.5, “Scalable implementation” on page 9
- ▶ 1.6, “Current models” on page 12
- ▶ 1.7, “Front and rear layouts” on page 15
- ▶ 1.8, “Processors” on page 18
- ▶ 1.9, “XceL4v cache” on page 27
- ▶ 1.10, “Memory” on page 28
- ▶ 1.11, “Serial Attached SCSI” on page 31
- ▶ 1.12, “PCI-X slots” on page 34
- ▶ 1.13, “Remote Supervisor Adapter II SlimLine” on page 36
- ▶ 1.14, “Baseboard Management Controller” on page 37
- ▶ 1.15, “Light path diagnostics and PFA” on page 38

1.1 The X3 Architecture servers

The three servers based on the X3 Architecture are the x460, x366 and x260. They have a common set of technical specifications and features, but the key differences are:

The *xSeries 460* is the flagship server with the following characteristics:

- ▶ Each chassis occupies 3U of rack space and supports four CPUs and 64 GB of RAM.
- ▶ Six hot-swap drive bays and six 266 MHz PCI-X 2.0 hot-swap slots.
- ▶ Targeted at eight-way and above configurations where effective scale-up options are essential.
- ▶ Up to eight systems can be connected together to form one single 32-way complex with up to 512 GB RAM.



The *xSeries 366* is a high-performance four-way server with the following characteristics:

- ▶ The same mechanical design as the x460, offering up to 64 GB of RAM and up to four Xeon MP processors.
- ▶ Targeted at two-way and four-way high performance commercial computing, such as database, e-mail, and e-commerce applications.



The *xSeries 260* is also a high-performance four-way server:

- ▶ Same central electronics as the x366 and x460.
- ▶ Larger 7U chassis to house up to 12 hot-swap disk drives and a full-height internal tape drive.
- ▶ Targeted at two-way and four-way high performance commercial computing applications where more internal disk storage is required.



1.2 Key features

The x460, x366, and x260 have a number of common features

- ▶ IBM @server X3 Architecture, featuring the XA-64e third-generation chipset.
- ▶ Common system boards: the CPU/memory board, the I/O board, and the PCI-X board.
- ▶ Up to four Intel Xeon MP processors. These Processors support 64-bit addressing with the Intel Extended Memory 64 Technology (EM64T) architecture.
- ▶ Existing or planned support for dual-core processors. The x460 and x366 now support the Xeon 7020 and 7040 (“Paxville”) processors.
- ▶ Up to 64 GB of RAM, using high performance PC2-3200 ECC DDR2 DIMMs.
- ▶ Active Memory with Memory ProteXion, memory mirroring, memory hot-swap and hot-add, and ChipKill.
- ▶ Six full-length 64-bit 266 MHz PCI-X 2.0 Active PCI slots.
- ▶ Integrated Adaptec AIC-9410 serial-attached SCSI (SAS) controller.
- ▶ Support for internal RAID arrays using an optional ServeRAID™-8i adapter. ServeRAID-6M is also supported for external SCSI storage with the EXP400 enclosure.
- ▶ Integrated dual-port Broadcom 5704 PCI-X Gigabit Ethernet.
- ▶ Integrated Baseboard Management Controller. Remote Supervisor Adapter II SlimLine adapter standard (x460) or optional (x366 and x260).
- ▶ Support for the IBM Integrated xSeries Adapter for iSeries™ (IXA) for a direct high speed link to an iSeries server (x460 and x366 only).
- ▶ Hot-swap fans and power supply.
- ▶ Light path diagnostics to identify any failed components.
- ▶ Three-year warranty on-site, nine hours per day, five days per week, with a next business day response.

The key features that differ between the three servers are shown in Table 1-1.

Table 1-1 Key feature differences

	xSeries 460	xSeries 366	xSeries 260
Processors	Intel Xeon MP "Potomac" or Xeon 7020/7040 "Paxville" processors	Intel Xeon MP "Cranford" or Xeon 7020/7040 "Paxville" processors	Intel Xeon MP "Cranford" processors
Installed / max processors	2 / 4	1 / 4	1 / 4
Memory standard / max	2 / 64 GB	2 / 64 GB	1 or 2 / 64 GB
Largest configuration	Eight nodes (32-way)	One node (4-way)	One node (4-way)
Rack height	3U	3U	7U
Tower-to-rack conversion	No	No	Yes
Power supplies	2x 1300W supplies (650W at 110V), both standard	1300W supplies (650W at 110V), one standard, one optional	2x 775W supplies / 4 (220V or 110V), both standard
Remote Supervisor Adapter II SlimLine	Standard	Optional	Optional
Hot-swap disk drive bays	Six (2.5" bays, SAS)	Six (2.5" bays, SAS)	Six standard, additional six optional (3.5" bays)
Optical media	8x DVD-ROM	8x DVD-ROM	40x CD-ROM
Diskette drive	Optional (external USB)	Optional (external USB)	Standard (internal)
Tape drive bay	No	No	Two half-high 5.25" bays, can be used as a single full-height bay

1.3 Multi-node capabilities and partitioning

The x460 offers the ability to expand the server by connecting multiple chassis together to form a larger complex. These chassis (or *nodes*) can be either other x460s or MXE-460 modular expansion enclosures.

Note: The x260 and x366 cannot expand beyond the single chassis. Only the x460 and MXE-460 support multi-node configurations.

The MXE-460 modular expansion enclosure is a system used to extend an x460 configuration. Like the x460, it contains microprocessors, memory, disks, and PCI-X adapters. However, unlike the x460, the MXE-460 can only be used to expand an x460 configuration.

The MXE-460 is functionally identical to the x460 and supports the same hardware options. The key differences between the x460 and MXE-460 are:

- ▶ The MXE-460 is for expansion purposes and cannot be used as the primary node of a multi-node complex, and it cannot be used as the primary node in a partition.
- ▶ The MXE-460 does not have a DVD-ROM drive installed.
- ▶ The MXE-460 has no processors installed as standard.
- ▶ The MXE-460 has no memory installed (although two memory cards are installed).

An x460 server can be configured together with one, three, or seven MXE-460s to form a single 8-way, 16-way, or 32-way complex. This is shown in Figure 1-1 on page 6. Other multi-node configurations are not supported.

Important: All configurations with 32 CPU cores (that is, 16-way for dual-core x460 systems) will require US government approval regardless of where the complex will be installed.




Figure 1-1 The four multi-node configurations supported

Additional x460s can substitute for the MXE-460s in Figure 1-1. While this does increase the cost of the complex, it does mean you can implement partitioning. This is discussed further in 3.3.3, “Partitioning” on page 116.

multi-node complexes support partitioning on node boundaries. This means, for example, you can logically partition up your 16-way as two 8-way systems, while still leaving the complex cabled as a 16-way. This increases flexibility. You can reconfigure the complex using the Web interface without changing the systems or cabling.

However, as discussed in 3.3.3, “Partitioning” on page 116, every partition must have an x460 as the primary node. This restriction means that clients must decided up front if they wish to use the partitioning feature and order an appropriate combination of x460 and MXE-460 systems.

For more information about multi-node complexes, see 3.3, “Multi-node configurations” on page 104.

A multi-node configuration can be partitioned using the RSA II Web interface. However, each partition must have an x460 as the primary node. You cannot

partition a multi-node configuration so that any partition has only MXE-460 nodes. If you want maximum partitioning flexibility, all your nodes should be x460 systems.

Important: If you plan to implement partitioning on your multi-node configuration, then you must have one x460 for every partition and to be the primary node for that partition. The secondary nodes can be MXE-460 or x460 units. MXE-460s are not supported as primary nodes.

Figure 1-2 shows some of the supported partitioning options. The key messages from this chart are:

- ▶ Partitioning is always at node boundaries. You cannot have, for example, two microprocessors in a node in one partition and the other microprocessors in that node in another partition.
- ▶ The primary node for every partition must be an x460. The use of MXE-460s limits the partitioning you can do.

An eight-node configuration is not shown in the figure, but the same rules apply.




Figure 1-2 Supported partitioning for 2-node and 4-node complexes

1.4 IBM XA-64e third generation chipset

The x260, x366, and x460 use the third generation IBM XA-64e chipset. The architecture consists of the following components:

- ▶ One to four Xeon MP processors
- ▶ One Hurricane Memory and I/O Controller (MIOC)
- ▶ Two Calgary PCI Bridges

Figure 1-3 shows the block diagram of the X3 Architecture.

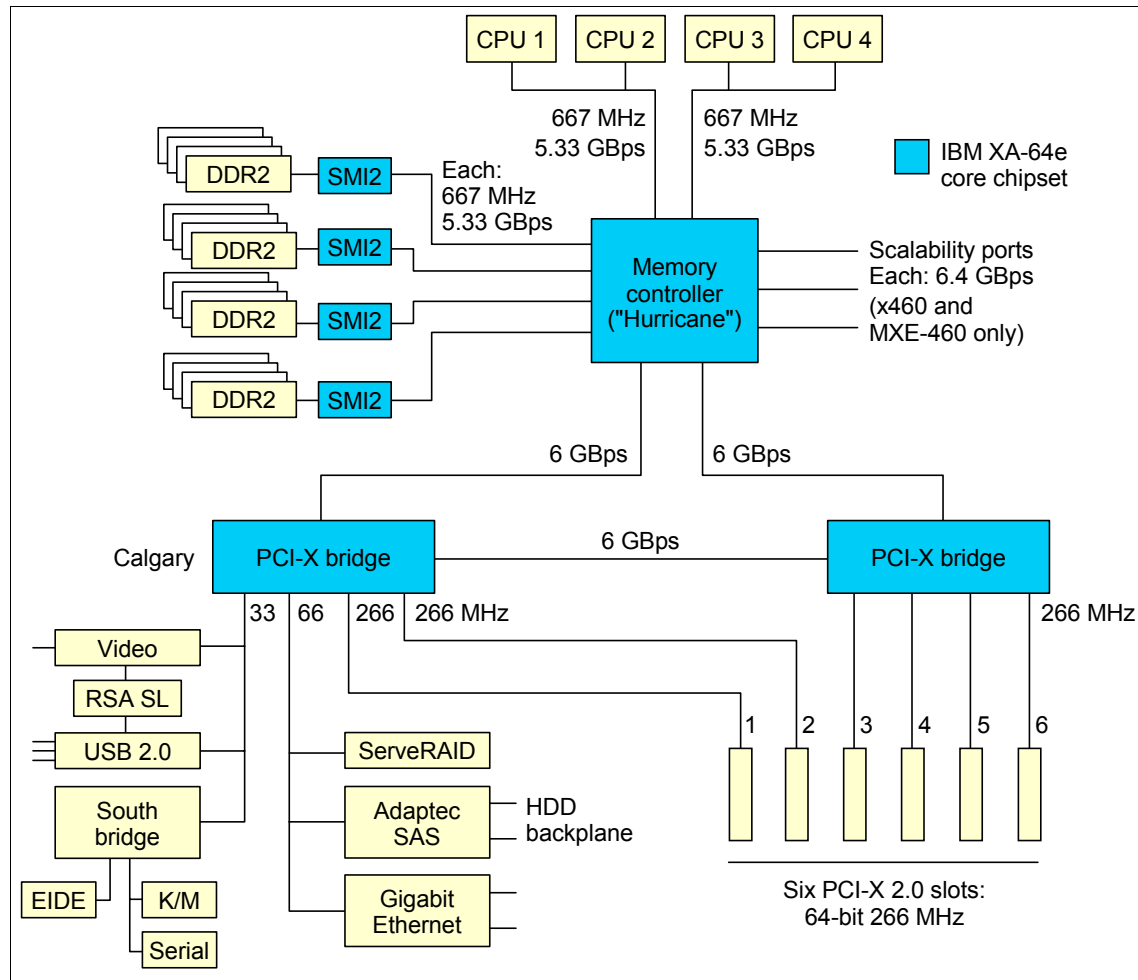


Figure 1-3 X3 Architecture system block diagram

Each memory port out of the memory controller has a peak throughput of 5.33 GBps. DIMMs are installed in matched pairs, two-way interleaving, to ensure the memory port is fully utilized. Peak throughput for each PC2-3200 DDR2 DIMM is 2.67 GBps. The DIMMs are run at 333 MHz to remain in sync with the throughput of the front-side bus.

Because there are four memory ports, spreading installed DIMMs across all four memory ports can improve performance. The four independent memory ports, or memory cards, provide simultaneous access to memory. With four memory cards installed, and DIMMs in each card, peak memory bandwidth is 21.33 GBps.

The memory controller routes all traffic from the four memory ports, two microprocessor ports, and the two PCI bridge ports. The memory controller also has embedded DRAM, which holds a snoop filter lookup table. This filter ensures that snoop requests for cache lines go to the appropriate microprocessor bus and not both of them, thereby improving performance.

One PCI bridge supplies four of the six 64-bit 266 MHz PCI-X slots on four independent PCI-X buses. The other PCI bridge supplies the other two PCI-X slots (also 64-bit, 266 MHz), plus all the onboard PCI devices, including the optional ServeRAID-8i and Remote Supervisor Adapter II SlimLine daughter cards.

1.5 Scalable implementation

In the standard single chassis 4-way configuration, the X3 Architecture servers act as an industry standard symmetric multiprocessor (SMP) system. Each processor has equal access to all system resources.

However, with multi-node x460 configurations, a NUMA-like architecture (non-uniform memory architecture) is implemented by connecting the scalability ports of each node together (see Figure 1-3 on page 8). These ports are directly connected to the Hurricane memory controller and allow high speed communication between processors located in different nodes. The ports act like hardware extensions to the CPU local buses. They direct read and write cycles to the appropriate memory or I/O resources, as well as maintain cache coherency between the processors.

In such multi-node configurations, the physical memory in each node is combined to form a single coherent physical address space. The resulting system has the property such that for any given region of physical memory, some processors are closer to it than other processors. Conversely, for any processor, some memory is considered local and other memory is remote.

Memory can be described in one of three ways, depending on the relation to a given processor:

- ▶ *Local*: Memory is in the same node as the processor.
- ▶ *Remote*: Memory is installed in another node that is directly connected via scalability cables.
- ▶ *Far*: In 8-node configurations, some nodes are not directly connected together (nodes 1 and 7, for example, as shown in Figure 3-9 on page 109). In such situations, memory that is two “hops” (two cables) away from a given processor is “far”.

The term NUMA is not completely correct since not only can memory be accessed in a non-uniform manner, but also I/O resources. PCI-X and USB devices may be associated with nodes. The exceptions to this situation are existing I/O devices, such as diskette and CD-ROM drives, which are disabled because the classic PC architecture precludes multiple copies of these existing items.

The key to this type of memory configuration is to limit the number of processors that directly access a piece of memory, thereby improving performance because of the much shorter queue of requests. The objective of the operating system is to ensure that memory requests be fulfilled by local memory whenever possible.

However, an application running on CPUs in node 1 may still need to access memory physically located in node 2 (a remote access). This access incurs longer latency because the travel time to access remote memory on another expansion module is clearly greater. Many people think this is the problem with NUMA. But this focus on latency misses the actual problem NUMA is attempting to solve: shorten memory request queues.

Another way to think about this is to ask yourself this question: You are checking out in your favorite grocery store, with a shopping cart full of groceries. Directly in front of you is a check-out lane with 10 customers standing in line but 20 meters to your left is another check-out lane with only two customers standing in line. Which would you go to? The check-out lane closest to your position has the lowest latency because you don't have far to travel. But the check-out lane 20 meters away has much greater latency because you have to walk 50 feet.

Clearly most people would walk the 20 meters and thereby incurring a delay to arrive at a check-out lane with only 2 customers instead of 10. We think this way because our experience tells us that the time waiting to check-out with 10 people ahead of us (the request queue) is far longer than the time needed to walk to the “remote” check-out lane (latency) and wait for only two people ahead.

This analogy clearly communicates the performance effects of queuing time versus latency. In a computer server, with many concurrent outstanding memory requests, we would gladly incur some additional latency (walking) to spread memory transactions (check-out process) across multiple memory controllers (check-out lanes), because this greatly improves performance by reducing the queuing time.

The performance implications of such a configuration are significant. It is essential that the operating system recognize which processors and ranges of memory are local, which are remote, and which are far.

Clearly, we do not want to walk 20 meters to a check-out lane that has 20 customers checking out, when one is directly in front of us with only two customers. So to reduce unnecessary remote access, the x460 maintain a table of data in the firmware called the Static Resource Allocation Table (SRAT). The data in this table is accessible by operating systems such as Windows Server 2003 (Windows 2000 Server does not support it) and current Linux kernels.

These modern operating systems attempt to allocate resources that are local to the processors being used by each process. So when a process and its threads start on node 1, all execution and memory access will be local to node 1. As more processes are added to the system, the operating system will balance them across the nodes. In this case, most memory accesses will be evenly distributed across the multiple memory controllers, reducing remote access, greatly reducing queuing delays, and improving performance.

1.6 Current models

The following models were announced at the time of publication.

1.6.1 x460 and MXE-460 models

The following are the currently available models of the x460 (Table 1-2) and the MXE-460 modular expansion enclosure (Table 1-3). The MXE-460 is similar to the x460, except the MXE-460 has no memory or processors standard and no DVD-ROM drive. The MXE-460 is used to form x460 multi-node configurations.

Table 1-2 x460 models

Model	Standard/max CPU	Cores/socket	L2 cache	L3 cache	Std/max memory
8872-1RY	2x 2.83 GHz Xeon MP / 4	Single-core	1 MB	4 MB	2 GB (2x 1 GB) / 64
8872-2RY	2x 3.00 GHz Xeon MP / 4	Single-core	1 MB	8 MB	2 GB (2x 1 GB) / 64
8872-3RY	2x 3.33 GHz Xeon MP / 4	Single-core	1 MB	8 MB	2 GB (2x 1 GB) / 64
8872-5RU	2x 2.67 GHz Xeon 7020 / 4	Dual-core	1+1 MB	0 MB	2 GB (2x 1 GB) / 64
8872-6RU	2x 3.00 GHz Xeon 7040 / 4	Dual-core	2+2 MB	0 MB	2 GB (2x 1 GB) / 64

The x460 supports two or four processors per node, up to eight nodes.

Table 1-3 lists the MXE-460. The type of microprocessors installed in the MXE-460s must match the primary x460 in the multi-node configuration.

Table 1-3 MXE-460 models

Model	Standard/max CPU	Cores/socket	L2 cache	L3 cache	Std/max memory
8874-1RY	0 / 4 (match the attached x460)	Single-core	1 MB	Varies	0 / 64 GB
8874-2RU	0 / 4 (match the attached x460)	Dual-core	Varies	0 MB	0 / 64 GB

Note: The U or Y in the model numbers is for countries in North and South America. For EMEA, substitute G, for example 1RG. For Asia-Pacific countries, the letter varies from country to country. Consult the announcement letter or the *xSeries Configuration and Option Guide*, found at:

<http://www.pc.ibm.com/support?page=SCOD-3ZVQ5W>

In addition to these models, IBM has also announced models that include Windows Server 2003, Datacenter Edition. These are described in 4.5, “Datacenter offerings” on page 147.

Both the x460 and MXE-460 support a maximum of 64 GB using 4 GB DIMMs in 16 sockets. The MXE-460 has no standard memory installed, but requires at least 2 GB of memory. The amount of memory installed does not have to match the other nodes in a multi-node complex, but for performance reasons, this is recommended.

Both the x460 and MXE-460 have six internal PCI-X 2.0 slots. Unlike the x445, the x460 does not support the attachment of an RXE-100 Remote Expansion Enclosure. The RSA II SlimLine and ServeRAID-8i options do not occupy any of these six PCI-X slots.

1.6.2 x366 models

The models shown in Table 1-4 are available.

Table 1-4 x366 models announced in March 2005

Model	CPU (std/max)	Cores/socket	L2 cache	L3 cache	Memory (std/max)
8863-1RY	1x Xeon MP 3.16 GHz / 4	Single-core	1 MB	None	2 GB (2x 1 GB) / 64
8863-2RY	1x Xeon MP 3.66 GHz / 4	Single-core	1 MB	None	2 GB (2x 1 GB) / 64
8863-3RU	1x 2.67 GHz Xeon 7020 / 4	Dual-core	1+1 MB	None	2 GB (2x 1 GB) / 64
8863-4RU	1x 3.00 GHz Xeon 7040 / 4	Dual-core	2+2 MB	None	2 GB (2x 1 GB) / 64

The x366 supports one, two, three, or four processors.

Both models support a maximum of 64 GB using 4 GB DIMMs in 16 sockets. To achieve the maximum, you will need to install three additional memory cards (one four-socket card is standard), remove the standard pair of 1 GB DIMMs, and insert 16 DIMMs.

The x366 has six internal PCI-X 2.0 slots. Unlike the x365, the x366 does not support the attachment of an RXE-100 Remote Expansion Enclosure. The RSA II SlimLine and ServeRAID-8i options do not use any of these six PCI-X slots.

Express models: Certain countries also offer preconfigured systems called Express models. These are preconfigured with more processors, memory, disk, or a ServeRAID controller. The advantage of these models is fewer part numbers to order.

1.6.3 x260 models

The following models of the x260 (Table 1-5) were announced in July 2005.

Table 1-5 x260 models announced in July 2005

Model	Configuration	Standard/max CPU	L2 cache	Standard/max memory
8865-12U	Tower	1x 3.16 GHz Xeon MP / 4	1 MB	1 GB (2x 512 MB) / 64 GB
8865-1SU	Rack	1x 3.16 GHz Xeon MP / 4	1 MB	1 GB (2x 512 MB) / 64 GB
8865-11U	Tower	1x 3.16 GHz Xeon MP / 4	1 MB	2 GB (2x 1 GB) / 64 GB
8865-1RU	Rack	1x 3.16 GHz Xeon MP / 4	1 MB	2 GB (2x 1 GB) / 64 GB
8865-21U	Tower	1x 3.66 GHz Xeon MP / 4	1 MB	2 GB (2x 1 GB) / 64 GB
8865-2RU	Rack	1x 3.66 GHz Xeon MP / 4	1 MB	2 GB (2x 1 GB) / 64 GB

The x260 supports one to four processors.

Note: The U in the model numbers is for countries in North and South America. For EMEA, substitute G for the U, for example, 1RG. For Asia-Pacific countries, the letter varies from country to country. Consult the announcement letter or the *xSeries Configuration and Option Guide*, found at: <http://www.pc.ibm.com/support?page=SCOD-3ZVQ5W>

All of the x260 series support a maximum of 64 GB using 4 GB DIMMs in 16 sockets. x260s have four standard DIMMS installed, but have 12 expandable sockets.

The x260 has six internal PCI-X 2.0 slots. The optional RSA II SlimLine and optional ServeRAID-8i adapters do not occupy any of these six PCI-X slots.

For clients who purchase a tower model and later wish to install the server into a suitable rack, a Tower To Rack Conversion Kit, part number 32R0719, is also available and provides the hardware needed to convert the tower unit to a 7U-high, rack-mounted server. It includes slides, cable management arm, latches, and miscellaneous mounting hardware.

Express models: Certain countries also offer preconfigured systems called Express models. These are preconfigured with various combinations of more processors, memory, disk, and a ServeRAID controller. The advantage of these is fewer part numbers to order.

1.7 Front and rear layouts

The three X3 Architecture servers share many components. This section describes the components accessible from the front and rear of each server.

1.7.1 x460 and MXE-460

The front panel of the x460 and MXE-460 is shown in Figure 1-4.




Figure 1-4 Front panel of the x460 and MXE-460

The rear panel of the x460 and MXE-460 is shown in Figure 1-5.




Figure 1-5 Rear panel of the x460 and MXE-460

1.7.2 x366

Figure 1-6 shows the x366 and major components on the front of the unit.




Figure 1-6 Front panel of the x366

The rear panel of the x366 is shown in Figure 1-7 on page 17.




Figure 1-7 Rear view of the x366

1.7.3 x260

Figure 1-8 shows the front panel of the x260.




Figure 1-8 Front panel of the x260

Figure 1-9 shows the rear panel of the x260.




Figure 1-9 Rear panel of the x260

1.8 Processors

The X3 Architecture servers all use Intel Xeon processors with EM64T extensions.

The x460 models use either *Potomac* single-core processors or *Paxville* dual-core processors. Models of the x460 have two processors installed. Two or four processors are supported. Installed processors must be identical in speed and cache size. You can connect MXE-460s to the x460 to form larger configurations. The x460 and each MXE-460 must have four identical processors installed. Multi-node capabilities are discussed in detail in 3.3, “Multi-node configurations” on page 104.

The x366 uses either *Cranford* single-core processors or *Paxville* dual-core processors. The x260 uses *Cranford* processors. Models of the x260 and x366 have one processor installed. One, two, three, or four processors are supported. Installed processors must be identical in speed and cache size.

Dual-core upgrade: The Dual Core X3 Upgrade Kit, part 39Y6580, can be used to upgrade an existing single-core x366 or x460 system to dual-core capability. The kit does not include the dual-core processors.

The processors are accessible from the front of the server on a sliding tray. The tray can be pulled out once the memory cards and fans are removed (see Figure 1-10).




Figure 1-10 Processor tray (x460 shown)

The VRMs for processors 1 and 2 are integrated on the microprocessor board; the VRMs for processors 3 and 4 come with the processor options.

See 3.1, “Processor subsystem” on page 90 for a further discussion about what you should consider before implementing a solution.

The *Cranford* and *Paxville* processors have two levels of cache on the processor die:

- ▶ L2 cache is 1 MB in size. The L2 cache implements the Advanced Transfer Cache technology, which means L2-to-processor transfers occur across a 256-bit bus in only one clock cycle.
- ▶ L1 execution trace cache is used to store micro-operations (that is, decoded executable machine instructions); it serves those to the processor at rated speed. This additional level of cache saves decode time on cache hits.

The *Potomac* Xeon MP processor used in the x460 and MXE-460 has three levels of cache on the processor die:

- ▶ The L3 cache is 4 MB or 8 MB, depending on the processor.
- ▶ The L2 cache is 1 MB in size. The L2 cache implements the Advanced Transfer Cache technology, which means L2-to-processor transfers occur across a 256-bit bus in only one clock cycle.
- ▶ The L1 execution trace cache is used to store micro-operations and decoded executable machine instructions. It serves them to the processor at rated speed. This additional level of cache saves decode time on cache hits.

Key features of both processors are described in the following subsections.

1.8.1 Models with dual-core processors

The new Paxville dual-core processors are a concept similar to a two-way system except that the two processors, or *cores*, are integrated into one silicon die. This brings the benefits of two-way SMP with less power consumption and faster data throughput between the two cores. To keep power consumption down, the resulting core frequency is lower, but the additional processing capacity means an overall gain in performance.

Figure 1-11 compares the basic building blocks of the Xeon MP single-core processor (Potomac) and dual-core processor (Paxville).

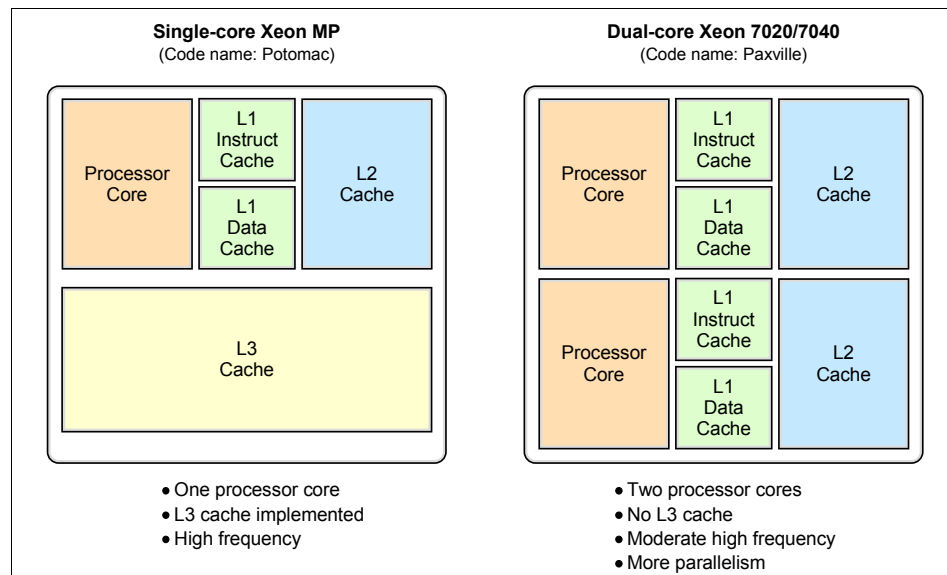


Figure 1-11 Feature of single core and dual core processors

In addition to the two cores, the dual-core processor has separate L1 instruction and data caches for each core as well as separate execution units (integer, floating point, and so on), registers, issue ports, and pipelines for each core. A dual core processor achieves more parallelism than Hyper-Threading Technology because these resource are not shared between the two cores. Estimates are that there is a 1.2 to 1.5 times improvement comparing the dual-core Xeon MP with current single-core Xeon MP.

With double the number of cores for the same number of sockets, it is even more important that the memory subsystem is able to meet the demand for data throughput. The 21 GBps peak throughput of the X3 Architecture of the x460 with four memory cards is well suited to dual-core processors.

1.8.2 667 MHz front-side bus

The Pentium® III Xeon processor in older servers had a 100 MHz front-side bus that equated a burst throughput of 800 MBps. With protocols such as TCP/IP, this has been shown to be a bottleneck in high-throughput situations.

Both the Cranford and Potomac Xeon MP improves on this by using two 133 MHz clocks, out of phase with each other by 90°, and using both edges of each clock to transmit data. This is shown in Figure 1-12.




Figure 1-12 Quad-pumped front-side bus

This increases the performance of the front-side bus without the difficulty of high-speed clock signal integrity issues. Because the bus is eight bytes wide, the end result is an effective burst throughput of 5.33 GBps, which can have a substantial impact especially on TCP/IP-based LAN traffic.

1.8.3 Hyper-Threading

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

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

The basic layout of a Hyper-Threading-enabled microprocessor is outlined in Figure 1-13, where you can clearly see that only the components for the architectural state of the microprocessor have doubled.

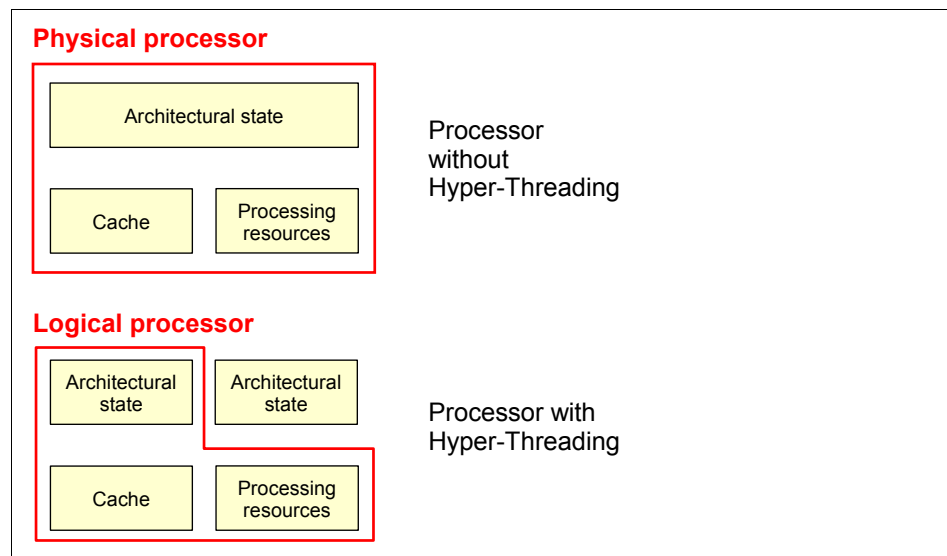


Figure 1-13 Architectural differences associated with Hyper-Threading

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

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

- ▶ Two physical processors: up to about 25% performance gain
- ▶ Four physical processors: up to about 15% performance gain
- ▶ Eight physical processors: up to about 10% performance gain

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

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

Note that Microsoft licensing of the Windows 2000 Server operating systems is by number of processors: four-way for Server, eight-way for Advanced Server, and 32-way for Datacenter Server. Therefore, the appearance of twice as many logical processors can potentially affect the installation of the operating system. See 4.1, “Supported operating systems” on page 136 for details.

Windows Server 2003 understands the concept of physical processors versus logical processors.

1.8.4 Intel 64-bit Extension Memory Technology (EM64T)

First introduced in the Xeon DP *Nocona* processor, EM64T is a 64-bit extension to the industry standard IA32 32-bit architecture. EM64T adds:

- ▶ A set of new 64-bit general purpose registers (GPR)
- ▶ 64-bit instruction pointers
- ▶ The ability to process data in 64-bit chunks

Even though the names of these extensions suggest that the improvements are simply in memory addressability, Intel EM64T is, in fact, a fully functional 64-bit processor.

A 64-bit processor is defined as one that is able to address 64-bits of virtual address space. A 64-bit processor can store data in 64-bit format and perform arithmetic operations on 64-bit operands. In addition, a 64-bit processor has general purpose registers (GPRs) and arithmetic logical units (ALUs) that are 64-bit wide.

There are now three 64-bit implementations in the “Intel compatible processor” marketplace:

- ▶ Intel IA64, as implemented on the Itanium® 2 processor
- ▶ Intel EM64T, as implemented on the Xeon DP Nocona and Irwindale processors and Xeon MP Cranford and Potomac processors.
- ▶ AMD AMD64, as implemented on the Opteron processor

The discussion whether or not EM64T really are 64-bit processor can be confusing to some because Intel calls its architecture “Extended Memory 64 Technology”. We know that it extends the IA32 instruction set. Therefore, are EM64T processors “real” 64-bit chips? The answer is yes. When these processors operate in 64-bit mode, the addresses are 64-bit, the GPRs are 64 bits wide, and the ALUs are able to process data in 64-bit chunks. Therefore, these processors are full-fledged 64-bit processors in this mode.

The Xeon MP with EM64T extends the previous Xeon MP in three ways:

- ▶ 64-bit addressing and 64-bit registers
- ▶ Additional 128-bit Streaming SIMD Extensions (SSE) registers and 64-bit general purpose registers
- ▶ Double precision (64-bit) integer support

There are three distinct operation modes available in EM64T:

- ▶ 32-bit mode

The first and, in the near future, probably most widely used mode will be 32-bit mode. In this mode, any EM64T processors will act just like any other IA32 compatible processor. You can install your 32-bit OS on such a system and run 32-bit applications; however, you will not be able to make use of the new features, such as the flat memory addressing above 4 GB or the additional 64-bit General Purpose Registers (GPRs). 32-bit applications will run just as fast as they would on any current 32-bit processor.

Most of the time, IA32 applications will run even faster, since there are numerous other improvements that boost performance regardless of the maximum address size.

- ▶ Compatibility mode

The second mode supported by the EM64T is compatibility mode, which is an intermediate mode of the full 64-bit mode described below. In order to run in compatibility mode, you will need to install a 64-bit operating system and 64-bit drivers. If a 64-bit OS and drivers are installed, the Xeon processor will be enabled to support a 64-bit operating system with both 32-bit applications or 64-bit applications.

Compatibility mode gives you the ability to run a 64-bit operating system while still being able to run unmodified 32-bit applications. Each 32-bit application will still be limited to a maximum of 4 GB of physical memory. However, the 4 GB limit is now imposed on a per-process level, not at a system-wide level. This means that every 32-bit process on this system gets its very own 4 GB of physical memory space (assuming sufficient physical memory is installed).

This is a big improvement compared to IA32, where the operating system kernel and the application have to share 4 GB of physical memory and

applications usually get no more than 3 GB of memory. 32-bit applications modification does not require, but BIOS updates will be needed

Additionally, compatibility mode does not support the virtual 8086 mode, so real-mode existing applications are not supported. 16-bit protected mode applications are, however, supported.

► Full 64-bit mode

The final mode is the full 64-bit mode. Intel refers to it as IA-32e mode. This mode is when a 64-bit operating system and 64-bit application are used. In the full 64-bit operating mode, an application can have a virtual address space of up to 40-bits (that equates to 1 TB of addressable memory). The amount of physical memory will be determined by the server chipset, how many DIMM slots the server has, and the maximum DIMM capacity supported and available at the time.

Applications that run in full 64-bit mode will get access to the full physical memory range (depending on the operating system), and will also get access to the new GPRs as well as to the expanded GPRs. However, it is important to understand that this mode of operation requires not only a 64-bit operating system (and of course 64-bit drivers), but also requires a 64-bit application that has been recompiled to take full advantage of the various enhancements of the 64-bit addressing architecture

Figure 1-14 gives an overview of the supported modes.

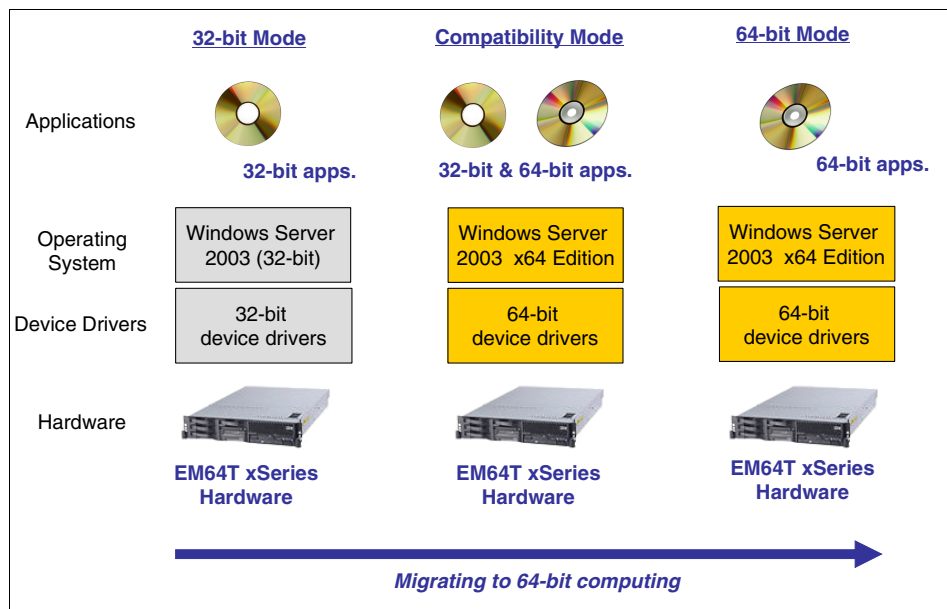


Figure 1-14 Intel 64-bit Extension Technology Supported Modes

For more information about EM64T, see:

<http://www.intel.com/technology/64bitextensions/>

1.8.5 The benefit of 64-bit computing

In the same way that 16-bit processors and 16-bit applications are no longer used in this space, it is likely that at some point in the future 64-bit processors and applications will replace their 32-bit counterparts.

Processors using the EM64T architectures are making this transition very smooth by offering 32-bit and 64-bit modes. This means that the hardware support for 64-bit may likely be in place before you upgrade or replace your applications with 64-bit versions. The X3 Architecture servers all use EM64T-enabled processors.

The question you should be asking is whether the benefit of 64-bit processing is worth the effort of upgrading or replacing your 32-bit applications. The answer is that it depends on the application. The following are examples of applications that will benefit from 64-bit computing:

- ▶ Encryption applications: Most encryption algorithms are based on very large integers and would benefit greatly with the use of 64-bit GPRs and ALUs.
- ▶ Scientific applications: Integer-based scientific applications will benefit. Floating-point operations do not benefit from the larger integer size since floating-point registers are already 80 or 128 bits wide even in 32-bit processors.
- ▶ Applications requiring more than 4 GB of memory: This is the biggest advantage of 64-bit computing for commercial applications: the flat, potentially massive, address space.

Enterprise applications, such as databases, are currently implementing Page Addressing Extensions (PAE) and Addressing Windows Extensions (AWE) addressing schemes to access memory above the 4 GB limit imposed by 32-bit address limited processors. These addressing extension schemes support access to memory up to 64 GB in size. One constraint with PAE and AWE, however, is that memory above 4 GB can only be used to store data; it cannot be used to store or execute code. So these addressing schemes only make sense for applications such as databases, where large data caches are needed.

In contrast, a 64-bit virtual address space provides for direct access of up to 2 Exabytes (EB). And even though we call these processors 64-bit, none of the current 64-bit processors actually supports the full 64 bits of physical memory addressing, simply because this is such an enormous amount of memory.

Table 1-6 lists the memory addressability by processor. These values are the upper limits imposed by the processors. Memory addressing is usually limited further by the chipset implemented in the server and by the DIMM technology available.

Table 1-6 Memory supported by processors

Processor	Flat addressing	Addressing with PAE-36
Intel Xeon MP “Gallatin” (32-bit)	4 GB (32-bit)	64 GB
Intel EM64T “Nocona” (64-bit)	64 GB (36-bit)	64 GB in compatibility mode
Intel EM64T “Cranford” (64-bit)	64 GB (36-bit)	64 GB in compatibility mode
Intel EM64T “Potomac” (64-bit)	1 TB (40-bit)	64 GB in compatibility mode
Intel EM64T “Paxville” (64-bit)	1 TB (40-bit)	64 GB in compatibility mode
Intel Itanium (64-bit)	18 TB (44-bit)	Not applicable
AMD Opteron (64-bit)	256 TB (48-bit)	64 GB in compatibility mode
Intel Itanium 2 (64-bit)	1024 TB (50-bit)	Not applicable

1.9 Xcel4v cache

The Xcel4v dynamic server cache serves two purposes in the X3 Architecture servers:

- ▶ As a single, 4-way server (x260, x366 or x460), the Xcel4v and its embedded DRAM (eDRAM) is used as a snoop filter to reduce traffic on the front side bus. It stores a directory of all processor cache lines to minimize snoop traffic on the dual front side buses and minimize cache misses.
- ▶ When the x460 is configured as a multi-node server, this technology dynamically allocates 256 MB of main memory in each node for use as an L4 cache directory and scalability directory. In a 32-way configuration, this means there will be 2 GB of Xcel4v cache.

With advances in chip design, IBM has now reduced the latency of main memory to below that of the Xcel4 cache in the x445. In other words, the time it takes to access data directly from memory is almost as fast as accessing it from L3. As a result, on a four-way system such as the x260 and x366, there is little/no need for either a L3 cache or L4 cache.

Since the L3 cache is inline, when cache misses do occur, it adds significant overhead to memory access. The L3 cache rate has to be very high for it to keep up with the 3.66 GHz processor. In most server applications with multiple users, the threads competing for L3 cache generate a lower hit rate, and the latency of the L3 drops performance. The same applies to any L4 cache.

As a result, there is no performance benefit in implementing either an L3 or L4 cache on the four-way x260 or x366. For these reasons, there is 0 MB of Xcel4v cache on the x260 and x366 servers.

1.10 Memory

Memory in the X3 Architecture servers is PC2-3200 ECC DDR2 DIMMs. Standard memory is as follows:

- ▶ x260: 1 GB or 2 GB of RAM standard (using 512 MB DIMMs)
- ▶ x366: 2 GB standard (using 512 MB DIMMs)
- ▶ x460: 2 GB standard (using 512 MB DIMMs)
- ▶ MXE-460: None standard

Memory is implemented in the servers using memory cards. Each server supports up to four memory cards, and each card has four DIMM sockets. Memory cards are standard as follows:

- ▶ x260: One card standard
- ▶ x366: One card standard
- ▶ x460: Two cards standard
- ▶ MXE-460: Two cards standard

Using 4 GB DIMMs in every socket, a total of 16 DIMMs, the servers can hold 64 GB of RAM.

The memory is two-way interleaved, meaning that memory DIMMs are installed in pairs. There are four ports to memory, with each supporting up to 5.33 GBps data transfers.

The DIMMs operate at 333 MHz instead of 400 MHz, as per the PC2-3200 spec, so that throughput is 2.67 GBps, or 333 MHz x 8 bytes. At 2.67 GBps and two-way interleaving, the throughput matches that of the front-side bus at 5.3 GBps so that bus transfers remain in sync.

There are a number of advanced features implemented in the X3 Architecture memory subsystem, collectively known as *Active Memory*:

► Memory ProteXion

The Memory ProteXion feature (also known as *redundant bit steering*) provides the equivalent of a hot-spare drive in a RAID array. It is based in the memory controller, and it enables the server to sense when a chip on a DIMM has failed and to route the data around the failed chip.

Normally, 128 bits out of every 144 are used for data and the remaining 16 bits are used for ECC functions. However, X3 Architecture needs only 12 bits to perform the same ECC functions, thus leaving four bits free. These four bits are equivalent to an x4 memory chip on the DIMM that Memory ProteXion uses. In the event that a chip failure on the DIMM is detected by memory scrubbing, the memory controller can reroute data around that failed chip through these spare bits.

It can do this automatically without issuing a Predictive Failure Analysis® (PFA) or light path diagnostics alert to the administrator, although an event is logged to the service processor log. After the second DIMM failure, PFA and light path diagnostics alerts would occur on that DIMM as normal.

Note: In BIOS, there is a Memory Array setting in Advanced Settings. When you set it to **High Performance Memory Array**, BIOS reconfigures the server for maximum performance at the expense of some fault tolerance features, including Memory ProteXion. For a production environment, we recommend that you not select this setting, thereby keeping your system protected from memory failures with Memory ProteXion.

► Memory scrubbing

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

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

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

If the error is not recoverable, then memory scrubbing sends an error message to the light path diagnostics, which then turns on the proper lights and LEDs to guide you to the damaged DIMM. If memory mirroring is enabled, then the mirrored copy of the data from the damaged DIMM is used until the system is powered down and the DIMM replaced.

▶ Memory mirroring

Memory mirroring is roughly equivalent to RAID 1 in disk arrays, in that usable memory is halved and a second copy of data is written to the other half. If eight GB is installed, then the operating system sees four GB once memory mirroring is enabled. It is disabled in the BIOS by default. Because all mirroring activities are handled by the hardware, memory mirroring is operating system independent.

When memory mirroring is enabled, certain restrictions exist with respect to placement and size of memory DIMMs and the placement and removal of memory cards.

▶ Chipkill™ memory

Chipkill is integrated into the XA-64e chipset, so it does not require special Chipkill DIMMs and is transparent to the operating system. When combining Chipkill with Memory ProteXion and Active Memory, X3 Architecture provides very high reliability in the memory subsystem.

When a memory chip failure occurs, Memory ProteXion transparently handles the rerouting of data around the failed component, as described above. However, if a further failure occurs, the Chipkill component in the memory controller reroutes data. 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. One of these additional failures can be handled for each memory port for a total of four Chipkill recoveries.

▶ Hot-add and hot-swap memory

The X3 Architecture servers support the replacing of failed DIMMs while the server is still running. This hot-swap support works in conjunction with memory mirroring. The server also supports adding additional memory while the server is running. Adding memory requires operating system support.

These two features are mutually exclusive. Hot-add requires that memory mirroring be disabled and hot-swap requires that memory mirroring be enabled.

In addition, to maintain the highest levels of system availability, if a memory error is detected during POST or memory configuration, the server can automatically disable the failing memory bank and continue operating with reduced memory

capacity. You can manually re-enable the memory bank after the problem is corrected by using the Setup menu in the BIOS.

Memory mirroring, Chipkill, and Memory ProteXion provide multiple levels of redundancy to the memory subsystem. Combining Chipkill with Memory ProteXion allows up to two memory chip failures for each memory port on the system, for a total of eight failures sustained. For example:

1. The first failure detected by the Chipkill algorithm on each port does not generate a light path diagnostics error because 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.

1.11 Serial Attached SCSI

Serial Attached SCSI (SAS) is the logical evolution of SCSI. SAS uses much smaller interconnects than SCSI, while offering SCSI compatibility, reliability, performance, and manageability. In addition, SAS offers longer cabling distances, smaller form factors, and greater addressability.

SAS 1.0 technology is replacing Ultra320 SCSI in SCSI and RAID controllers. The X3 Architecture family of servers have a disk subsystem that is based on an Adaptec AIC-9410 SAS controller. The servers do not support RAID as standard, but these servers also support the addition of the ServeRAID-8i daughter card, which in turn supports a variety of RAID levels and stripe sizes up to 512 KB.

Beyond the upgrades in I/O processor and memory speeds, SAS-based products will differ from SCSI-based products in the following ways:

- ▶ Higher bandwidth

Ultra320 SCSI supports 320 MBps of bandwidth per channel. SAS 1.0 supports three Gbps, approximately 300 MBps, of bandwidth for each port. So while the two Ultra320 SCSI channels of the ServeRAID-6M can potentially support 640 MBps of bandwidth, the onboard SAS controller with its eight ports could support up to 24 Gbps, approximately 2.4 GBps, of bandwidth. Therefore, bandwidth will be limited by PCI-X or PCI Express bus speeds.

- ▶ Greater drive support

SCSI-based products support 14 drives for each channel. By cascading drive enclosures, SAS-based products will support up to 72 drives for four ports.

The AIC-9410 SAS controller is central to the SAS implementation on the X3 Architecture servers. It has eight SAS ports, each supporting a maximum transfer rate of 3 Gbps.

Depending on the server, the drives are either directly attached to the SAS ports (x366 and x460) or they are connected via SAS Expanders (x260). This is shown in the following two figures.

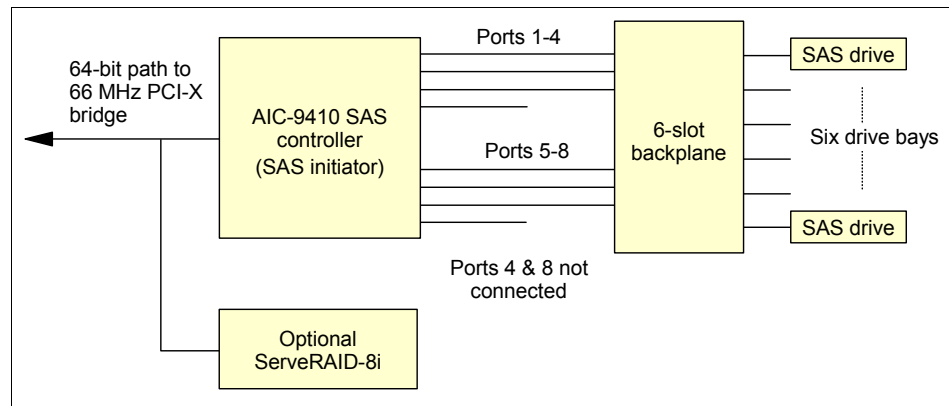


Figure 1-15 SAS block diagram (x460 and x366)

With the x336 and x460 (Figure 1-15), only six of the eight ports are used from the SAS controller and these are connected directly to the six drive bays in the server via the backplane.

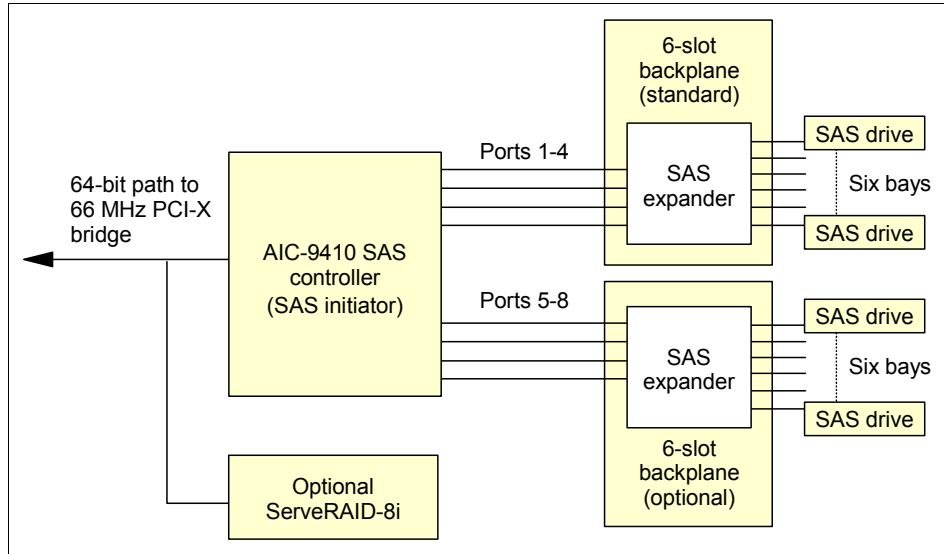


Figure 1-16 SAS block diagram (x260)

For the x260 (Figure 1-16), four ports from the SAS controller chip are fed into the expander and the signals are expanded out to six devices, similar in concept to a router. The x260 comes standard with only one of the two drive backplanes. In such a configuration, only ports 1-4 are used. Ports 5-8 are used only when the second backplane is installed.

In terms of performance, there is no benefit in inserting drives in any particular order (for example, to spread the drive load across ports). We recommend you simply install the drives sequentially. The peak transfer rate is 3 Gbps per port or 12 Gbps per SAS expander; however, it is more likely that the peak drive transfer rate will be less than this value.

All three servers support hot-swap drives. The number and type of drive bays in each server are as follows:

- ▶ x260: Six 3.5" hot-swap bays standard with another six optional
- ▶ x366: Six 2.5" hot-swap bays
- ▶ x460 and MXE-460: Six 2.5" hot-swap bays

For details on supported drives, see 3.5.3, "Internal storage" on page 124.

When the optional ServeRAID-8i is installed, you can see from the above diagrams that the flow of data still goes through the SAS controller. With the ServeRAID-8i installed, the adapter controls the RAID data layout and instructs the controller to write the data to the requested drives. For details about the ServeRAID-8i, see 3.5.2, "ServeRAID-8i" on page 123.

For drivers, if you do not have the ServeRAID-8i installed, then use the SAS driver. If you have a ServeRAID-8i installed, then use these ServeRAID-8i driver. Both are available from the xSeries driver matrix:

<http://www.pc.ibm.com/support?page=MIGR-4JTS2T>

1.12 PCI-X slots

As shown in Figure 1-17, there are six full-length hot-swap PCI-X 2.0 slots internal to the x260 and all are vacant in the standard models. The six slots all support hot-plug PCI-X 3.3 V, 32-bit or 64-bit PCI and PCI-X 2.0 adapters.




Figure 1-17 PCI-X slots

All six slots have the following characteristics:

- ▶ Separate bus from the other slots and devices. This means the speed of the adapter does not affect the other adapters.
- ▶ PCI-X 2.0
- ▶ 266 MHz, each supporting lower speed adapters
- ▶ 64-bit, each supporting 32-bit adapters as well
- ▶ 3.3 V

One additional dedicated PCI slot is reserved for the ServerRAID-8i adapter.

The PCI subsystem also supplies these I/O devices:

- ▶ Adaptec AIC-9410 Serial-attached SCSI (SAS) controller
- ▶ Broadcom dual port 5704 10/100/1000 Ethernet
- ▶ ATI 7000-M video controller with 16 MB video memory
- ▶ Three USB ports, one on the front panel, two on the rear
- ▶ Remote Supervisor Adapter II SlimLine adapter (optional in a dedicated socket on the I/O board)
- ▶ EIDE interface for the CD-ROM drive
- ▶ Serial port

Note: There is no parallel port on the x260. For parallel port connections, use the NetVista USB Parallel Printer Cable, part number 19K4164.

Here is some further configuration information:

- ▶ Video adapters are not supported.
- ▶ The PCI slots support 3.3V adapters only. 5V adapters, such as the Adaptec 2944UW, are not supported.
- ▶ The system scans PCI-X slots to assign system resources. The system attempts to start the first device found, with the search order as follows:
 - a. DVD-ROM.
 - b. Integrated dual Gigabit Ethernet controller.
 - c. Integrated SAS devices.
 - d. Internal PCI and PCI-X slots (in the order 1, 2, 3, 4, 5, and 6).
 - e. If the x460 is attached to a MXE-460 Modular Expansion Enclosure, the ordering continues as 7, 8, 9, 10, 11, 12, and so forth.

Restriction: The x460 does not support connectivity to an external PCI slot enclosure such as the RXE-100.

1.13 Remote Supervisor Adapter II SlimLine

The x460 and MXE-460 both have the Remote Supervisor Adapter II SlimLine service processor (shown in Figure 1-18 on page 36) as a standard component, while it is an optional upgrade in the x260 and x366. This adapter is installed in a dedicated slot, and provides the similar functionality as the Remote Supervisor Adapter II PCI option available for other xSeries servers. The processor can be configured using the service processor Web interface.




Figure 1-18 Remote Supervisor Adapter II SlimLine

Key features of the Remote Supervisor Adapter II SlimLine include:

- ▶ IBM ASIC with an integrated PowerPC® 405 core executing at 200 MHz
- ▶ 16 MB SDRAM and 4 MB flash ROM
- ▶ System-independent graphical console redirection
 - Built-in video compression hardware eliminates drivers.
 - Graphics response, up to five times faster than with the original RSA, makes monitoring and control more efficient.
 - System-independent installation eliminates the need to install service processor drivers, helps save IT staff time, and reduces installation complexity.
- ▶ Remote control with remote diskette and CD-ROM drive support
 - Enables remote booting and software loading of the server for application and operating system installation and updates.
 - Performs configuration remotely, helping to save IT time and money by reducing on-site presence and server downtime.
- ▶ Scriptable command-line interface and text-based serial console redirect
 - Command-line interface supports program control of server management functions using scripts.

- Serial text redirect provides access to text-mode BIOS and text-based system consoles, such as Linux, NetWare, and Windows EMS (Emergency Management Services).
- Program control of text-based console uses scripts.
- Uses point-to-point protocol support.
- ▶ User authentication and authority features
 - User IDs, passwords, and login permission attributes can be stored in an LDAP server.
 - Enhanced user authority levels set the access rights for users to match job responsibilities for managing your xSeries servers.
 - Secure Sockets Layer (SSL) encrypts the data transmitted between the LDAP server and the Remote Supervisor Adapter II.
- ▶ Investment protection
 - Integrates with IBM Director and Director Agent.

1.14 Baseboard Management Controller

The Baseboard Management Controller (BMC) is a small, independent micro-controller used to perform low-level system monitoring and control functions, as well as remote IPMI interface functions. It uses multiple I2C bus connections to communicate out-of-band with other onboard devices. The BMC provides environmental monitoring for the server. If environmental conditions exceed thresholds or if system components fail, the BMC lights the light path diagnostic LEDs to help you diagnose the problem and also records the error in the BMC system event log.

The BMC functions are as follows:

- ▶ Initial system check at A/C on

The BMC monitors critical I2C devices in standby power mode to determine if the system configuration is safe for power on.
- ▶ BMC Event log maintenance

The BMC maintains and updates an IPMI-specified event log in non-volatile storage. Critical system information is recorded and made available for external viewing.
- ▶ System power state tracking

The BMC monitors the system power state and logs transitions into the system event log.

- ▶ System initialization
The BMC has I2C access to certain system components that might require initialization before power-up.
- ▶ System software state tracking
The BMC monitors the system and reports when the BIOS and POST phases are complete and the operating system has booted.
- ▶ System event monitoring
During runtime, the BMC continually monitors critical system items, such as fans, power supplies, temperatures, and voltages. The system status is logged and reported to the service processor, if present.
- ▶ System fan speed control
The BMC monitors system temperatures and adjusts the fan speed accordingly.

The BMC also provides the following remote server management capabilities through the OSA SMBridge management utility program (see 5.5, “OSA SMBridge utility” on page 194 for details).

1.15 Light path diagnostics and PFA

To limit the client’s need to slide the server out of the rack to diagnose problems, a light path diagnostics panel is located at the front of the x260, x366, and x460. This panel slides out from the front of the server so the client can view all light path diagnostics-monitored server subsystems. In the event that maintenance is required, the client can slide the server out of the rack and, using the LEDs, find the failed or failing component.

Light path diagnostics can monitor and report on the health of microprocessors, main memory, hard disk drives, PCI-X and PCI adapters, fans, power supplies, VRMs, and the internal system temperature (see Figure 1-19).




Figure 1-19 Light path diagnostic panel

PFA is a mechanism developed by IBM that periodically measures critical components, such as hard drives, power supplies, fans, processors, and memory to reduce unscheduled system downtime.

If a predefined threshold of one of these components is exceeded, for example, the number of single bit errors on a memory module, a PFA alert will be issued. This enables you to replace the component before a failure actually occurs.

In case a PFA alert is raised, the Error LED on the light path diagnostic panel will be lit and the RSA logs an event. The RSA can also be configured to forward alerts via e-mail or to an IBM Director server.



Positioning

IBM @server X3 Architecture, the third generation of the IBM Enterprise X-Architecture technology, delivers high performance with 64-bit memory addressability, outstanding availability, and manageability required for the next generation of industry standard servers.

This chapter positions the servers in the marketplace and describes the target applications for them. Topics covered are:

- ▶ 2.1, “Positioning the servers” on page 42
- ▶ 2.2, “Focus market segments and target applications” on page 52
- ▶ 2.3, “Server consolidation” on page 56
- ▶ 2.4, “XpandOnDemand” on page 64
- ▶ 2.5, “Positioning the x460 and the MXE-460” on page 66
- ▶ 2.6, “Application scalability” on page 68
- ▶ 2.7, “Scale-up versus scale-out” on page 82

2.1 Positioning the servers

X3 Architecture delivers a formidable combination of 64-bit performance, mission-critical availability, unmatched modular scalability, and investment protection not previously available in the industry-standard x86 server market. With extensive high-performance server chipset development experience, IBM is uniquely positioned to offer a robust and powerful server, offering innovation that delivers real business IT results.

IBM has invested over \$100M bringing the new chip set and a completely refreshed set of products to market (the xSeries 260, xSeries 366, and the flagship xSeries 460), reinforcing the future of 64-bit x86 computing and the need for single system image scalability (known as *scale-up*).

X3 architecture servers were specifically designed for the demanding commercial transaction processing marketplace where clients value fast performance and rock-solid reliability. IBM engineers expanded on the XA-32 chipset designs of the x440 (first generation) and x445 (second generation) to produce a highly scalable design, as shown in Table 2-1.

Table 2-1 Enterprise X-Architecture generations

Item	XA-32 First generation	XA-32 Second generation	XA-64e Third generation
First available	2002	2003	2005
Servers	x360, x440	x365, x445	x260, x366, x460
Processors	Xeon MP, Xeon DP Clock speeds to 2.0 GHz	Xeon MP, Xeon DP Clock speeds to 3.0 GHz	EM64T 64-bit Clock speeds to 3.66 GHz Dual core processors
Cache L2	512 KB	512 KB	1 MB or 2 MB per core
Cache L3	512 KB - 2 MB	1 - 4 MB	0 MB (x366, x260, dual core) 4 or 8 MB (single-core x460)
Cache L4	32 MB XceL4	64 MB XceL4 12% lower latency	256 MB XceL4v cache for multi-node only (x460)
Front side bus	Single-bus 32-bit Architecture, 400 MHz	Single-bus 32-bit Architecture, 400 MHz	Dual-bus x86-64 Architecture, 667 MHz

Item	XA-32 First generation	XA-32 Second generation	XA-64e Third generation
Memory	<ul style="list-style-type: none"> ▶ 6.4 GBps peak bandwidth ▶ Memory Mirroring ▶ ChipKill ▶ Memory ProteXion 	<ul style="list-style-type: none"> ▶ 6.4 GBps peak bandwidth ▶ Memory Mirroring ▶ ChipKill ▶ Memory ProteXion ▶ Hot-swap memory ▶ 40% more aggregate I/O 	<ul style="list-style-type: none"> ▶ 21.3 GBps peak bandwidth ▶ Memory Mirroring ▶ ChipKill ▶ Memory ProteXion ▶ Hot-swap memory ▶ Hot-add memory ▶ Support for EM64T ▶ x86 64-bit Extensions
PCI bridge I/O controller	<ul style="list-style-type: none"> ▶ PCI-X 66 -133 MHz slots ▶ Remote I/O Support 	<ul style="list-style-type: none"> ▶ PCI-X 66 -133 MHz slots ▶ Remote I/O Support 	<ul style="list-style-type: none"> ▶ PCI-X 2.0 support ▶ All slots 266 MHz

These servers are designed for the demands of the application and database serving tiers offering leadership performance and reliability to power enterprise workloads such as ERP, CRM, and in-house developed applications, including Web services implementations.

The X3 Architecture offers investment protection for both software and hardware to meet the demands of this changing marketplace. With 32-bit compatibility on a 64-bit platform, the X3 Architecture allows you to migrate according to your business needs. The servers supports both 32-bit and 64-bit operating systems and applications.

And now with the addition of dual-core processors and an upgrade path for existing x366 and x460 clients, these servers are an excellent platform for business growth.

With the xSeries 460, IBM XpandOnDemand scalability combined with Xcel4v Dynamic Server Cache and the latest 64-bit Intel Xeon MP processor, the X3 servers offer unmatched levels of flexibility and expandability, up to 32 processors and 512 GB of RAM.

As the third-generation of Enterprise X-Architecture, the IBM @server X3 Architecture powers servers with the latest performance technologies to reduce latency and increase bandwidth for your commercial x86 applications. First to market with each of these high-performance server advancements, X3 Architecture servers feature Active PCI-X 2.0 up to 266 MHz, DDR2-based Active Memory, and highly reliable Serial Attached SCSI (SAS) hard drives.

Figure 2-1 shows the architectural differences between the x445 and the x460 servers and the improvements in bandwidth

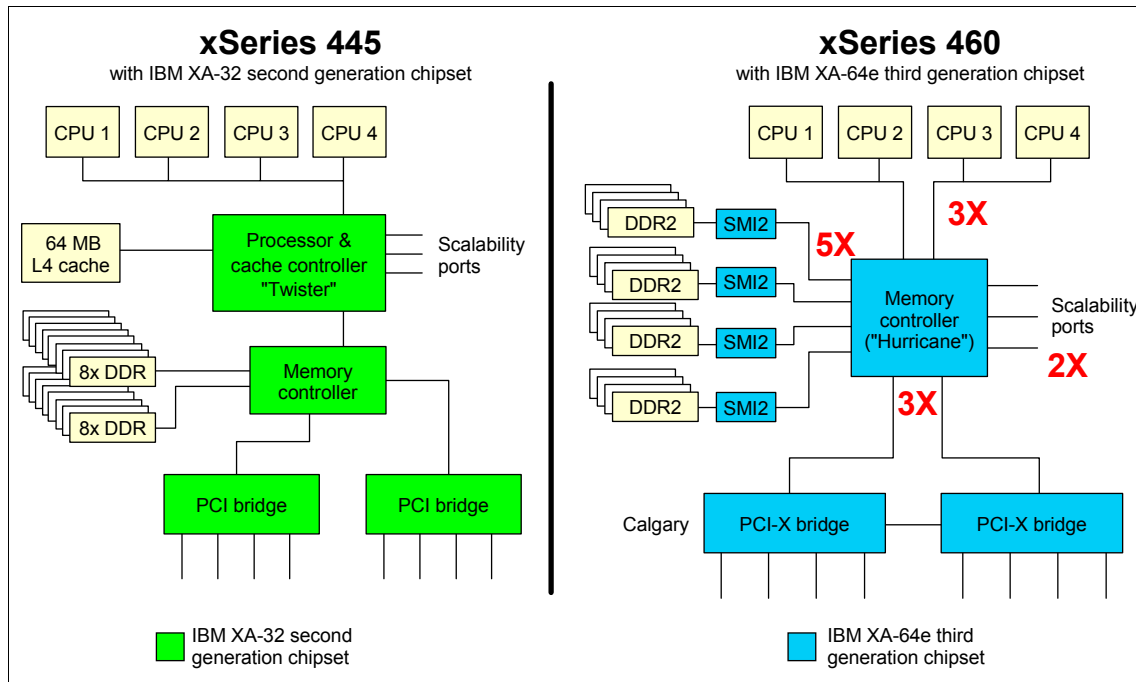


Figure 2-1 EXA2 versus EXA3 chipset comparison

The features and benefits of the IBM @server X3 Architecture include the following:

- ▶ Scalability
 - Highly modular: incremental CPU, I/O, and Memory.
 - Building block architecture.
 - Optimized for virtualization.
 - Pay-as-you-grow.
- ▶ Performance
 - 4-way x86 Performance.
 - 8-way to 32-way scalable performance with the x460 and MXE-460.
 - 32-bit and 64-bit, Intel EM64T.
 - Dual core models and upgrade options.
 - Active PCI-X 2.0.

- ▶ Availability
 - Active Memory (mirroring, hot-swap + Memory ProteXion).
 - Light path diagnostics.
 - PFA with IBM Director.
 - Redundant hardware components.
- ▶ Manageability
 - Remote deployment.
 - Multi-chassis partitioning.
 - Integrated hardware monitoring.
 - Integrated security.
- ▶ Competitive advantages
 - The x366 and x260 offer higher performance than competitive Potomac-based servers using lower priced Cranford processors.
 - The industry's only greater than 4-socket EM64T server.
 - The industry's highest performing x86-64 4-way, 8-way, 16-way, and 32-way.
 - Three years and over 80 number 1 Benchmarks! Enterprise X-Architecture has delivered more number 1 benchmarks than any other industry-standard server architecture:
 - x360: 15 number one benchmarks
 - x440: 35 number one benchmarks
 - x365: 11 number one benchmarks
 - x445: 21 number one benchmarks
 - x366: Six number one benchmarks

The following list will help you decide which of the three X3 Architecture servers to choose:

- ▶ If you plan to implement a system with eight or more processors, the only choice is the x460. Similarly, if your application can benefit from more than 64 GB of RAM, then choose a multi-node x460 configuration.
- ▶ If you plan to implement a 4-way system, and do not plan to scale-up further than 8-way, then the x366 offers better price/performance than the x460 in the 4-way space
- ▶ If you are planning 4-way or fewer processors, and a priority is rack density, then choose the x366 over the x260. Conversely, if you require more internal disk drivers or an internal tape drive, then choose the x260.

2.1.1 x366 versus x360

The x366 is capable of expanding to four CPUs and offers many new features over its predecessor, the x360. It is based on the IBM XA-64e chipset and utilizes new Xeon MP processors from Intel, which incorporate Intel EM64T 64-bit extensions, making this both an excellent 32-bit or 64-bit platform as well as a great environment to migrate from 32-bit to 64-bit.

The Active memory subsystem of the x366 is the most highly available memory subsystem in this level of server in the market today. In addition to incorporating both ECC and Chipkill memory, the x366 also incorporates memory mirroring, hot-swap memory, and Memory ProteXion, an IBM exclusive technology. Memory ProteXion provides a highly available memory configuration without having to resort to the significant expense associated with a mirrored memory configuration.

Memory on the x366 scales all the way from 2 GB to 64 GB, which is double that of the previous generation. All six I/O slots have been enhanced to support adapters up to 266 MHz. Neither the optional ServeRAID 8i or RSA II cards occupy PCI-X slots, since they connect via dedicated sockets on the system planar.

The x366 also incorporates the new serial attach SCSI (SAS) technology, with six 2.5" hot-swap disk bays available to accommodate disk expansion all the way up to 440 GB of internal disk storage. Both integrated dual Gigabit Ethernet and a easily serviced 3U mechanical package carry over from the previous generation of the 4-way rack product.

Table 2-2 compares the x366 with the x360.

Table 2-2 Major differences between x366 and x365:

Component	x365	x366
Chipset	XA-32 second generation chipset	XA-64e third generation chipset
SMP	1-way to 4-way SMP	1-way to 4-way SMP single-core or dual-core processors
Processors	Xeon MP up to 3.0/4 MB	Xeon MP single-core to 3.66 GHz/1 MB Xeon dual-core to 3.00 GHz/4 MB
Front side bus	Single-bus 32-bit Architecture, 400 MHz FSB	Dual-bus x86-64 Architecture 667 MHz FSB
Cache	L2: 512 MB integrated on the CPU L3: 1-4 MB integrated on the CPU L4: None	L2: Integrated on the CPU L3: None L4: None (XceL4v acts as snoop filter)

Component	x365	x366
Memory	32 GB max memory 2-way interleaving 16 max DIMM slots (8 standard) DDR SDRAM PC2100 6.4 GBps	64GB max memory 2-way interleaving 16 max DIMM slots (4 standard) DDR2 SDRAM PC2-3200 21.3 GBps memory bandwidth
Active Memory features	ECC Memory ProteXion Chipkill Memory mirroring Hot-swap	ECC Memory ProteXion Chipkill Memory mirroring Hot-swap Hot-add
SCSI	LSI Ultra320 SCSI Integrated RAID-1	Adaptec Serial Attached SCSI (SAS) Optional RAID with ServeRAID-8i
Storage	Six hot-swap 3.5" bays Max internal storage = 876 GB	Six hot-swap 2.5" bays Max internal storage = 440 GB
PCI-X slots	Active PCI-X: six internal 4 @ 133 MHz 1 @ 100 MHz 1 @ 33 MHz (RSA adapter) RXE-100 for 6-12 additional slots	Active PCI-X 2.0: six internal All @ 266 MHz No support for the RXE-100
Video	ATI Radeon	ATI Radeon 16 MB
Optical drive	24X CD-ROM	8X DVD-ROM
Power supplies	2 x 950W hot-swap standard (only one standard in some models)	One 1300W hot-swap One additional PSU optional
Ethernet	Broadcom 5704 dual port Gigabit	Broadcom 5704 dual port Gigabit
Management	Remote Supervisor Adapter II standard	BMC controller standard RSA II SlimLine optional
Warranty	3-year next business day 9x5 warranty	3-year next business day 9x5 warranty
Rack height	3U	3U

2.1.2 x460 versus x445

The x460 is second of the new X3 products announced and it too adds many new features over its predecessor, the x445. Like the x366, it is based on the IBM XA-64e chipset and also utilizes the new Xeon MP processors from Intel, which incorporate Intel's EM64T 64-bit extensions. The x460 also supports dual-core processors.

Each x460 supports up to four processors compared with the eight processors of the x445. The x460 supports up to eight nodes connected together to form a single 32-way complex.

The x460 uses the same Active memory subsystem as the x366, which are the most highly available memory subsystem in this level of server in the market today. It also incorporates both ECC and Chipkill memory, as well as memory mirroring, hot-add memory, and Memory ProteXion.

The memory on the x460 scales to 64 GB in a single chassis. This scales up to 512 GB in a 32-way complex (eight nodes). More and more application providers and clients alike are looking for larger amounts of memory so they can load the entire application into memory to gain the highest possible performance. The peak memory bandwidth of the x460 is over 21 GBps, up from 6.4 GBps in the x445.

All six PCI-X 2.0 slots have been enhanced to 266 MHz. The standard RSA II SlimLine card and the optional ServeRAID-8i card have dedicated connectors and do not consume one of these PCI-X slots. As you add additional chassis to a base x460, you can take advantage of the additional I/O slots that each chassis provides up to a maximum of 48 slots across the maximum of eight chassis. Because of this significant I/O support, the RXE-100 that was used in the previous generation of the scalable product is no longer be supported.

The x460 also incorporates the new serial attach SCSI (SAS) technology, with six 2.5" hot-swap disk bays available to accommodate disk capacity up to 440 GB per node. As with the I/O slots, the disk bays of each additional chassis can be used as you scale-up a base x460 configuration by adding more chassis.

Dual-port Gigabit Ethernet carries over from the predecessor product, but the older 4U mechanical package has been replaced with a denser and more modular 3U mechanical package that has excellent serviceability features.

Table 2-3 on page 49 shows the major differences between the x460 and the x445.

Table 2-3 Major differences between x460 and x445

Component	x445	x460
Chipset	XA-32 second generation chipset	XA-64e third generation chipset
SMP	Xeon MP: 2-way to 16-way Xeon DP: 2-way to 4-way	2-way to 32-way
Processors	<ul style="list-style-type: none"> ▶ Intel Xeon MP to 3.0/4M ▶ Maximum eight per chassis ▶ Intel Xeon DP 3.0GHz up to 4-way; maximum of four per chassis 	<ul style="list-style-type: none"> ▶ Xeon MP to 3.33 GHz 8 MB ▶ Dual-core processors to 3.00 GHz ▶ Maximum of four per chassis
Front side bus	400 MHz front side bus	667 MHz front side bus
Cache	L3: 1-4 MB L4: 64 MB XceL4 per 4-way (dedicated cache memory)	L3: 4-8 MB on single-core processors L4: 256 MB XceL4v per 4-way (uses main memory)
Memory	64 GB maximum (32 sockets) 16-way maximum DDR SDRAM PC2100 2-way Interleaving 6.4 GBps memory bandwidth	64 GB max per chassis (16 sockets) 32-way maximum: 512 GB DDR2 SDRAM PC2-3200 2-way Interleaving 21.3 GBps memory bandwidth
Active Memory features	ECC Memory ProteXion Chipkill Memory mirroring Hot-swap (top CEC only) Hot-add (top CEC only)	ECC Memory ProteXion Chipkill Memory mirroring Hot-swap Hot-add
SCSI	LSI Ultra320 SCSI; supports RAID-1 Optional ServeRAID	Adaptec Serial Attached SCSI (SAS) Optional RAID with ServeRAID-8i
Storage	Two hot-swap 3.5" bays Maximum internal storage = 292 GB	Six hot-swap 2.5" bays Maximum internal storage = 440 GB
PCI-X slots	Active PCI-X: six internal 2 @ 133 MHz 2 @ 100 MHz 2 @ 66 MHz RXE-100 for 6-12 additional slots	Active PCI-X 2.0: six internal All @ 266 MHz No support for the RXE-100
Video	ATI Rage XL 8 MB	ATI Radeon 16 MB
Serial port	One, using the supplied cable and PCI-slot bracket assembly (consumes a PCI slot)	One serial port

Component	x445	x460
Removable media	24X-8X DVD-ROM Diskette drive	8X DVD-ROM No diskette drive
Power supplies	Two 1200W hot-swap power supplies	Two 1300W hot-swap power supplies
Ethernet	Broadcom 5704 dual port GbE	Broadcom 5704 dual port GbE
Service processor	RSA II for EXA standard	BMC integrated RSA II SlimLine standard
Warranty	3-year next business day 9x5 warranty	3-year next business day 9x5 warranty
Rack height	4U	3U

2.1.3 x260 versus x255

The x260 is the third of the new X3 products and offers many new features over its predecessor, the x255. The x260 is based on the IBM XA-64e chipset and like the x366, uses the Intel Xeon MP “Cranford” processors, which incorporate the EM64T 64-bit extensions. These extensions allow the addressing of much higher amount of physical memory and the execution of 64-bit instructions and 32-bit instructions simultaneously, making this both an excellent 32 or 64-bit platform as well as a great environment to migrate from 32-bit to 64-bit.

The x260 uses the same Active Memory subsystem as the x366, which is the most highly available memory subsystem in this level of server in the market today. It also incorporates both ECC and Chipkill memory as well as memory mirroring and Memory ProteXion.

Memory on the x260 scales all the way from 2 GB to 64 GB, which is more than double that of the previous generation. All six I/O slots have been enhanced to 266 MHz and remain available, even if you install a ServeRAID 8i or RSA II card; you do not need to utilize card slots to add these capabilities.

The x260 also implements Serial Attach SCSI (SAS) technology, and the x260 can have up to 12 internal 3.5” disk drives attached. The ability to insert six drives is standard and a further six can be installed with the optional 6-bay backplane. The x260 also supports either one or two internal tape backup devices through its 5.25” drive bays.

Both integrated dual Gigabit Ethernet and an easily serviced 7U mechanical package that is available in either tower or rack-mount versions carries over from the previous generation of the 4-way tower product. A tower-to-rack conversion kit is also available to convert an existing x260 tower model to be installable in a rack.

The remote Supervisor Adapter II card has been re-designed as a SlimLine option and occupies a slot on the Super I/O card, as opposed to requiring one of the PCI-X slots on the I/O planar board.

Table 2-4 shows the major differences between the x260 and the x255.

Table 2-4 Major differences between x260 and x255

Component	x255	x260
Chipset	ServerWorks GC-HE chipset	XA-64e third generation chipset
SMP	1-way to 4-way	1-way to 4-way
Processors	Xeon MP: 2.0/1M, 2.2/2M, 2.7/2M, 3.0/4M	Xeon MP: 3.16/1M or 3.6/1M L2
Front side bus	400 MHz	667 MHz front side bus
Memory	24 GB max per chassis 12 DIMM sockets (all standard) DDR PC1600 SDRAM 2-way or 4-way interleaved	64 GB max per chassis 16 DIMM sockets max, 4 standard DDR2 SDRAM PC2-3200 2-way Interleaving 21.3 GBps memory bandwidth
Active Memory features	ECC Chipkill Memory mirroring Hot-spare memory	ECC Memory ProteXion Chipkill Memory mirroring Hot-swap Hot-add
Cache	L2: 512 KB integrated on the CPU L3: 1-2 MB integrated on the CPU L4: None	L2: 1 MB integrated on the CPU L3: None L4: None (XceL4v acts as a snoop filter only)
SCSI	Dual-channel Adaptec Ultra160 SCSI Optional RAID with ServeRAID	Adaptec Serial Attached SCSI (SAS) Optional RAID with ServeRAID-8i
Storage	Six hot-swap 3.5" bays standard Six additional bays with optional backplane Maximum internal storage = 1.7 TB	Six hot-swap 3.5" bays standard Six additional bays with optional backplane Maximum internal storage = 876 GB
PCI-X slots	Active PCI-X: Seven internal 6 @ 100 MHz 1 @ 33 MHz for RSA II card	Active PCI-X 2.0: six internal All @ 266 MHz No support for the RXE-100
Video	ATI Rage XL 8 MB	ATI Radeon 16 MB

Component	x255	x260
Ethernet	Single Port Broadcom 5703 Gigabit Ethernet	Dual Port Broadcom 5704 Gigabit Ethernet
Serial port	One serial port	One serial port
Removable media	48X CD-ROM Diskette drive Tape drive optional	40x CD-ROM Diskette drive Tape drive optional
Power supplies	Two hot-swap 370W standard, four maximum	Two hot-swap 775W standard, 4 maximum
Service processor	Optional RSA II	BMC integrated Optional RSA II SlimLine
Warranty	3-year next business day 9x5 warranty	3-year next business day 9x5 warranty
Mechanical	7U rack or tower available tower-to-rack conversion kit	7U rack or tower available tower-to-rack conversion kit

2.2 Focus market segments and target applications

The X3 Architecture servers from IBM are designed for the demands of the application and database serving tiers offering leadership performance and the proven reliability of the Intel Xeon MP processor architecture to power mission-critical stateful workloads, such as:

- ▶ Enterprise Resource Planning (ERP)

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

These applications today use a Web-based infrastructure with interfaces to suppliers, clients, and internal company employees. There are three general architectures used by enterprise solutions:

- Four-tier architecture (often referred to as an Internet architecture) with client systems, Web servers, application servers, and database servers
- Three-tier architecture, which includes client systems, Web/application servers and database servers

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

Key ERP software vendors are SAP (SAP Business Suite and My SAP All-in-one), Oracle (PeopleSoft and JD Edwards), Microsoft (Axapta®) and Baan

- ▶ Customer Relationship Management (CRM)

CRM is an IT-industry term for methodologies, software, and usually Internet capabilities that help an enterprise manage client relationships in an organized way. The application can use a four-tier, three-tier, or two-tier architecture similar to ERP applications.

Key CRM software vendors are Siebel, Oracle (PeopleSoft and JD Edwards), SAP (SAP Business Suite and My SAP All-in-one), Baan, and Onyx

- ▶ Supply Chain Management (SCM)

SCM is the oversight of materials, information, and finances as they move, through a process, from supplier to manufacturer to wholesaler to retailer to consumer. SCM involves coordinating and integrating these flows both within and among companies. The application also can use a four-tier, three-tier, or two-tier architecture.

Key SCM software vendors are I2, SAP (SAP Business Suite and My SAP All-in-one), Oracle (JD Edwards and PeopleSoft) and International Business System (IBS)

- ▶ Business Intelligence (BI)

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.

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

- ▶ Database

The X3 Architecture servers are ideal as database servers or application servers, with their fast 4-way and above processors and their large and very fast memory subsystems. The x460 in particular provides an extremely scalable platform with room to scale to additional nodes. These configurations use an external storage enclosure or SAN, depending on the size of the database, which is driven by the number of users.

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

Key database software vendors are IBM (DB2), Microsoft (SQL server), and Oracle

▶ Server Consolidation

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

Server consolidation is discussed in detail in 2.3, “Server consolidation” on page 56

Key server consolidation software vendors are VMware (ESX Server and GSX Server) and Microsoft (Virtual Server)

▶ eCommerce

eCommerce 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 clients. In order to achieve these goals, e-business requires a highly scalable, reliable, and secure server platform.

Key software vendors are IBM (WebSphere®) and BEA.

2.2.1 x260 target applications

The x260 is targeted at the following applications:

- ▶ Database: SQL Server, DB2, and Oracle
- ▶ ERP/CRM/SCM: SAP, Siebel, and i2
- ▶ E-mail and collaboration: Exchange and Notes
- ▶ Web services: WebSphere



The x260 is the departmental workhorse of the new generation of X3 servers. With its outstanding 64-bit performance and high availability features, the x260 server is well suited to a wide variety of different computing environments.

The x260 should be targeted to departmental and server consolidation applications as well as ERP/CRM/SCM, e-mail, and Collaboration and Web Services. It is also an excellent server for infrastructure applications, such as

departmental databases, where the application is not one that needs the additional scalability for which the 4-way scalable offering is better suited.

2.2.2 x366 target applications

The x366 is targeted at the following applications:

- ▶ Database: SQL Server, DB2, and Oracle
- ▶ ERP/CRM/SCM: SAP, Siebel, and i2
- ▶ E-mail and collaboration: Exchange and Notes
- ▶ Web services: WebSphere
- ▶ Server consolidation: VMware ESX Server



With its outstanding 64-bit performance and high availability features, the x366 server is targeted at the application tier in a multi-tier computing environment. The x366 should be targeted to applications such as ERP/CRM/SCM, e-mail and Collaboration, Web Services, and Server Consolidation.

With the elimination of any premium for buying a 4-way over two a 2-way, the x366 should be thought of as an upgradeable 2-way with more memory address slots (16 versus 8). The extra DIMM capacity of the x366 is perfect for 64-bit applications, and with twice the slots, clients can buy lower density DIMM sticks, saving even more money. The x366 is the better mid-tier commercial application server and offers application performance headroom with no price penalty.

2.2.3 x460 target applications

The x460 is targeted at the following applications:

- ▶ Database: SQL Server, DB2, and Oracle
- ▶ ERP/CRM/SCM: SAP, Siebel, and i2
- ▶ Server consolidation: VMware ESX Server



With its outstanding 64-bit performance, high availability features, and scalability, the x460 server is targeted primarily at the database tier or the high-end of the application tier in a multi-tier computing environment.

The x460 is targeted to very demanding applications, such as ERP/CRM/SCM and Server Consolidation. It is also an excellent server for database applications, providing the additional scalability of up to 32 processors and 512 GB of RAM that the x366 offering cannot.

2.3 Server consolidation

Server consolidation is the first step towards designing a more rational, efficient, and flexible IT environment that delivers consistently high levels of services to the entire organization. It is a lot more than simply replacing smaller boxes with fewer bigger boxes. It encompasses not just hardware, but software, services, and, most importantly, the systems management procedures that tie it all together.

Server consolidation is the discipline of simplifying end-to-end IT infrastructures, including servers, databases, applications, networks, and systems management processes, with the goal of reducing complexities and costs and establishing a stable foundation for growth and new solution deployment.

Over the last decade, Intel-based servers were purchased as a cost-effective alternative to mainframes and proprietary UNIX® servers. As demand grew, both servers and IT administrators were added to distributed departments. The frequent addition of a small number of inexpensive servers was not viewed as a major budget item when companies were flourishing. However, as server numbers grew, it triggered an enormous amount of hidden costs in software licensing, maintenance, management, and data center floor space, power, and cooling.

An unpredictable growth of the number of Intel base servers and applications in organization or departments happens due to number of reasons:

- ▶ It is a practice to run one application on one server.
- ▶ Application incompatibility.
- ▶ Reliable workload management techniques for splitting physical resources of one server by applications is missing.
- ▶ Lack of applications scalability.

2.3.1 Why consolidate servers?

The main objectives of server consolidation are:

- ▶ Need to make your company more responsive, more flexible, and more resilient in the face of changing market conditions and client demands

- ▶ Reduced IT budget for hardware and IT staff

Consolidating enterprise resources is a critical task for every organization, since by consolidating resources, many systems can be utilized more efficiently, maximizing resource utilization, increasing the efficiency of system management, security, and reliability, and reducing server count and complexity.

The goals of server consolidation often include:

- ▶ Reduce computing costs and complexities
- ▶ Improve productivity through better management of information, data, and resources
- ▶ Improve service levels, and establish a flexible, stable foundation for growth and new solution deployment
- ▶ Reduce hardware, software costs
- ▶ Better use existing facility
- ▶ Reduce the maintenance, upgrade, and licensing costs
- ▶ Optimize server performance
- ▶ Faster deployment of applications in a more secure environment
- ▶ Standardize procedures and operations
- ▶ Reduce total cost of ownership

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

2.3.2 Types of server consolidation

There are four general types of server consolidation projects that clients are implementing, offering a wide range of business value through varying degrees of solution complexity and investment. The most effective approach for a specific enterprise could fall into any one of these four categories, or it may require work across all four areas:

- ▶ Centralization
- ▶ Physical consolidation
- ▶ Data integration
- ▶ Application integration

These areas are shown in Figure 2-2.

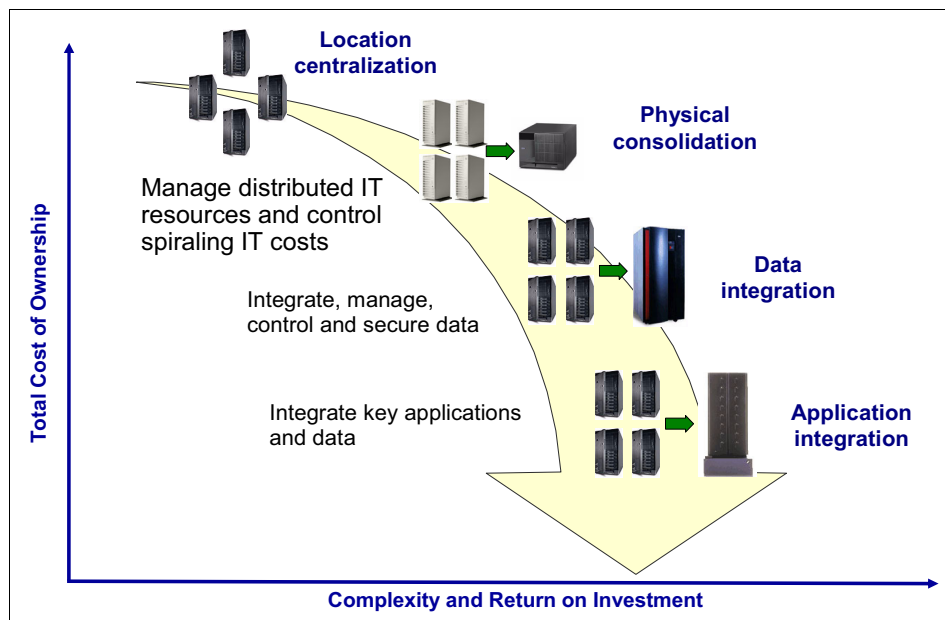


Figure 2-2 Types of server consolidation

The four types are implemented as follows:

- ▶ Location centralization

Centralization is commonly called datacenter consolidation. By simply relocating existing servers to fewer numbers of IT sites, economies of scale of operation can provide simplified management and cost improvement. Centralization is typically the initial step most companies take towards doing something broader. And this is almost always where organizations begin to

support mergers and acquisitions, by co-locating data centers before merging operations and business processes.

► Physical consolidation

Physical consolidation is what most people think of first when they consider server consolidation. It is the process of replacing or reducing the actual number of servers by replacing many small servers with fewer, more powerful servers or clustered systems.

This can take place within the same architecture or across architectural boundaries. For example, replacing several two-way Windows servers with one 16-way VMware ESX Server-based system with several virtual machines, or with a few BladeCenter chassis.

This approach is typically appropriate for implementations of key packaged applications such as SAP, PeopleSoft, and Siebel, where minimal integration with other applications and data is required. Additional benefits can be gained through data integration and application integration. While these are often more complex projects that require extensive analysis, planning, and implementation, they can provide significant return-on-investment.

► Data integration

Data integration involves physically combining data from different sources across the enterprise into a single repository or format. The result is that the merged data can reside on fewer servers and more centralized and consistent storage devices, greatly lowering the total costs associated with managing the data. The data can, on file servers across disparate operating systems, be consolidated to a single system. Also, multiple types of databases, such as DB2, SQL Server, Oracle, and Sybase can be converged to a single database architectures.

► Application Integration

This approach involves consolidating and combining multiple applications and data to fewer server architectures for not only reduced complexity, but also for business process integration and automation. It allows for co-locating of mixed workloads within an unified infrastructure so that applications can communicate with each other and work together seamlessly.

For example, some business processes can be achieved through integration of order entry, inventory, fulfillment, and accounting systems (Web serving, ERP, business intelligence, and other applications) that represent the backbone of the core business.

2.3.3 Benefits of server consolidation

There are numerous benefits to server consolidation, depending on the type of consolidation you plan to implement:

- ▶ Centralization:
 - Reduction in administration costs
 - Increased reliability and availability
 - Lower operation costs
 - Improved security and management
- ▶ Physical server consolidation:
 - Reduced hardware, operating system, and database types
 - Reduced number of servers per application (mail and file server)
 - Reduced hardware and software costs
 - Improved processor utilization
 - Reduced facilities costs (space, power, and A/C)
 - Lower operations costs
 - Improved manageability
 - Improved reliability
- ▶ Data integration:
 - Improved resource utilization
 - Reduction in administration costs
 - Reduced duplicated data
 - Enhanced data access and integrity
 - Reduced storage management costs
 - Improved backup/recovery capabilities
 - Improved reliability
- ▶ Application integration
 - Reduction in administration costs
 - Increased reliability and availability
 - Reduced facilities costs (space, power, and A/C)
 - Lower operation costs
 - Scalability

The main benefits of server consolidation, generally, are:

- ▶ Single point of control.

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

A single point of control allows enterprises to:

 - Reduce or eliminate department operational costs.
 - Reduce some software licenses.

- Reduce the number of the systems and disk storage costs.
 - Reduce maintenance charges.
 - Avoid multiple copies of the same application on distributed systems.
 - Reduce owner operational costs.
 - Offer better availability of service.
 - Improve systems management.
 - Have better version control management.
 - Have better software distribution.
 - Reduce risk and increase security.
- ▶ Giving users better services.
- With a consolidated infrastructure, it is easier to offer 24x7 support to end users. The response time is much better than with an overly distributed environment, and the data is more easily accessible while being highly protected. The control procedures are simpler, while security becomes even higher. And information sharing is improved, giving end users increased data consistency. The availability of service is improved mainly due to a reduction in the time needed to communicate between clients and servers in a single location.
- ▶ Regaining flexibility.
- The standardization of procedures, releases, and servers also makes it easier to install new application software. Computing resource consolidation enables a trouble-free upgrade of the information system and less costly adaptation to organization or environment changes. Enterprises can react more quickly to market changes, since storage is readily available and can easily be reallocated.
- ▶ Avoid floor space constraints.
- While a small server may be easily fit into a closet, as compute demands increase, enterprises find that suitable floor space is hard to find for proliferating small servers. The solution is a central site outfitted with appropriate power, cooling, access to communications links, and so on, and populated with more powerful systems, each giving more performance in the same footprint.
- ▶ Reduction of the Total Cost of Ownership (TCO).
- There are several costs associated with server consolidation, including:
- Hardware costs: New servers and infrastructure, and upgrades.
 - Software costs: Fewer software licenses are required with fewer servers.
 - Disruption costs: Migration and change management.

► Manageability and availability

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

- Enterprise management: Integrated operations allows for consistent management of all facilities and IT services.
- Consistent performance: Providing consistent response time at peak load periods is very important.
- Dependability: Commonly cited problems of distributed environments include frequency of outages and excessive requirements for manual intervention by the IT staff. In addition, it provides the following benefits:
 - It is easier to enforce consistent user policies in a consolidated environment.
 - Fewer servers lead to a simpler network structure that is easier to manage.
 - Reorganization following mergers or acquisitions is easier in a well-controlled environment.
 - Consolidation encourages standardization of tools, processes, and technologies to provide a stable and consistent application platform.

► Data access and protection.

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

- Network technology: The growth of networking and network speeds is enabling the centralization of IT networks today and will continue and expand into the future.
- Fragmentation and duplication of data: This is a core issue in most organizations with large numbers of distributed servers.
- Physical security: Consolidation of servers in a central data center can restrict unwanted access and ensure a more secure environment.
- Integrity, local backup, and recovery: Enterprises are concerned about the dangers of business disruption, client lawsuits, and regulatory action in the event of severe data loss, and they need to implement effective disaster recovery procedures.

► Use your existing investment in IT.

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

- Expand existing servers: Add new capabilities to the existing installation rather than deploy new dedicated servers.

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

► Scalability and workload growth.

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

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

► Service level.

Most companies spend much of their IT budget for services. They need services for hardware, software, and infrastructure maintenance. Server consolidation can help you to reduce the increasing service costs in the following ways:

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

▶ Business continuity.

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

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

In most cases, developing a sound and visionary IT realignment and consolidation strategy will not only provide IT benefits, but can also serve to improve the economic and operational conditions of an enterprise for strategic business achievements.

Table 2-5 shows how benefits in IT can translate into real business benefits.

Table 2-5 Business benefits

IT benefits to consolidation	Business benefits to consolidation
<ul style="list-style-type: none"> ▶ Reduce complexity ▶ Lower costs and dramatically improve TCO ▶ Simplify systems management ▶ Enhance resources utilization ▶ Drive open adoption of standards end to end ▶ Improve uptime/availability and increase recoverability ▶ Improve performance and optimize scalability ▶ Accelerate e-business integration 	<ul style="list-style-type: none"> ▶ Attract and retain clients more efficiently ▶ Maximize revenue per client ▶ Bring new services online quickly ▶ Respond to new business challenges quickly by establishing a foundation for future growth ▶ Consolidate operations and overall systems control ▶ Reduce or eliminate redundancy in infrastructure and personnel ▶ Better service client, employees, suppliers, and partners

2.4 XpandOnDemand

XpandOnDemand is the term given to the ability of the x460 to scale-up as required by adding nodes to an x460 complex. This ability ensures that clients only pay for the computing resources they need at the present time, without sacrificing the investment they have made in hardware and software should they choose to upgrade.

The investment protection features of the x460 include the following:

- ▶ **Processors:** Clients can upgrade from 2-way to 4-way, 8-way, 16-way, and ultimately 32-way processing should they require it. Models with single-core processors can be upgraded to dual-core processors.
- ▶ **Memory:** The standard 2 GB can be expanded to 512 GB
- ▶ **PCI-X slots:** The x460 has six PCI-X 2.0 slots standard, but an x460 complex can be expanded to 48 slots if required.
- ▶ **Drives and USB devices.** The number of internal disk drives and other internal resources is also increased as a result of such expansions.

These “pay-as-you-grow” design options are achieved by connecting x460 servers (or x460s with MXE-460 expansion units) together to form multi-chassis (or multi-node) complexes. Each node contains processors, memory, and PCI-X slots and other components in the x460. By joining the nodes together, the single operating system running on the entire complex has full access to all resources in all attached compute nodes.

In addition, once the nodes are connected together to form a larger complex, should the business need arise, you can divide the complex back into partitions. These partitions are formed on node boundaries. For example, an 8-node x460 complex could be divided into a 4-node partition and two 2-node partitions, each running its own operating system. This flexibility ensures that the computing resources match the need of the business.

All the X3 Architecture servers support 64-bit operating systems using the Intel Xeon MP processors with EM64T extensions. This allows you to grow more easily as the needs of your business change over time, transitioning to 64-bit applications or adding incremental performance capacity when you need it without the penalty of paying for costly up-front infrastructure.

The X3 Architecture servers are also ready for the next wave of processor technology, with the support of the Intel dual-core CPUs.

The main advantages of XpandOnDemand capability are:

- ▶ **Investment protection:** You pay only for performance that is necessary today. You do not need to pay more than you need now. Supports 32-bit and 64-bit applications on the same platform. You can migrate to 64-bit as needed or when required 64-bit versions of commercial applications will be available.
- ▶ **“Pay-as-you-grow”:** Server performance is growing with your company. You do not need to buy a new, more powerful server; just expand the server you already have.

- ▶ Configuration and performance flexibility: Server configuration and performance can be simply modified to smooth peak loading or to periodically perform some more processor intensive tasks, such as weekly accounting and billing.
- ▶ Near linear performance increase: By adding additional nodes to the xSeries 460 based on the Enterprise X-Architecture, you will get not only additional processors and power, you will also get more memory capabilities, more PCI-X slots, more internal storage, more front side buses, more memory controllers with additional memory buses, and more chipsets on the motherboard, which, in some cases, gives you practically linear performance increase.

2.5 Positioning the x460 and the MXE-460

The MXE-460 is almost identical to the x460. The purpose of the MXE-460 is to act as an expansion node for multi-node configurations. The MXE-460 is less expensive than the x460 and can be used as a secondary node in partitions.

When building a multi-node configuration, you can use a combination of x460 and MXE-460 systems as the nodes. It is the number of x460s you have that determines how you can partition the complex. For example, if you require a 8-node complex, you can do the following:

- ▶ If you configure one x460 and seven MXE-460, then you can only create one partition (a 32-way partition).
- ▶ If you configure eight x460s and no MXE-460s, then you will have maximum flexibility as to what partitions you can make, from a single 32-way partition, to eight 4-way partitions, and combinations in between.

Tip: The key is that for every partition you wish to create, the primary node in that partition must be an x460. Having an MXE-460 as a primary partition is not supported.

The technical differences between the x460 and the MXE-460 are as follows:

- ▶ The x460 comes standard with processors and memory. The MXE-460 ships standard with neither installed. Clients will need to install matching CPUs and the appropriate amount of memory (we recommend you also match the amount of memory installed).
- ▶ The x460 has a DVD-ROM as standard. The MXE-460 has no optical drive standard.

- ▶ Some components of the country kit are different. For example, the MXE-460 country kit does not include a ServerGuide™ CD.
- ▶ The MXE-460 is not supported as a primary node in any partition.

Once you have cabled and configured a multi-node complex, you need to partition the complex. Most clients will not create more than one single partition in the complex. However, there are some advantages to creating multiple partitions:

- ▶ You can run different operating systems or versions on different partitions without the need for products such as VMware ESX Server.
- ▶ You can easily reconfigure the partitions if you need to perform some periodic processor-intensive or memory-intensive tasks, such as a weekly accounting or a Business Intelligence analysis task.

Figure 2-3 shows the configuration alternatives for 2-node and 4-node complexes, and the partitioning options available to you as a result. Note that the 8-node configuration is not shown in the diagram, but the same rules apply.

If you want maximum partitioning flexibility, all your nodes should be x460 systems.



Figure 2-3 Supported partitioning for 2-node and 4-node complexes

The key rules about partitioning are as follows:

- ▶ Partitioning is always at node boundaries. You cannot form, for example, a partition using two processors in a node and two processors in another node.
- ▶ The primary node for every partition must be an x460. The use of MXE-460s limits the partitioning you can do.
- ▶ All nodes in a multi-node partition must have four processors installed.

2.6 Application scalability

If you plan to increase performance of your system, there several issues that you should consider:

- ▶ Application scalability
- ▶ Operating system scalability
- ▶ Server scalability
- ▶ Storage scalability

Adding processors can improve server performance under certain circumstances because software instruction execution can be shared among the additional processors. However, both the operating system and, more importantly, the applications must be designed to take advantage of the extra processors. Merely adding processors does not guarantee a performance benefit.

For example, not all applications can use the full power of four processors in the one server. File and print servers often only take advantage of one or two processors and popular mail systems typically only scale well to four processors. Table 2-6 shows the suitability of multi-processor systems to application types.

Table 2-6 Processor scalability by application type

Processors	File & print	Web server	E-mail collaboration	Business logic	Database	Server consolidation
1 way	Suitable	Suitable	Suitable	Suitable		
2 way	Suitable	Suitable	Suitable	Suitable	Suitable	Suitable
4 way			Suitable	Suitable	Suitable	Suitable
8 way					Suitable	Suitable
16 way					Suitable	Suitable
32 way					Suitable	Suitable

Processors are only one part of the scalability story. It is typically very important that the memory, disk, and networking subsystems are examined for the effect on scalability. It is typical that the performance gains from adding processors are only realized when memory is added in parallel. For disk intensive applications, such as OLTP-type applications, it is essential to have a large disk array to stream data to the CPU and memory subsystems so that any disk-related delays are kept to a minimum.

It is important to plan your system in advance according to your business requirements so that you will not have to replace your server, operating system, or storage subsystem, because your server no longer meets your processing needs, for example, the operating system does not support more than four processors, or your server is not able to hold more than six PCI-X adapters.

Table 2-7 shows how application types scale and what is required to achieve peak performance. This table lists the server configurations used to produce some recent benchmark results. As you can see, the amount of memory and disks varies widely depending on the application.

Table 2-7 Benchmarks configurations of processors, memory, and disk

Benchmark	Processors	Memory (GB)	Disk drives
TPC-C (Database)	8	128	500
	4	64	351
TPC-H (Decision support)	8	32	138
	4	16	160
SAP SD (ERP)	16	128	14
	8	16	15
Oracle OASB 11.5.6	8	32	42
PeopleSoft GL 8.4	8	8	92
Exchange MMB3	4	8	303
SPECjbb2000 (Java)	8	32	1
SPEC CPU2000 (CPU-intensive)	8	4	2

The different server configurations reflect the different workloads of these benchmarks. The workload that the benchmark generates causes the server to bottleneck in a particular subsystem. To alleviate the problem, resources are added to that subsystem until the bottleneck is moved to another subsystem.

The last row in Table 2-7 on page 69 also highlights the component-focused nature of the SPEC benchmarks and the CPU-intensive applications they serve. This 8-way server required only 4 GB of memory and two disks. Clearly, the workload isolates the CPU with very little dependency on other subsystems. This means that the benchmark may be very good for comparing raw CPU performance, but it provides limited information regarding the performance of the entire system. The CPUs in a system may be very fast, but performance remains poor if the memory or I/O subsystems cannot supply data to them quickly enough.

2.6.1 Enterprise application scaling

Enterprise applications enable you to run your business more effectively and are often referred to as back-office applications. They bring together four major application groups to create integrated end-to-end solutions:

- ▶ Business Intelligence
- ▶ Customer Relationship Management
- ▶ Enterprise Resource Planning
- ▶ Supply Chain Management

Enterprise applications work with your most critical business data so it is important that these applications are highly available and secure. There are three general architectures used by these applications, as shown in Figure 2-4 on page 71:

- ▶ A three-tier architecture (often referred to as an Internet architecture) includes client systems, Web servers, application servers, and database servers.
- ▶ A two-tier architecture includes client systems, Web servers, and database/application servers.
- ▶ A three-in-one tier architecture includes client systems and database servers.

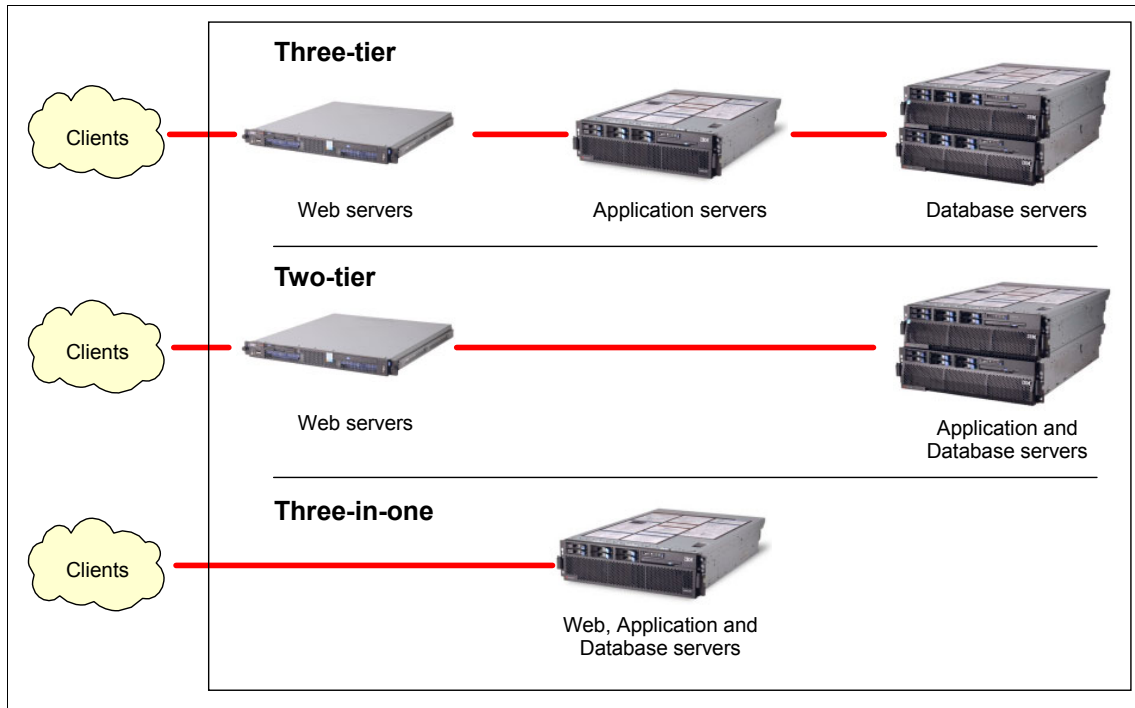


Figure 2-4 Enterprise solution architectures

While three-tier architecture has far greater complexity, it also allows for greater scalability. The architecture selected for a solution will depend on your business needs, the type of application deployed, and the number of planned users. In most cases, if you need to scale your applications, use a two-tier or three-tier architecture. Smaller clients may prefer to implement a “three-in-one” implementation, simply because it is easier to manage and the number of users supported can be handled by the “three-in-one” solution.

See also 2.2, “Focus market segments and target applications” on page 52.

2.6.2 SAP product offering

SAP is the largest supplier of integrated business solutions in the world. SAP is traditionally seen as an ERP vendor. Today, through its suite of modular mySAP.com solution offerings, SAP can provide integrated end-to-end solution applications that span ERP, CRM, SCM, and BI. In addition, SAP offers industry-specific solutions based on mySAP.com.

SAP offerings are:

- ▶ SAP Business One. For clients with 10-250 employees. For wholesale distribution and retail industries.
- ▶ mySAP All-in-One. For clients with 10-100 employees. For these industries:
 - Fabrication and Assembly
 - Electronics and Electrical
 - Automotive
 - Food and Beverage
 - Wholesale Distribution
- ▶ mySAP Business Suite. For high-end SMB and ISU clients in all industries.

Typical SAP solutions are usually implemented as two-tier or three-tier configurations, as shown in Figure 2-4 on page 71.

- ▶ Two-tier
 - Simple network
 - Simple to upgrade
 - Simple and inexpensive to manage
 - Easy transition to Three-Tier configuration
- ▶ Three-tier architecture
 - Nearly unlimited scalability
 - Improved availability and multiple networks

You can use the xSeries product selection criteria shown in Table 2-8 for SAP Solutions.

Table 2-8 xSeries product positioning for SAP solutions

xSeries	Suggested fit within SAP solution
x460	<ul style="list-style-type: none">▶ DB Server in large 3-tier installations with external storage▶ Large 2-tier installations with external storage▶ Powerful application server in front of a pSeries® or zSeries® database in a 3-tier deployment
x366	<ul style="list-style-type: none">▶ Medium to large 3-tier installation as back-end database or application server▶ Medium 2-tier installation with external storage
x260	<ul style="list-style-type: none">▶ Small to medium 2-tier installation▶ Database server within a medium 3-tier DB server▶ Test and development platform

xSeries	Suggested fit within SAP solution
x346/x236	<ul style="list-style-type: none"> ▶ Application server within a medium 3-tier implementation ▶ Database server within small to medium 3-tier implementation ▶ Citrix or Terminal Server ▶ CRM Comm Station or J2EE™ application Server ▶ Test and development platform ▶ Small 2-tier installation (75 users)
x336	<ul style="list-style-type: none"> ▶ Application server within a medium 3-tier SAP implementation ▶ Web server ▶ Citrix or Terminal Server ▶ ITS or J2EE application server ▶ CRM Comm Stations or CRM repository server

For information about SAP solutions including sizing, refer to the following Web Pages:

- ▶ IBM and SAP Home Page
<http://www.ibm.com/solutions/sap/us/en/>
- ▶ IBM hardware sizing portal
<http://www.ibm.com/servers/solutions/finder/portal/hw-sizing.html>
- ▶ IBM Sizing and Planning Questionnaire for SAP Solutions
<http://www.ibm.com/support/techdocs/atmsastr.nsf/PubAllNum/PRS261>
- ▶ SAP Benchmark results
<http://www.sap.com/solutions/benchmark>

For IBM employees, the following links are relevant resources:

- ▶ IBM internal SAP portal on the Sales Support Advisor
<http://w3.ncs.ibm.com/solutions/SAP/>
- ▶ IBM SAP International Competence Center
<http://w3.ncs.ibm.com/solution.nsf/SOL/BTIY-5PBJ73>

2.6.3 Oracle/J.D. Edwards product offering

The J.D. Edwards product on the Intel platform is EnterpriseOne, which is a suite of integrated industry-specific business application modules:

- ▶ Asset life cycle management

- ▶ Customer Relationship Management (CRM)
- ▶ Financial management
- ▶ Human Capital Management
- ▶ Manufacturing and supply chain management (SCM)
- ▶ Project management
- ▶ Supply management
- ▶ EnterpriseOne tools and technology

Typical EnterpriseOne solutions are based on the following architectures, as shown in Figure 2-4 on page 71:

- ▶ All-in-one, where the application server, application logic, batch workload, and database server all run on one server.
 - This is a perfect solution for the majority of J.D. Edwards clients, from small to mid-market.
 - Consolidated and simplified IT architecture.
 - The x460 is ideally suited to run all of J.D. Edwards' server functions on one server.
- ▶ Three-tier configuration, when DB and application servers are separate servers.

EnterpriseOne supports the following

- ▶ Databases:
 - IBM DB2 UDB
 - Microsoft SQL Server
 - Oracle
- ▶ Operating systems:
 - Microsoft Windows
 - Red Hat Linux

A typical J.D. Edwards solution includes:

- ▶ Citrix for Windows client over WAN
- ▶ DB2 UDB
- ▶ WebSphere Application Server
- ▶ WebSphere Portal Server

You can use the xSeries product selection criteria shown in Table 2-9 on page 75 for J.D. Edwards EnterpriseOne solutions.

Table 2-9 xSeries Product Positioning for J.D. Edwards Solutions

xSeries	Suggested Fit within J.D. Edwards Solution
x460	<ul style="list-style-type: none"> ▶ Powerful all-in-one solution ▶ Excellent database server ▶ Powerful application server in front of a pSeries or zSeries database in a 3-tier deployment
x366	<ul style="list-style-type: none"> ▶ Database node or medium-sized implementation ▶ Application server in front of x460 back-end database ▶ Citrix or terminal server
x260	<ul style="list-style-type: none"> ▶ All-in-one solutions for small to medium implementations ▶ Excellent choice for clients wanting to maximize internal storage
x346/x236	<ul style="list-style-type: none"> ▶ Front-end Web server ▶ Application server for small to medium Implementations ▶ Citrix, terminal, deployment, or Crystal Reports server
x336	<ul style="list-style-type: none"> ▶ Front-end Web server ▶ Citrix or terminal server ▶ Crystal Reports server

IBM has built a strong foundation and alliance with J.D. Edwards based on a 25-year relationship serving mid-market clients. Approximately 85% of all J.D. Edwards clients are running on IBM @server hardware.

2.6.4 Siebel product offering

Siebel offers solutions in the CRM arena: an integrated suite of sales, marketing and client service applications for field sales and service, customer service, telesales, telemarketing, third-party resellers, and internet-based e-commerce. The core design of Siebel CRM solutions is modularity to meet the needs of all businesses.

Siebel offers a broad range of product sets for enterprise clients and for mid-market clients:

- ▶ Business Analytics is a suite of advanced Business Intelligence applications suitable for sales, services, marketing, finance, human resources, and supply chain requirements.
- ▶ Siebel Professional CRM Edition: The product line formerly known as Siebel eBusiness Applications provides modular CRM functionality, allowing

organizations to choose the specific CRM capabilities needed to run their business. This product offers an easy path to Siebel enterprise applications and can be deployed with Siebel CRM OnDemand.

- ▶ Siebel CRM OnDemand: This hosted CRM software offering is delivered via the Web and accessible from a Web browser at a fixed price per user, per month. Clients can deploy Siebel CRM OnDemand quickly, easily, and affordably, without any upfront IT investments. Siebel CRM OnDemand delivers complete sales force automation, marketing automation and customer service functionality, built-in client analytics, virtual call center technology, embedded CRM best practices, and world class hosting services and support.

Typical Siebel solution based on two following architectures, as shown in Figure 2-4 on page 71:

- ▶ All-in-one, where all the application servers, Web servers, and DB servers run on one server
- ▶ Three-tier configuration, where the DB, application server, and Web server are separate servers

You can use the xSeries product selection criteria shown in Table 2-10 for Siebel solutions.

Table 2-10 xSeries Product Positioning for Siebel Solutions

xSeries	Suggested Fit within Siebel Solution
x460	<ul style="list-style-type: none"> ▶ Excellent Siebel database server node ▶ Powerful application server in front of a large pSeries, zSeries, or even a 16-way x460 back-end database
x366/x260	<ul style="list-style-type: none"> ▶ Small/medium Siebel database server ▶ Excellent remote Siebel server or application server
x346/x236	<ul style="list-style-type: none"> ▶ Remote Siebel server ▶ Small application server ▶ Web server
x336	<ul style="list-style-type: none"> ▶ Excellent fit as Web server ▶ Gateway server

IBM and Siebel Systems have created an International Competency Center (ICC) in San Mateo, California, dedicated to joint solution development. We have extended our collaborative efforts by opening and jointly staffing a support center in Montpellier/La Gaude, France, where we validate Siebel architecture on IBM hardware and software and showcase advanced technologies. Clients, business

partners, consultants, and integrators alike are invited to observe these efforts firsthand. Solution centers also demonstrate and provide a wide array of sales and technical support options.

Some useful links are listed below:

- ▶ IBM-Siebel portal
<http://www.ibm.com/solutions/businesssolutions/siebel>
- ▶ IBM Sizing and Planning Questionnaire for Siebel
<http://www.ibm.com/support/techdocs/atmastr.nsf/PubAllNum/PRS593>
- ▶ For IBM employees, the internal IBM-Siebel portal is at:
<http://w3.ncs.ibm.com/solutions/Siebel>

2.6.5 Oracle/PeopleSoft product offering

PeopleSoft was originally a human resources solution. The current solution now includes components that help link clients, suppliers, partners, and employees to provide a competitive advantage.

PeopleSoft uses IBM @server servers to run their own business and has chosen DB2 as their strategic database platform. In addition, PeopleSoft has embedded WebSphere Application Server into PeopleSoft applications. IBM can offer the most comprehensive and integrated solution stack for PeopleSoft ranging from hardware, middleware, financing, and services.

Together, IBM and PeopleSoft share over 6,000 clients worldwide using PeopleSoft products. These clients are using a variety of solutions from classic financial management tools to advanced supply chain, procurement, customer relationship management, and fully Web-enabled offerings.

There are three product lines to address the varying needs of the market:

- ▶ PeopleSoft Enterprise is targeted at large enterprises in industries, such as financial services and health care.
- ▶ PeopleSoft EnterpriseOne is focused on product-related industries, such as construction, manufacturing and distribution, and asset intensive industries
- ▶ PeopleSoft World is the leading platform for iSeries clients, especially those who want a low risk solution, a single database, and Web enablement

Typical PeopleSoft solutions are based on an Internet application server architecture with a “no code on the client” philosophy that allows rapid deployment with PeopleSoft systems using standard protocols.

PeopleSoft solutions implement one of three architectures, as shown in Figure 2-4 on page 71:

- ▶ All-in-one, where all the application servers, Web servers, and database servers run on one server
- ▶ Three-tier configuration, where the database, application server, and Web server are separate servers
- ▶ Two-tier configuration, where the DB server runs on one server and the application server and Web server are on a separate server

People Soft supports:

- ▶ Databases
 - DB2 zOS
 - DB2 UDB
 - Informix®
 - SQL Server
 - Oracle
 - Sybase
- ▶ Operating systems for the application servers
 - Linux
 - Windows
 - AIX®

You can use the xSeries product selection criteria shown in Table 2-11 for PeopleSoft solutions.

Table 2-11 xSeries product positioning for PeopleSoft solutions

xSeries	Suggested fit for PeopleSoft solutions
x460	<ul style="list-style-type: none"> ▶ Excellent scalable database server node ▶ Powerful server for an all-in-one PeopleSoft solution ▶ Database server for 500+ users for PeopleSoft SCM solutions or 1000 users for PeopleSoft CRM solutions ▶ Powerful application server in front of a large-scale pSeries or zSeries back-end database

xSeries	Suggested fit for PeopleSoft solutions
x366/x260	<ul style="list-style-type: none"> ▶ x366: Small/medium PeopleSoft database server ▶ Good as an application server in-front of an x460, pSeries, or zSeries back-end database ▶ x260: Excels as a test and deployment platform when a large server is required by the client ▶ All-in-one server when internal storage is required
x346	<ul style="list-style-type: none"> ▶ Powerful Web application server ▶ Small application server ▶ Reporting server for PeopleSoft applications
BladeCenter	<ul style="list-style-type: none"> ▶ Can host multiple application servers in an extremely rack dense form factor ▶ Opportunity to have Web server, application server, and other point solutions servers in a single chassis
x336/x236	<ul style="list-style-type: none"> ▶ x236: Affordable test and deployment server ▶ x336: Excellent Web server

2.6.6 Microsoft SQL product offering

Table 2-12 on page 80 shows the Microsoft SQL Server product offering, which ranges from the Enterprise Edition down to the Personal edition.

- ▶ The Enterprise Edition is the most scalable from both a CPU and memory perspective. In addition, it offers some unique features, such as failover clustering, distributed partition views, log shipping, SAN support, ROLAP, and VL dimensions, and other advanced OLAP and analysis functions.
- ▶ The Standard edition is one that is often used and targeted towards small to medium sized businesses. It includes all of the core SQL server features except those mentioned above and it does include analysis services.
- ▶ Developer is the same as the Enterprise, but it is only for use for development and test purposes only. In addition, you will have to download SQL Server CE and get distribution rights for it as well.
- ▶ The Personal is a subset of the Standard, which again is ideal for mobile computers and is not available for purchase, but it ships with the Enterprise and Standard editions.

Table 2-12 Microsoft SQL Server product offering

Edition	Max CPU	Max RAM	Unique features
Enterprise	32	64 GB	<ul style="list-style-type: none"> ▶ Failover clustering ▶ Distributed partitioned views ▶ Log shipping ▶ Parallel operation for SMP systems ▶ SAN support ▶ ROLAP and VL dimensions ▶ Distributed partitioned cubes ▶ Other advanced OLAP/analysis
Standard	4	2 GB	<ul style="list-style-type: none"> ▶ All Core SQL Servers features except those listed above ▶ Includes Analysis Services
Developer	32	64 GB	<ul style="list-style-type: none"> ▶ For development and test use ▶ Includes SQL Server CE download/distribution rights
Personal	Subset of Standard, ideal for mobile computers; not available for purchase, ships with Enterprise Edition and Standard Edition		

2.6.7 DB2 Universal Database

DB2® UDB for Windows is available in several editions that can scale from hand-held devices running Windows CE to server editions running on Windows Server 2003, Datacenter Edition.

DB2 UDB V8.1 is certified for Microsoft Windows Server 2003 for the Standard, Enterprise, and Datacenter editions of the operating system, both the 32-bit and 64-bit versions.

There are several server editions of DB2 UDB V8.1 that can be deployed on the Windows Server 2003 operating system. These include:

- ▶ Express Edition
- ▶ Workgroup Server Edition
- ▶ Workgroup Server Unlimited Edition
- ▶ Enterprise Server Edition

Table 2-13 on page 81 shows the various DB2 UDB server editions.

Table 2-13 DB2 UDB server editions

Edition	Max processor	Description
Express	2-way (32-bit)	A full-function database server designed to be embedded as a transparent part of a business solution with a named user licensing model.
Workgroup Server	4-way (32-bit)	A full-function database server designed for small organizations or departmental use with a per-user licensing model.
Workgroup Server Unlimited	4-way (32-bit)	A full-function database server designed for small organizations or departmental use with a per-processor licensing model.
Enterprise Server	Unlimited (32-bit and 64-bit)	A full-function database server designed for the enterprise with flexible scale-up and scale-out capability.

► Express Edition

Express Edition is the newest member of the product family and is a specially tailored full-feature relational database for small and medium businesses designed to be embedded as a transparent part of a business solution. This edition can be deployed on servers with up to two processors and is only available for 32-bit platforms.

► Workgroup Server Edition

Workgroup Server Edition is designed for small organizations and departmental use that supports a small number of connected users. Workgroup Server Edition uses a licensing model designed to provide an attractive price point for smaller installations while still providing a full-function database server.

This edition is ideal for deployment on servers with up to four 32-bit processors and when the number of concurrent or registered users is small (25–30 users), as a license is required for each concurrent or registered user in addition to the base server license.

► Workgroup Server Unlimited Edition

The Workgroup Server Unlimited Edition offers the same features provided with Workgroup Server Edition, but with a simpler per-processor licensing model. This edition is ideal for deployment on 32-bit servers with up to four processors when the number of concurrent or registered users is large (more than 25–30 users) or unknown (Web servers), since no additional license is required for each concurrent or registered user.

- ▶ **DB2 UDB Enterprise Server Edition**

DB2 UDB Enterprise Server Edition is designed for large enterprises that require a database management system that can scale to support multi-terabyte databases. It is available in both 32-bit and 64-bit versions. There is a migration path that allows you to migrate databases from 32-bit to 64-bit versions without having to rebuild the database from scratch.

The Enterprise Server Edition is ideal for deployment on servers of all sizes, as it provides a number of features that are not available in the Workgroup Server Edition of the product.

In general, this edition is required when you need support for more than four processors, have a large number of concurrent or registered users, or you require the scalability provided by 64-bit platforms and the Database Partitioning Feature (DPF).

DPF allows you to partition a database within a single system or across a cluster of systems. This feature provides support for very large databases (VLDBs) by partitioning the database into smaller parts that can more easily be managed by individual partitioned database servers (engines). This approach not only supports very large databases with complex workloads, it also provides great scalability for database administration tasks, such as loads, backups, and recovery.

2.7 Scale-up versus scale-out

The goal of system scalability is to increase performance at a rate that is proportional to increases in system resources for a given workload.

There are two methods to achieving system scalability. They are:

- ▶ **Scale-up:** Increasing the capacity of the single system image by adding (in particular) processors, memory, and disk.
- ▶ **Scale-out:** Adding systems that can be managed and run together.

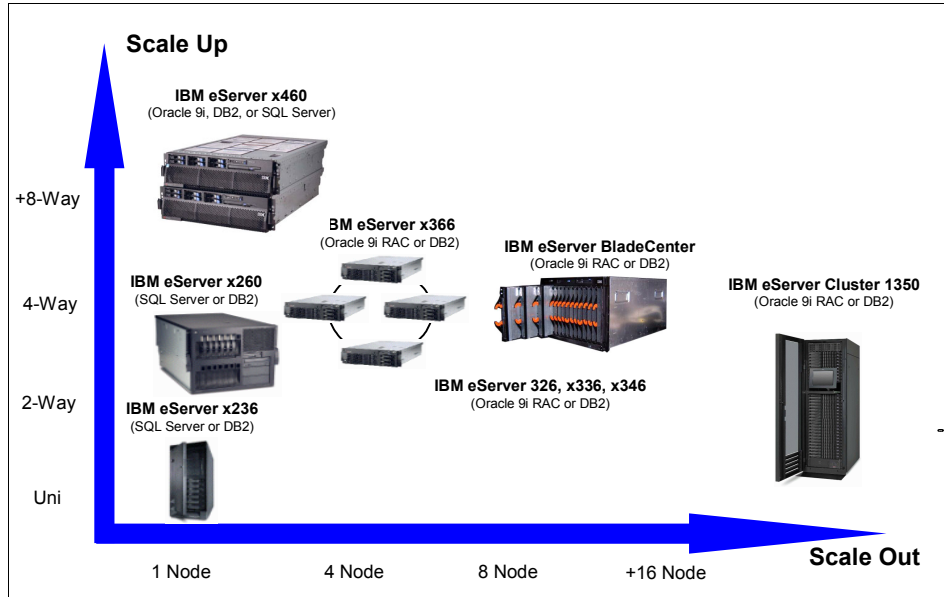


Figure 2-5 Scale-up versus scale-out

Scale-up

Scaling up is achieved by adding resources, such as memory, processors, and storage, to an existing system that runs on a single server. It is also referred to as vertical scaling. The benefit to scaling up is that it is relatively easy, as in general it only requires hardware or software that is designed to take advantage of additional memory, processor, and storage resources.

For example, your database server may start out on a 2-way SMP system with four GB of memory and six hard drives. As the database grows in size or the number of users increase, you can easily scale-up by adding additional processors, memory, and disk resources to maintain the same level of performance. You may eventually need to replace the server with one that is capable of supporting more resources, but today on xSeries servers, you can scale-up to systems that support 32 processors and 64 GB of memory on 32-bit versions of operation systems and 512 GB of memory on operation systems with 64-bit extension EM64T.

NUMA architecture, when compared to other architectures, provides near linear scalability and minimum overhead in resource management that limits the scalability of a single large systems when you add processors, memory, and storage.

The xSeries x460 server is a good example of an SMP system based on Enterprise X-Architecture and NUMA technologies. The server starts with a base 2-way configuration and you can add incremental capacity as your needs grow to a maximum of 32 processors. Likewise, memory can be expanded from 2 GB to 512 GB. This modular server architecture delivers investment protection without the upfront costs of expensive switch-based alternatives.

Advantages of scale-up:

- ▶ Easier to configure and administer
- ▶ Good when most queries access small blocks of data
- ▶ Best for applications that maintain state (OLTP)
- ▶ Add CPU and memory as required (scheduled downtime especially for CPUs)
- ▶ All tools and queries work as expected
- ▶ Can be maintained by lesser skilled DBAs

Disadvantages:

- ▶ Requires higher cost hardware
- ▶ The database has finite capacity limits tied to hardware
- ▶ Must balance CPU, memory, and I/O to achieve peak performance
- ▶ Fail-over cluster server usually configured equal to primary

Databases such as Microsoft SQL Server are well suited to a scale-up approach using a multi-node x460 configuration.

Scale-out

Scale-out means add discrete servers to your server “farm” to gain more processing power. Although many options exist for implementing a farm comprised of small low-end servers, we consider the use of the IBM @server BladeCenter or 1U rack servers such as the xSeries 336 as the most viable alternative when discussing this requirement.

Scale-out is sometimes called horizontal scaling, as well as in general being referred to as clustering. However, clustering can sometimes be ambiguous in that there are distinct types of clusters, which include high availability, load balancing, and high-performance computing.

Load balancing is the goal of scaling out. That is to say, we scale-out by adding one or more servers to an existing system to balance the system load as we add additional demands on the system. For example, your database server may start out on a 2-way system with 4 GB of memory and six hard drives. As the database grows in size or the number of users increase you scale-out by adding another server with two processors, 4 GB of memory, and six disk drives to maintain the same level of performance. Although you do not necessarily have to

add an additional server with the exact specifications, it does reduce the complexity of scaling out.

The benefit to scaling out is that you can achieve near linear scalability. That is, as you add each additional server to the system, you effectively increase your system capacity proportionally. Thus, scaling out provides much better returns in terms of the additional costs associated with adding additional servers to the system. Another benefit inherent with scaling out is that a cluster of smaller servers generally costs less than a single large system.

The drawback to scaling out is that it requires system and database administrators that understand the technology well enough so that it can be implemented effectively. Another drawback is that clustering requires software that is specifically designed for the task.

Advantages of scale-out:

- ▶ Uses lower cost hardware.
- ▶ Scaling is near linear.
- ▶ The database size is not gated by hardware.
- ▶ Preferred when queries access large blocks of data.
- ▶ Best for serving stateless applications (Web).

Disadvantages:

- ▶ Requires more skilled DBA to maintain cluster.
- ▶ Management and scheduling more complex.
- ▶ Depends upon intelligent data partitioning.
- ▶ Introduces query overhead.
- ▶ Maintenance activities require downtime.
- ▶ Cluster aware applications can be much more expensive than stand-alone version.

Architectures for scaling out

There are two distinct approaches to scaling out database management systems. These are generally referred to as a *shared architecture* and a *shared-nothing architecture*. Both architectures attempt to achieve the same goal. That is, to implement a database management system that consists of a cluster of servers, provides linear scalability, and appears as single database to the end users.

A *shared* architecture attempts to accomplish this goal while sharing the database. As additional servers are added to the system, they all share or attempt to share the same database, which resides on shared storage, hence the name shared architecture.

Oracle is an example of a database application that implements a shared-disk approach.

Figure 2-6 shows an example of a scale-out shared architecture.




Figure 2-6 Scale-out shared architecture

On the other hand, a *shared-nothing* architecture accomplishes the same goal by dividing a large database into smaller and more manageable “parts,” called partitions. The term shared-nothing simply refers to the fact that as additional servers are added to the system, each server manages a clearly defined portion of the database.

Figure 2-7 shows an example of a scale-out shared nothing architecture.




Figure 2-7 Scale-out shared nothing architecture

The fact that the database is partitioned should not imply that the system cannot be implemented on shared storage.

IBM DB2 and Microsoft SQL Server both implement a shared-nothing approach.

Choosing scale-up or scale-out

Microsoft SQL Server is well-suited for scale-up configurations, such as a multi-node x460 configuration. It follows a single server, shared nothing approach and it is a high performance solution for Windows environments

Oracle uses a shared-disk approach and is suited to scale-up or scale-out. It is a leading solution for middle market UNIX, Windows, and Linux environments. Scale-out capabilities can be extended with Oracle 9i RAC.

DB2 is suited to scale-up or scale-out. It is developed following a multi-server, shared nothing approach, and is the highest performing database environment for mainframe, UNIX, and Linux environments. Scale-up is preferred for smaller databases (150-200 GB). For larger databases, large block I/O, data warehousing and decision support applications, use a scale-out deployment.



Hardware planning

In this chapter, we discuss the topics you need to understand before you finalize the configuration of your x460, x366, or x260 solution. The topics covered are:

- ▶ 3.1, “Processor subsystem” on page 90
- ▶ 3.2, “Memory subsystem” on page 95
- ▶ 3.3, “Multi-node configurations” on page 104
- ▶ 3.4, “RSA II SL and BMC initial setup” on page 119
- ▶ 3.5, “Storage options” on page 122
- ▶ 3.6, “Power considerations” on page 126
- ▶ 3.7, “Performance tuning and optimization” on page 129
- ▶ 3.8, “ServerProven” on page 130
- ▶ 3.9, “Solution Assurance” on page 131

3.1 Processor subsystem

All three X3 Architecture servers use Intel Xeon MP processors, featuring the 64-bit EM64T technology and driven by a 667 MHz dual front side bus.

The x260 and the x366 are available with up to four Intel “Cranford” processors. The x460 support up to four Intel “Potomac” processors, but can scale-up further by connecting up to seven MXE-460 expansion enclosures to form a single system image of up to 32 processors. These multi-node complexes are discussed further in 3.3, “Multi-node configurations” on page 104.

Table 3-1 lists the processors supported by each server.

Table 3-1 Xeon MP processors used in the X3 Architecture servers

Speed	L2 / L3 cache	x260	x366	x460	MXE-460	Part number
Intel “Cranford” processors						
3.16 GHz	1 / 0 MB	Supported	Supported	No	No	13N0694
3.66 GHz	1 / 0 MB	Supported	Supported	No	No	13N0695
Intel “Potomac” processors						
2.83 GHz	1 / 4 MB	No	No	Supported	Supported	13N0715
3.00 GHz	1 / 8 MB	No	No	Supported	Supported	13N0714
3.33 GHz	1 / 8 MB	No	No	Supported	Supported	13N0713
Intel “Paxville” dual-core processors						
2.67 GHz	1+1 / 0 MB	No	Supported	Supported	Supported	25R8941
3.00 GHz	2+2 / 0 MB	No	Supported	Supported	Supported	25R8942

Table 3-2 lists the processors standard in each server.

Table 3-2 Processors standard in each model

Server	Server	Core speed	Cores/socket	Installed	Maximum	Part number
x260	8863-1RY	3.16 GHz	Single-core	1	4	13N0694
x260	8863-2RY	3.66 GHz	Single-core	1	4	13N0695
x366	8863-1RY	3.16 GHz	Single-core	1	4	13N0694
x366	8863-2RY	3.66 GHz	Single-core	1	4	13N0695
x366	8863-3RU	2.67 GHz	Dual-core	1	4	25R8941

Server	Server	Core speed	Cores/socket	Installed	Maximum	Part number
x366	8863-4RU	3.00 GHz	Dual-core	1	4	25R8942
x460	8872-1RY	2.83 GHz	Single-core	2	4	13N0715
x460	8872-2RY	3.00 GHz	Single-core	2	4	13N0714
x460	8872-3RY	3.33 GHz	Single-core	2	4	13N0713
x460	8872-5RU	2.67 GHz	Dual-core	2	4	25R8941
x460	8872-6RU	3.00 GHz	Dual-core	2	4	25R8942
MXE-460	8874-1RY	Match the x460	Single-core	0	4	Match the x460
MXE-460	8874-2RU	Match the x460	Dual-core	0	4	Match the x460

Note: The Y and U in the model numbers is for countries in North and South America. For EMEA, substitute G (for example, 1RG). For Asia-Pacific countries, the letter varies from country to country. Consult the announcement letter or the *xSeries Configuration and Option Guide*, found at:

<http://www.pc.ibm.com/support?page=SCOD-3ZVQ5W>

The following guidelines apply to each of these servers:

- ▶ All processors must be the same type, speed, and L2/L3 cache size.
- ▶ The Voltage Regulator Modules (VRM) for processors 1 and 2 are integrated on the microprocessor board. The VRMs for processor 3 and 4 are shipped with the microprocessor option and must be installed in the appropriate sockets. If you use this option to install a CPU in sockets 1 or 2 (in the MXE-460 for example), the VRM in the package is not used.
- ▶ The installation order for the microprocessors and the position of the VRM sockets are depicted in Figure 1-10 on page 19.
- ▶ To ensure proper airflow and cooling, each processor socket must be populated with either an processor air baffle, or with a heat sink when a processor is installed.

3.1.1 Dual-core upgrades

The Dual Core X3 Upgrade Kit, part 39Y6580, is available to upgrade any single-core x460 and MXE-460 servers to be dual-core capable. This kit contains a new processor planar.

Processors are not part of the kit and must be purchased separately. All previously installed processors will need to be replaced with dual-core processors, as a mix of processors is not supported.

Note: If a client has implemented a four-node or eight-node configuration using servers with single-core processors, and plans to upgrade their complex to dual-core processors, then they will need to comply with export regulations. In addition, they will need to replace one memory card in each node with part number 40K2450.

This is described in more detail in 3.3.4, “Export regulations for 32-core complexes” on page 118.

3.1.2 Processor configuration options

BIOS includes a number of settings that clients can set related to processor configuration and performance. To access these settings, do the following:

1. Press F1 during system startup when prompted.
2. From the Main Menu, choose **Advanced Setup**.
3. Select **CPU Options**. Figure 3-1 appears.




Figure 3-1 CPU options in BIOS

Hyper-Threading Technology

As described in 1.8, “Processors” on page 18, the processors used in these servers include a technology called Hyper-Threading, which makes a single CPU appear to the operating system as two logical processors, so that they can receive and process two data/instruction streams simultaneously.

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

Note: In x460 multi-node configurations, all nodes in a partition should have the same Hyper-Threading setting. See 3.3.2, “Scalable system setup” on page 110 for more information.

Clustering Technology

For certain operating systems, it is necessary to configure how the routing of processor interrupts in a multi-processor system is handled. It is a low level setting that sets a multi-processor interrupt communication protocol (XAPIC). The settings are functional only, and do not affect performance.

In the Clustering Technology menu, choose the appropriate mode for your operating system, as advised in the operating system installation instructions in Chapter 4, “Operating system installation” on page 135.

Clustering here refers to the ability of the x260, x366, and x460 to have CPUs across multiple processor buses: the processors are “clustered” into pairs of processors, one pair per bus. Each server has two processor buses, and each additional node in an x460 multi-node complex has two extra processor buses.

Note: The term *clustering* here does not refer to the cluster technology provided by such services as Microsoft Cluster Service.

The choices are Logical Mode, Physical Mode, and Special Mode.

- ▶ The Logical Mode is the default mode for the system. It can be used by Windows and Linux.
- ▶ Physical Mode is the required setting for Linux to boot systems when there are 16 or more processor cores in the complex (that is 16-way or larger with the current processors, or 8-way or more when dual-core processors become available).
- ▶ The Special Mode is required for the 64-bit edition of Red Hat Enterprise Linux 3.0 and only when single-core processors are installed. This is the only operating system that needs it and Windows Server 2003 with Service Pack 1 will not even boot in this mode.

When dual-core *Paxville* processors are used, the necessary BIOS update will remove Special Mode from the choices and, if it was previously set, will force it to Logical Mode.

Processor Adjacent Sector Prefetch

When this setting is enabled, which is the default, the processor retrieves both sectors of a cache line when it requires data that is not currently in its cache. When it is disabled, the processor will only fetch the sector of the cache line that contains the data requested.

This setting may affect performance, depending on the application running on the server and memory bandwidth utilization. Typically, it will affect certain benchmarks by a few percent, although in most real applications it will be negligible. This control is provided for benchmark users that wish to fine-tune configurations and settings.

Processor Hardware Prefetcher

By default, hardware prefetching is enabled, which enables the processors to prefetch extra cache lines for every memory request. Recent tests in the performance lab have shown that you will get the best performance for most commercial application types if you disable this feature. The performance gain can be as much as 20% depending on the application.

To disable prefetch, go to BIOS Setup (press F1 when prompted at boot) and select **Advanced Settings** → **CPU** and set HW Prefetch to **Disabled**. For high-performance computing (HPC) applications, we recommend you leave HW Prefetch enabled. Future releases of BIOS that ship to enable dual-core will have HW Prefetch disabled by default.

Processor Execute Disable Bit

Processor Execute Disable Bit is a function of new Intel processors with lets you prevent the execution of data that is in memory as though it were code. When enabled (the default is disable), this will prevent viruses or the like from gaining unauthorized access to applications by exploiting buffer overruns in those applications.

Note: This function is only used for 32-bit operating environments where the processor is in one of the following modes:

- ▶ Legacy protected mode, if Physical Address Extension (PAE) is enabled on a 32-bit operating system.
- ▶ IA-32e mode, when EM64T is enabled on a 64-bit operating system.

The operating system must also implement this function.

If this feature is enabled, and provided the operating system has marked the memory segment as containing data, then the processor will not execute any code in the segment.

This parameter in BIOS is disabled in BIOS by default just in case the applications you run on your server are affected by it (at least one Java application is known to fail if this setting is enabled). For added protection, you may wish to enable it, but you should first test your applications to ensure they will still run as expected before enabling this in a production environment.

For more details of the Execute Disable Bit function, see:

<http://www.intel.com/cd/ids/developer/asmo-na/eng/149308.htm>

3.2 Memory subsystem

This section discusses the memory subsystem: available options, installation rules, and configuration recommendations.

The X3 Architecture servers implement memory using 1-4 memory cards, each of which holds four DIMMs, as shown in Figure 3-2. The servers have one or two memory cards installed as standard (model dependant); additional ones are part number 13M7409.




Figure 3-2 Memory card locations (x460 and x366 chassis shown)

The memory cards connect directly into the system planar and are powered by two separate memory power buses (two memory cards on each bus). As shown in Figure 3-3, memory cards 1 and 2 are on power bus 1, and memory cards 3 and 4 are on power bus 2. The power arrangement is of particular importance with the hot-swap and hot-add functions, as discussed in 3.2.2, “Hot-swap memory” on page 99 and 3.2.3, “Hot-add memory” on page 100 respectively.

The system memory uses ECC DDR2 DIMMs meeting the PC2-3200 standard. The DIMM layout of the memory cards is also shown in Figure 3-3.




Figure 3-3 Memory implementation

The standard memory cards and DIMMs are shown in Table 3-3.

Table 3-3 Standard memory cards and DIMMs

Server	Memory cards	Standard DIMMs
x260	One (card 1)	Two 1 GB DIMMs or two 512 MB DIMMs (model dependant), installed in memory card 1.
x366	One (card 1)	Two 1 GB DIMMs, installed in memory card 1.
x460	One (cards 1 and 2)	Two 1 GB DIMMs, installed in memory card 1. Memory card 2 contains no DIMMs.
MXE-460	Two (cards 1 and 2)	None (all nodes should match).

Supported DIMM options are:

- ▶ 1 GB upgrade kit, part number 73P2865 (contains two 512 MB DIMMs)
- ▶ 2 GB upgrade kit, part number 73P2866 (contains two 1 GB DIMMs)
- ▶ 4 GB upgrade kit, part number 73P2867 (contains two 2 GB DIMMs)
- ▶ 8 GB upgrade kit, part number 30R5145 (contains two 4 GB DIMMs)

Note: At the time of writing, the 1 GB kit (two 512 MB DIMMs) was only supported on the x260. Check ServerProven® for the latest information.

With 4 GB DIMMs, a total amount of 64 GB RAM can be installed. In an x460 eight-node system, the maximum installable memory size is 512 GB (8 x 64 GB).

x460 Multi-node configurations: As discussed in 3.3, “Multi-node configurations” on page 104, with x460 multi-node configurations, 256 MB of memory in each node is allocated to the XceL4v cache. As a result, the memory seen by the operating system is reduced by 256 MB for each node (2 GB for an eight-node complex, for example).

The configuration rules are as follows:

- ▶ A minimum of one memory card containing two DIMMs is required for the server to operate. This also applies to the MXE-460.
- ▶ Memory is two-way interleaved to ensure maximum data throughput. As a result, the DIMMs are always installed in pairs of the same size and type to populate a memory bank. Banks are DIMM sockets 1 and 3 and 2 and 4, as shown in Figure 3-3 on page 96.
- ▶ The installation sequence for memory cards is 1-2-3-4. There are two ways to fill the DIMMs sockets, depending on whether cost or performance is the more important consideration:
 - Cost-effective configuration: To minimize cost, you can install the memory DIMMs by filling each memory card before adding DIMMs to the next memory card.
 - Performance-optimized configuration: As described in 1.4, “IBM XA-64e third generation chipset” on page 8, there are four independent memory ports. Therefore, to optimize performance, you can spread the DIMMs (still installed in matched pairs) across all four memory cards before filling each card with two more DIMMs.

For more performance tuning options, see also 3.7, “Performance tuning and optimization” on page 129.

A more detailed description and the exact sequence for installation can be found in the *User's Guide* of each server.

- ▶ If you are planning to implement a four-node or eight-node complex using models with dual-core processors, you will need to comply with export regulations and install one export-controlled memory card, part number 40K2450, in each node. This is described in more detail in 3.3.4, “Export regulations for 32-core complexes” on page 118.

3.2.1 Memory mirroring

Memory mirroring is available on all X3 Architecture servers for increased fault tolerance. Memory mirroring is operating system independent, since all mirroring activities are handled by the hardware.

To understand what memory mirroring and hot-swap capabilities exist with the server, you must first understand how the memory cards are powered. The system planar has two separate memory power buses that are split between the four memory cards. As shown in Figure 3-3 on page 96, memory cards 1 and 2 are on power bus 1, and memory cards 3 and 4 are on power bus 2.

Mirroring takes place across the two power buses. In other words, the memory DIMMs in cards 1 and 2 are mirrored to the memory DIMMs in cards 3 and 4. With memory mirroring enabled in BIOS, you can hot-swap one memory card at a time on each memory power bus. Once memory mirroring is enabled, the data that is written to memory will be stored in two locations. For read operations, data is read from the DIMMs with the least amount of reported memory errors through memory scrubbing.

Memory scrubbing is an automatic and regular test of all the system memory that detects and reports memory errors before they cause a server outage. If memory scrubbing determines that the DIMM is damaged beyond use, read and write operations are redirected to the remaining good DIMM. Memory scrubbing then reports the damaged DIMM and the light path diagnostic displays the error. If memory mirroring is enabled, then the mirrored copy of the data from the damaged DIMM is used until the system is powered down and the DIMM replaced. After the damaged DIMM is replaced, memory mirroring will copy the mirrored data back onto the new DIMM.

Key configuration rules relating to memory mirroring are as follows:

- ▶ Memory mirroring must be enabled in the BIOS (it is disabled by default).
- ▶ Both memory cards must have the same total amount of memory, and must have identical DIMMs. In other words, DIMMs must be installed in matched quads to support memory mirroring. Partial mirroring is not supported. See the server's *Installation Guide* for information about the exact installation order required.

Important: Because of memory mirroring, you will only have half of the total amount of memory available. If 8 GB is installed, for example, then the operating system sees 4 GB once memory mirroring is enabled (it is disabled in the BIOS by default).

3.2.2 Hot-swap memory

These servers supports hot-swap memory, which means that if a DIMM fails, it can be replaced with a new DIMM without powering down the server. This advanced feature allows for maximum system availability. Hot-swap memory requires that memory mirroring be enabled.

To easily identify whether hot-swap is enabled and the status of power to the memory card, each memory card has a green *memory hot-swap enable* LED, and a green *memory port power* LED on the top panel of the memory card, as shown in Figure 3-3 on page 96. The memory card has eject levers with sensors, so that the system can recognize when a memory card is being removed and power down that card's slot accordingly.

The overall process to hot-swap a failed DIMM is as follows:

1. Verify that memory mirroring and hot-swap are enabled by checking the memory hot-swap enabled LED on the memory cards.
2. When a DIMM fails, you will be alerted via the memory LED on the light path diagnostics panel (and by other means via the service processor if this has been configured).
3. Locate the memory card with the failed DIMM by using the Error LED on the memory card.
4. Remove the memory card containing the failed DIMM.
5. Press the button on the memory card to identify which DIMM has failed. The LED next to the failed DIMMs lights up.
6. Replace the failed DIMM and reinsert the memory card.

For a more detailed description of how to correctly hot-swap memory and which sequence to follow, see the server's *User's Guide*.

3.2.3 Hot-add memory

The hot-add memory feature enables you to add DIMMs without turning off the server. This section shows the requirements for enabling the hot-add memory feature on the server.

Note: Hot-add and hot-swap are mutually exclusive. You can only enable one of these features.

The requirements are as follows:

- ▶ Operating system support: Adding usable system memory to a running operating system requires operating system support. This is done via an ACPI sequence. Currently, the only operating systems that have this capability and are supported on the X3 Architecture servers are the Enterprise and Datacenter editions of Windows Server 2003.
- ▶ Memory hot-add must be specifically enabled in the BIOS setup. When this is done, the system allocates blank windows of memory space for future memory additions. By enabling hot-add, memory mirroring will automatically be disabled.
- ▶ Memory cards 3 and 4 must not be installed yet because these are the *only* ones that can be hot-added.
- ▶ If only one memory card is installed (memory card 1) prior to the hot-add operation, then *only* one additional memory card may be added in slot 3.
- ▶ If two memory cards (1 and 2) are already in the system, then two additional memory cards *must* be added in slots 3 and 4. The card in slot 4 must be inserted first. This is because power to the second bus will be applied once card 3 is inserted and closed.
- ▶ The DIMMs must be added two at a time (matched pairs) and they must also match the equivalent pair of DIMMs on the matching memory card on the other power bus.
- ▶ A minimum of 4 GB of memory must be installed in the server in order for hot-add memory to be available. Additionally, for 32-bit operating systems, the Physical Address Extension (PAE) mode has to be enabled to take advantage of the additional memory.

For information About how to perform a hot-add operation, and more information about the restrictions, see the server's *User's Guide*.

Notes:

- ▶ Once you have added a memory card with two DIMMs, you cannot add more memory to that same memory card without powering off the server.
- ▶ Enabling hot-add reserves a portion of the memory map for the memory that may be hot-added in the future. If you do not plan to use hot-add, we recommend that you not enable this feature in BIOS.

3.2.4 Memory ProteXion: redundant bit steering

Redundant bit steering is the technical term for Memory ProteXion.

When a single bit in a memory DIMM fails, the function known as redundant bit steering (RBS) automatically moves the affected bit to an unused bit in the memory array, removing the need to perform the ECC correction and thereby returning the memory subsystem to peak performance. The number of RBS actions that can be performed depends on the type of DIMMs installed in the server:

- ▶ A pair of single-sided DIMMs can perform one RBS action. These are the 512 MB and 1 GB DIMMs. A pair of single-sided DIMMs is also known as a single Chip Select Group (CSG).
- ▶ A pair of double-sided DIMMs (also known as stacked DIMMs) can perform two RBS actions. These are the 2 GB and 4 GB DIMMs. A pair of double-sided DIMMs is also known as two Chip Select Groups.

RBS is supported in both non-mirrored and mirrored configurations.

In the X3 Architecture servers, DIMMs are installed in matched pairs in banks. Each memory card installed in the server is comprised of two banks:

- ▶ DIMM sockets 1 and 3 form bank 1.
- ▶ DIMM sockets 2 and 4 form bank 2.

Memory errors are handled as follows:

- ▶ If a single-bit error occurs in a CSG, RBS is used to correct the error.
- ▶ If a second single-bit error occurs in the same CSG, the ECC circuitry is used to correct the error.

So, for example, if an x366 is configured with 16x 1 GB DIMMs (which are single-sided), then each pair corresponds to a single CSG, so the server has a total of eight CSGs. This means that the server can survive up to 16 single-bit memory failures, two in each pair of DIMMs (RBS recovery, and then ECC).

As a second and more complex example, if the same server is installed with eight 1 GB DIMMs (single-sided) and eight 2 GB DIMMs (double-sided), then this means there are a total of $4 + 8 = 12$ CSGs (four from the four pairs of 1 GB DIMMs and eight from the eight pairs of 2 GB DIMMs). This means the server can survive up to 24 single-bit memory errors.

The first single bit correctable error on a CSG will result in an RBS event and RBS log entry from the SMI Handler. The second will result in automatic hardware correction and a threshold log entry from the SMI Handler.

3.2.5 Memory configuration in BIOS

Depending on your needs, the system memory can be configured in four different ways:

- ▶ Redundant Bit Steering (RBS) (default)
- ▶ Full Array Memory Mirroring (FAMM)
- ▶ Hot Add Memory (HAM)
- ▶ High Performance Memory Array (HPMA)

You configure the memory subsystem in the server's BIOS Setup menu by selecting **Advanced Settings** → **Memory** → **Memory Array Setting**. The window in Figure 3-4 should appear.




Figure 3-4 Memory options in BIOS

The choices are as shown in Table 3-4.

Table 3-4 Memory configuration modes in BIOS

Mode	Memory ProteXion	Memory-mirroring	Hot-swap memory	Hot-add memory
HPMA (high performance memory array)	Disabled	Disabled	Disabled	Disabled
RBS (redundant bit steering) (default)	Yes	Disabled	Disabled	Disabled
FAMM (full array memory mirroring)	Yes	Yes	Yes	Disabled
HAM (hot-add memory)	Yes	Disabled	Disabled	Yes

The memory configuration mode you select depends on what memory features you want to use:

▶ Redundant Bit Steering (RBS):

This option enables Memory ProteXion and is the default/standard setting. Select **RBS** if you are not using mirroring, hot-swap, or hot-add. See 3.2.4, “Memory ProteXion: redundant bit steering” on page 101 for details of how RBS works.

▶ Full Array Memory Mirroring (FAMM)

Select **FAMM** to enable memory mirroring (and to enable hot-swap).

Memory mirroring reduces the amount of addressable memory by half on each chassis in the partition, but provides complete redundancy of all addressable memory. RBS is available in this mode. See 3.2.1, “Memory mirroring” on page 98 for more information.

▶ Hot-Add Memory (HAM)

Select **HAM** to enable the use of the hot-add in the future.

HAM provides an array layout that supports runtime hot memory add within an OS that supports that feature. This setting has lower performance and may also restrict the amount of memory that can be installed in each chassis, as addressable ranges must be reserved on each chassis for the hot add function. RBS is available in this mode. See 3.2.3, “Hot-add memory” on page 100 for more information.

▶ High Performance Memory Array (HPMA)

HPMA optimizes the installed memory array on each chassis in the partition for maximum memory performance. Hardware correction (ECC) of a single correctable error per chip select group (CSG) is provided, but RBS is not available. See 3.2.4, “Memory ProteXion: redundant bit steering” on page 101 for a discussion of chip select groups.

We recommend you do not select the HPMA setting in a production environment, as this disables Memory ProteXion.

3.3 Multi-node configurations

Note: This section only applies to the x460 and MXE-460. The x260 and x366 do not support multi-node configurations.

The IBM *XpandOnDemand* scalability features give the flexibility to expand the x460 server's capacity in terms of number of CPUs, memory, and I/O slots, as the demand grows.

Clients can expand the server in the following ways:

- ▶ CPUs from two-way up to 32-way
- ▶ Memory from 2 GB to 512 GB
- ▶ PCI-X slots from six to 48

This scalability is achieved by connecting multiple x460s or MXE-460s to the base x460. These *nodes* then form a single complex. The supported expansion steps are listed in Table 3-5.

An x460 server can be configured together with one, three, or seven MXE-460s to form a single 8-way, 16-way, or 32-way complex.

Table 3-5 x460 scalability options

Nodes	CPUs	Maximum RAM	PCI slots	Number of MXE-460s*
1	2-way	64 GB	Six	None
1	4-way	64 GB	Six	None
2	8-way	128 GB	12	One
4	16-way	256 GB	24	Three
8	32-way	512 GB	48	Seven
* Additional nodes can be either MXE-460 or x460 servers				

You can also form multi-node complexes using multiple x460s or combinations of x460s and MXE-460s. With these combinations, you can partition the complex, as described in 3.3.3, "Partitioning" on page 116.

Important: As we discuss in 3.3.4, "Export regulations for 32-core complexes" on page 118, all 32-core or 64-core configurations will require US government approval regardless of where the complex will be installed.

As shown in Figure 3-5, a scalable system comprises of an x460 server and one, three, or seven MXE-460 systems.

A fully configured, eight-node, and scalable system would have 32 processors, 512 GB of memory (using 4 GB DIMMs), 48 PCI-X 2.0 adapters, 3.5 TB of disk space (non-RAID), and 16 Gigabit Ethernet connections.

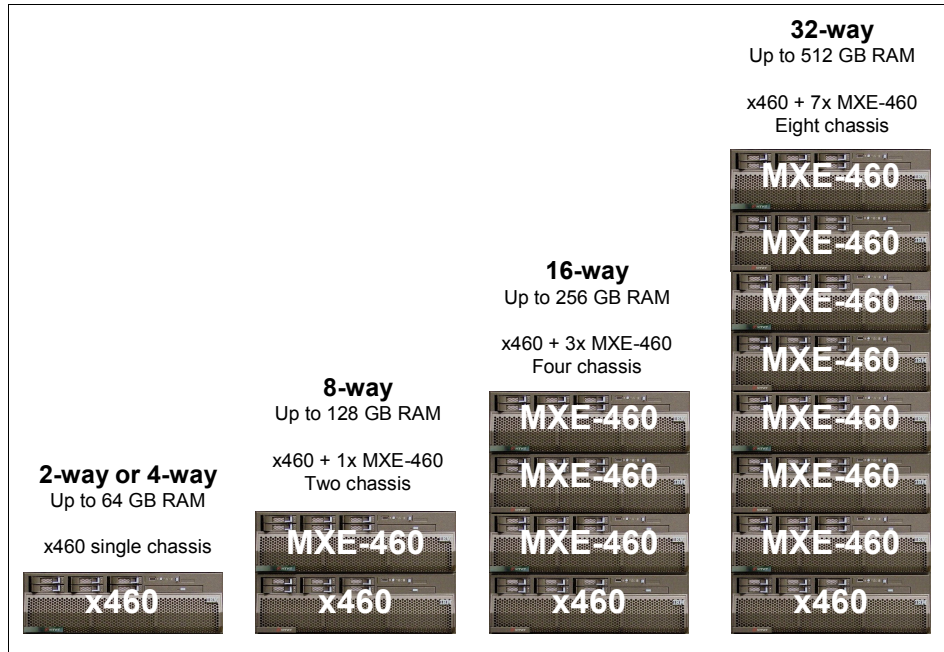


Figure 3-5 The four multi-node configurations supported

The following configuration rules apply:

► Multi-node configurations

The multi-node configurations can have more than one x460. Only one is shown in Figure 3-5 for simplicity. In fact, if you wish to use partitioning as described in 3.3.3, “Partitioning” on page 116, you will need one x460 for each partition to be the primary node.

The first x460 is known as the primary node; all other nodes in a complex are called secondary nodes. The primary node must always be an x460 (MXE-460s cannot be primary nodes).

Only one, two, four, or eight nodes are supported. Other combinations cannot be selected in the configuration utility.

- ▶ Processors

The x460 and all MXE-460s must each have four processors installed and all processors must be the same speed and cache size. The x460 and MXE-460 are technically capable of having less than four CPUs installed in a multi-node configuration, but this type of configuration will require SPORE approval before it can be supported.
- ▶ Memory

For performance reasons, you should have the same amount of memory in each node. A minimum of 2 GB of RAM is required in each node. The x460 and MXE-460 are technically capable of having less than 2 GB of RAM installed in a multi-node configuration, but this type of configuration will require SPORE approval before it can be supported.

In a multi-node configuration, 256 MB of RAM per node is allocated to the Xcel4v cache. In an 8-node, 32-way complex, this means that 2 GB of RAM is allocated to Xcel4v and is unavailable to the operating system.
- ▶ 32-core and 64-core configurations

As described in 3.3.4, “Export regulations for 32-core complexes” on page 118, partitions with 32 single-core processors or 16 or 32 dual-core processors are prevented from being configured unless a special memory card is installed.
- ▶ Firmware

All system firmware, including the system BIOS, diagnostics, BMC firmware, and RSA II SlimLine firmware, must be at the same level across all systems.

Updating the system BIOS in every node in a scalable system can be performed from the primary node. The BIOS code in all secondary nodes will automatically be updated as well. The server diagnostics, as well as the BMC and RSA II SlimLine firmware, must be individually updated on each node, but this can be performed remotely: The RSA II firmware can be updated using the RSA II Web interface or IBM Director. The BMC firmware can be updated with an RSA II remote console session using the remote diskette function.
- ▶ Disk drives installed in any of the MXE-460s are seen by the operating system as normal disk drives.
- ▶ If you install the optional ServeRAID 8i RAID adapter into each node, you will be able to form RAID arrays, but these arrays cannot span nodes. All drives in each array must be local to the RAID adapter.
- ▶ All PCI-X slots and onboard Gigabit Ethernet ports in the MXE-460 are visible to the operating system as well.

- ▶ The concept of forming a single system image from multiple nodes is called *merging* the nodes. Once a multi-node system is started and the nodes are merged, all existing devices on the secondary nodes (COM ports, keyboard, video, mouse, and DVD-ROM drives) are disabled. PCI devices are on all nodes (including the onboard Gigabit Ethernet and USB controllers) available to the operating system.
- ▶ Power control is as follows:
 - Pressing the power button of any node in a complex will power on or off the entire complex.
 - Pressing the reset button on any node in the complex will restart the entire complex.

3.3.1 Cabling

In order to create a multi-node complex, the servers and expansion modules are inter-connected using a high-speed data bus known as the *scalability bus*. The connection uses copper colored *scalability cables*, which are available in two different lengths:

- ▶ 2.3 m (7.5 ft) scalability cable, part number 13M7414
- ▶ 2.9 m (9.5 ft) scalability cable, part number 13M7416

The 2.3 m cable is usually used to connect nodes together except when the longer 2.9 m cable is preferred, such as in the 8-node complex, where it is used to connect nodes that are too far away for the 2.3 m cable (for example, connecting node 1 and node 5).

Note: These cables are not compatible with the x440 and x445 equivalent cables.

To build your multi-node complex, you will need to order the appropriate number of each type of scalability cable for your specific configuration, as in Table 3-6. The x460 and MXE-460 do not come with scalability cables.

Table 3-6 Scalability cable configuration table

Chassis configuration	2.3 m (short) cables needed	2.9 m (long) cables needed
4-way: 1-chassis	None	None
8-way: 2 chassis	Two	None
16-way: 4 chassis	Six	None

Chassis configuration	2.3 m (short) cables needed	2.9 m (long) cables needed
32-way: 8 chassis	Eight	Four

There are three different cabling schemes, depending on the number of nodes used in the complex. These are shown in the following diagrams. In a 2-node configuration, two scalability cables are used to connect both chassis. The second cable provides redundancy for the chassis interconnect as well as a slight performance benefit.

In any multi-node configuration, any *one* scalability cable may fail without impact on the server's operation. In this situation, a warning LED on the light path diagnostic panel will be lit and an event will be logged into the RSA event log, as shown in Figure 3-6

```

SMI reporting Scalability Event:Double Wide Link Down.Chassis Number = 1. Port
Number = 0.
SMI reporting Scalability Event:Link Down.Chassis Number = 1. Port Number = 1.
SMI reporting Scalability Event:Double Wide Link Down.Chassis Number = 2. Port
Number = 0.
SMI reporting Scalability Event:Link Down.Chassis Number = 2. Port Number = 1.

```

Figure 3-6 Messages in the RSA II event log when a scalability cable fails

Figure 3-7 depicts the scalability cabling plan for a 2-node / 8-way configuration.




Figure 3-7 Cabling for a two-node configuration

Figure 3-8 on page 109 depicts the scalability cabling plan for a 4-node / 16-way configuration. This uses just the short (2.3 m) cables.

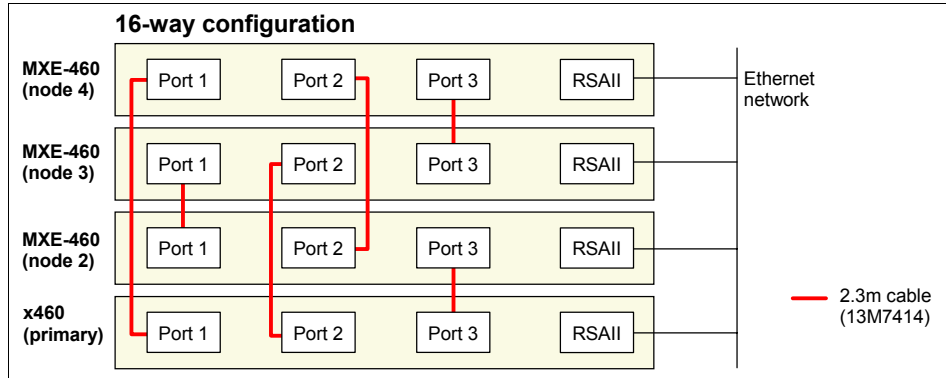


Figure 3-8 Cabling for a four-node configuration

Figure 3-9 depicts the scalability cabling plan for an 8-node / 32-way configuration. This one uses a combination of short (2.3 m) and long (2.9 m) cables.

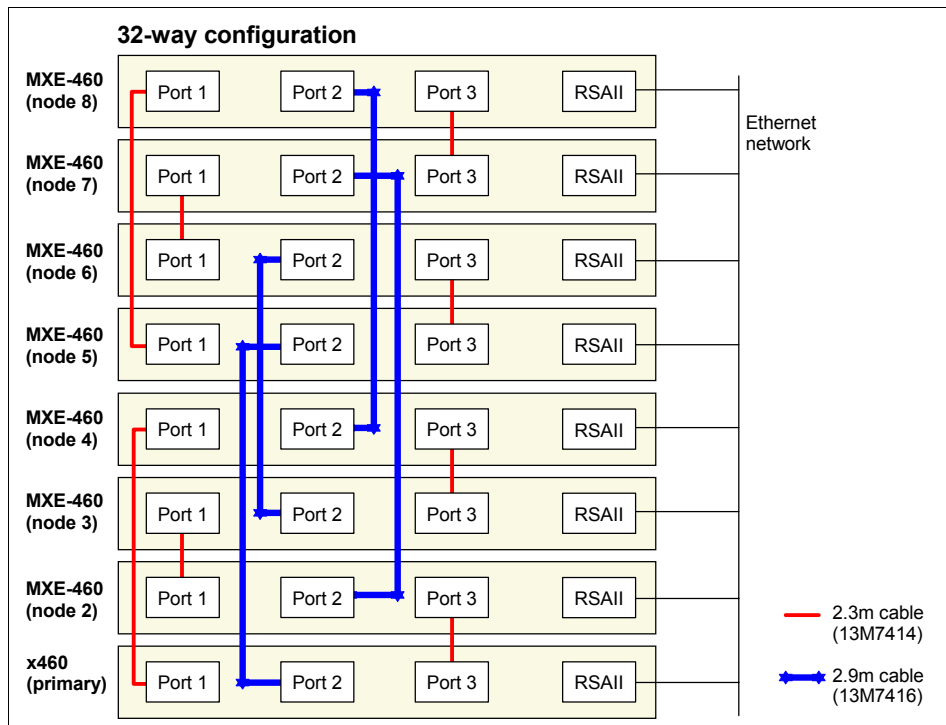


Figure 3-9 Cabling for a eight-node configuration

All Port 2 connectors use the longer 2.9 m cable, since they are further apart. One exception to this is the Port 2 cable between nodes 3 and 6. The recommendation is to use the longer cable so you can insert a KVM or other device in the middle of the complex. The choice of cable is purely a matter of physical connectivity; it does not affect signaling or performance.

3.3.2 Scalable system setup

The setup of a multi-node complex will be configured through the Web interface of the RSA II SlimLine adapter, which is standard with every x460 and MXE-460.

The multi-node configuration data is stored on the RSA in all chassis that are members of the scalable partition. Figure 3-10 on page 111 depicts the RSA Web interface for the scalable partition configuration.

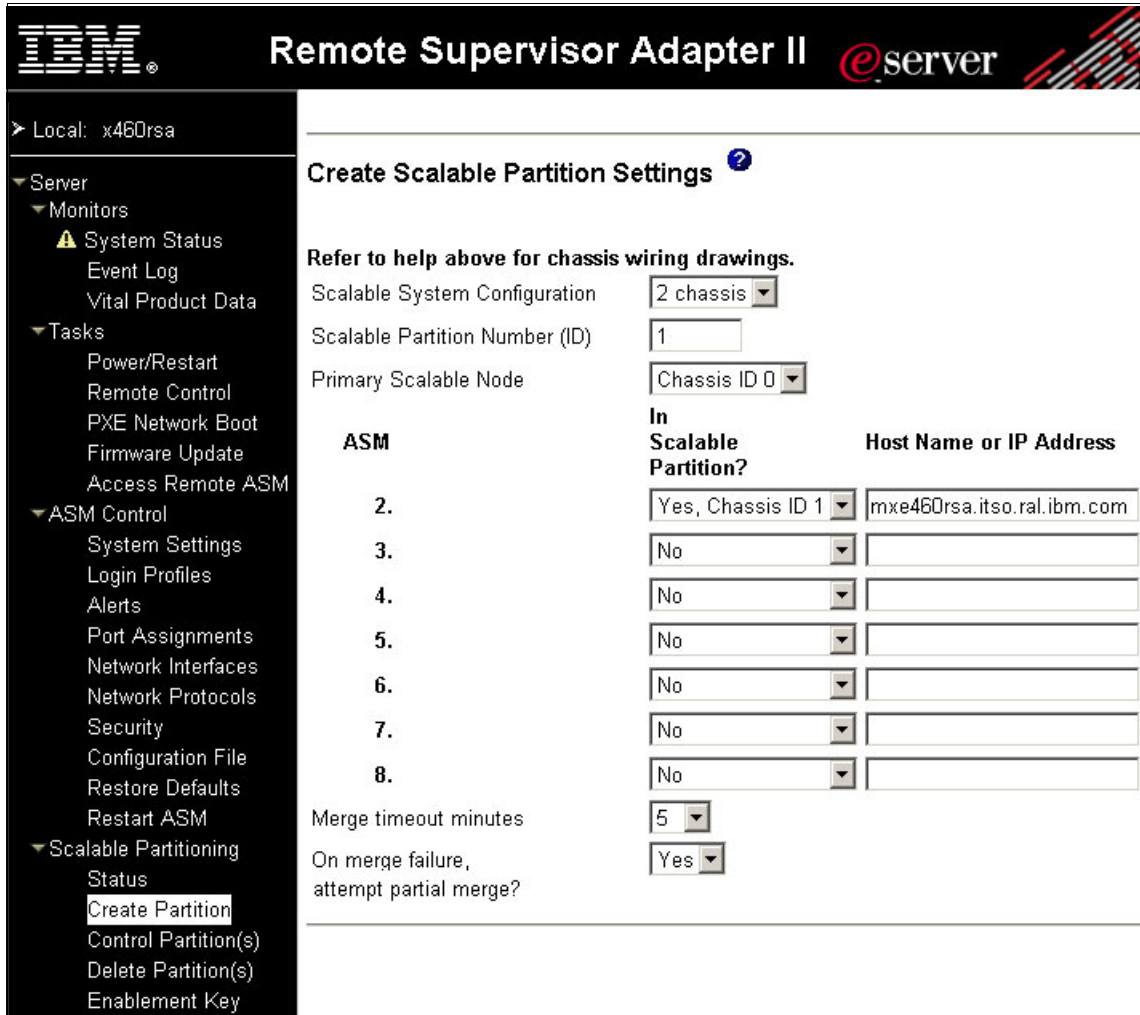


Figure 3-10 Create Scalable Partition Web interface

The communication between the RSA adapters is handled through the systems management Ethernet connections. Therefore, it is very important to ensure secure and reliable network connections. We recommend you connect all RSAs to a separate management network, which is exclusively used for management and not shared with the production or campus LAN.

For proper operation, it is necessary to maintain continuous connectivity between the RSAs. As a result, we recommend you assign static IP addresses to the RSA II cards. If you are using DHCP, you should configure address reservations to ensure the addresses never change.

3.4, “RSA II SL and BMC initial setup” on page 119 describes how to prepare the RSA II adapter.

To set up a scalable system, the following prerequisites must be fulfilled:

- ▶ The firmware on all nodes must be at the same level. This includes system BIOS, diagnostics, BMC firmware, and RSA firmware.
- ▶ The settings made in the Configuration/Setup utility for memory and CPUs must be identical.
- ▶ During the configuration via the RSA II Web interface, all participating nodes must be powered off, but connected to AC power.
- ▶ The operating system must reside on the primary node. This means either internal SAS hard disk drives or host bus adapters (HBAs) for SAN boot must be installed in the primary node.
- ▶ You can only create partitions out of set combinations of nodes. If you select any other combination, you will get an error message when you try to create the configuration. Valid node combinations are:
 - Nodes 0, 1
 - Nodes 2, 3
 - Nodes 4, 5
 - Nodes 6, 7
 - Nodes 0, 1, 2, 3
 - Nodes 4, 5, 6, 7

To create a scalable partition and then activate it, do the following:

1. Ensure all nodes are powered off but that they are connected to AC power.
2. Log in to the RSA of the server that will be the primary node in the complex. Open a Web browser and enter the Ethernet address of the primary node’s RSA II adapter.

Note: Scalable partitions can only be created using the primary node’s RSA interface. It is not possible to create, delete, start, or stop a scalable partition from a secondary node’s RSA interface.

3. In the RSA main menu, select **Scalable Partitioning** → **Create Partition**, as shown in Figure 3-10 on page 111.
4. Choose the number of chassis in the drop-down menu, as shown in Figure 3-11 on page 113.




Figure 3-11 Number of chassis

5. Enter a Scalable Partition Number (ID). This can be any numeric string from 1 up to 4294967295 (which is FFFFFFFF in hex). The scalable partition ID must be unique for each partition in a complex.

Tip: Since the partition ID will be converted into a string in hex format, which is displayed in the Scalable Partition Status windows, we recommend you just use the values 1-8.




Figure 3-12 Scalable Partition ID

6. Select an Chassis ID for the Primary Scalable Node in the drop-down menu, as shown in Figure 3-13. The Chassis ID uniquely identifies each RSA within a partition, although the actual value you choose does not represent any functional value. Within a scalable partition, only the following sets of Chassis IDs are valid: (0,1), (2,3), (4,5), (6,7), or (0,1,2,3) (4,5,6,7).

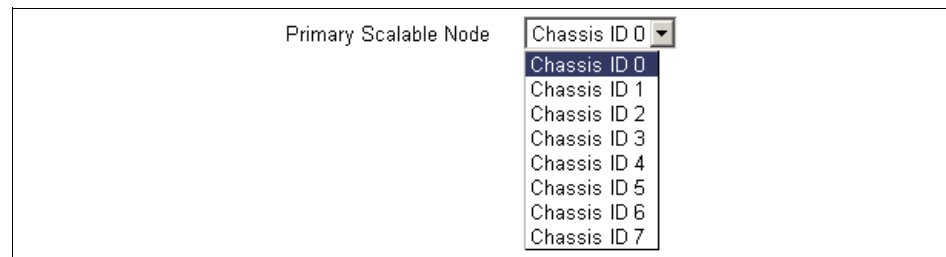


Figure 3-13 Primary node Chassis ID

- In Figure 3-14, choose all the nodes that will join the partition by selecting a valid Chassis ID in the “In Scalable Partition?” drop-down menu. Next to each where you select **Yes**, enter the IP address or host name of the RSA adapter in that node.

If you intend to use host names instead of IP addresses for the RSA adapters, note the following:

- Only fully qualified domain names (FQDN) may be used, for example, `mxe460rsa.itso.ral.ibm.com`.
- The RSA must have a DNS server specified in **ASM Control** → **Network Protocols** → **Domain Name System (DNS)**.

ASM	In Scalable Partition?	Host Name or IP Address
2.	Yes, Chassis ID 1	mxe460rsa.itso.ral.ibm.com
3.	No	

Figure 3-14 Secondary nodes settings

- Select a desired value in the **Merge timeout minutes** drop-down menu, which is shown in Figure 3-15. The default timeout is five minutes, which is suitable in most cases.

Merge timeout minutes	5
-----------------------	---

- 5
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10

Figure 3-15 Merge timeout

The merge timeout determines the time the primary node in the complex will wait for secondary nodes to merge. If for any reason one or more secondary nodes fail to merge, the primary node will boot either stand-alone or with the remaining nodes in a reduced configuration, after the merge timeout has expired.

If you wish to shorten this value to reduce the possible downtime of a multi-node complex for business critical applications, we recommend you wait

until you have finalized the configuration of the multi-node complex and measure the time it takes until the Merge Complete message appears on all nodes in the complex. You can then use a value larger than this time.

- The option On merge failure, attempt partial merge?, shown in Figure 3-16, gives you the choice of whether or not a partial merge attempt shall be made for the case when one or more secondary nodes failed to merge in a 4-node or an 8-node complex.

On merge failure, attempt partial merge?	<input type="button" value="Yes"/> <input type="button" value="No"/>
--	---

Figure 3-16 Partial merge

- Click the **Save** button on the bottom right to save the new, scalable partition configuration.

If any value should be invalid, the communication between the RSAs fails or any other problem occurs, saving the partition configuration will fail and a pop-up error message is shown. This ensures that no invalid configuration can be saved or even become active.

After saving the partition configuration, the status page will be displayed (also available by selecting **Scalable Partitioning** → **Status**) (see Figure 3-17).

Scalable Partition Status ?		
	Current Scalable Partition	New Scalable Partition
Scalable System Configuration	None	2 chassis
Scalable Partition Configuration	None	2 chassis
Scalable Partition Number (ID)		1
Chassis ID	None	0
Local Scalable Node		Primary
Merge Timeout Minutes		5
Attempt Partial Merge		Yes
Scalable Partition State		

Scalable Partition Members			
	SMP ID	Current Hostname	New Hostname
Primary	0		x460rsa
Other(s)	1		mxe460rsa.itso.ral.ibm.com

Figure 3-17 Status after creating a new scalable partition

11. Now that the partition has been created, it must be made active by “moving it” to be a current scalable partition. To do so, select **Scalable Partitioning** → **Control Partition(s)** → **Move New Scalable Partition to Current Scalable Partition**.

After moving the configuration information, the status will be displayed, as shown in Figure 3-18.

Scalable Partition Status ?		
	Current Scalable Partition	New Scalable Partition
Scalable System Configuration	2 chassis	None
Scalable Partition Configuration	2 chassis	None
Scalable Partition Number (ID)	1	
Chassis ID	0	None
Local Scalable Node	Primary	
Merge Timeout Minutes	5	
Attempt Partial Merge	Yes	
Scalable Partition State	Stopped	

Scalable Partition Members				
	SMP ID	Current Hostname	SMP ID	New Hostname
Primary	0	x460rsa		
Other(s)	1	mx460rsa.itso.ral.ibm.com		

Figure 3-18 Status after moving new scalable partition to current scalable partition

Note: Moving the configuration data from New Scalable Partition to Current Scalable Partition overwrites any existing partition configuration!

12. As soon as the partition configuration is finished successfully and looks as shown in Figure 3-18, it can be started by selecting **Scalable Partitioning** → **Control Partition(s)**.

3.3.3 Partitioning

As discussed in 1.3, “Multi-node capabilities and partitioning” on page 4, the complex can be configured as one scalable partition with two, four, or eight nodes. Alternatively, it is also possible to split this complex into multiple independent partitions. For example, an eight-node configuration can be split into two 4-node systems by changing the configuration without changes to the cabling.

The decision whether partitioning is required or not must be made during the planning stage of a multi-node system, since the primary node in a multi-node complex always must be an x460. Configuring multiple partitions on a complex that consists of one x460 with one, three, or seven MXE 460s attached is not supported. You must have one x460 as the primary node for every partition you create.

The actual steps to configure a server partition are exactly the same as described in 3.3.2, “Scalable system setup” on page 110.

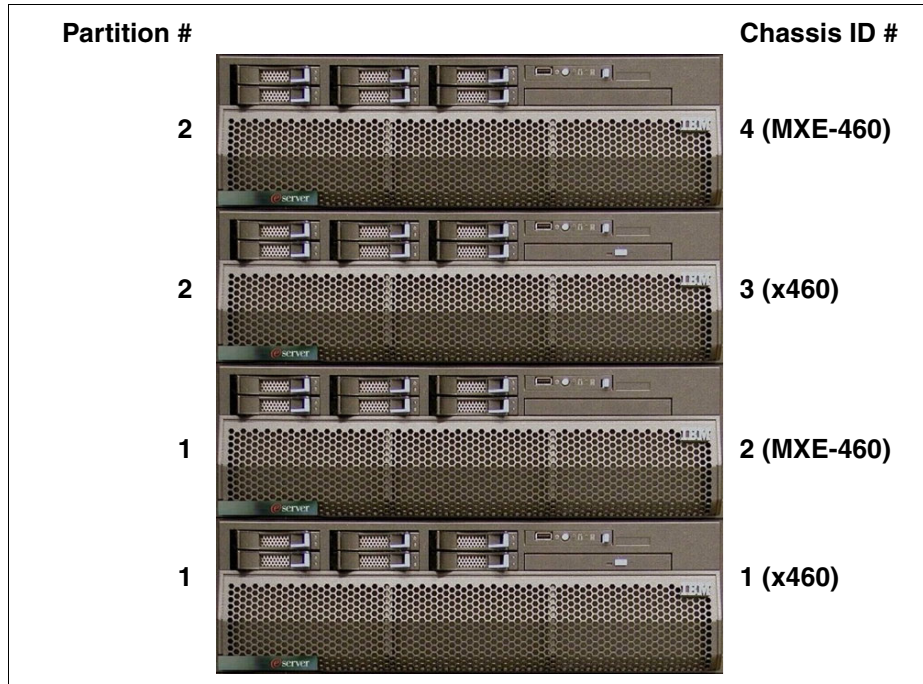


Figure 3-19 Four-node complex split into two partitions

As an example, to configure two 2-node partitions in a four-node complex, do the following:

1. Log in to the RSA Web GUI of chassis 1.
2. Create a new partition with chassis 2 as the secondary node, as described in 3.3.2, “Scalable system setup” on page 110
3. Select **Scalable Partitioning** → **Control Partition(s)** → **Move the new Scalable partition configuration to Current Partition configuration.**

4. Select **Scalable Partitioning** → **Control Partition(s)** → **Move New Scalable Partition to Current Scalable Partition** to make the new configuration active.
5. Log in to the RSA web GUI of chassis 3.
6. Create a new partition with chassis four as the secondary node, as described in 3.3.2, “Scalable system setup” on page 110.
7. Select **Scalable Partitioning** → **Control Partition(s)** → **Move the new Scalable partition configuration to Current Partition configuration**.
8. Select **Scalable Partitioning** → **Control Partition(s)** → **Move New Scalable Partition to Current Scalable Partition** to make the new configuration active.

The two partitions now behave like two independent 8-way servers. Each partition can be controlled independently via the RSA in each primary node. This means they can be powered on and off without any impact on each other. Though they are still wired as a 4-node complex, there will be no data transmitted between the partitions.

Note: As soon as the **New Scalable Partition** configuration is moved to the **Current Scalable Partition** configuration, all previously active settings are lost.

3.3.4 Export regulations for 32-core complexes

Due to the power of the xSeries 460 scalable server, x460/MXE-460 configurations with 32 or more processor cores exceed the US Government mandated composite theoretical performance (CTP) threshold of 190,000 millions composite theoretical operations per second (MTOPS) and require a U.S. export license before you can ship a 32-core or 64-core system into certain countries.

Note: Configurations with 32 or more processor cores are the following:

- ▶ An eight-node (32-way) x460 complex with servers that use single-core processors
- ▶ A four-node (16-way) or eight-node (32-way) x460 complex with servers that use dual-core processors

The scalable systems setup menu prevents such configurations from being created unless a special export-controlled memory card is installed in every node

in the complex. This memory card is called the xSeries 460 Scalability Enablement/Memory Adapter, part number 40K2450.

At least one 40K2450 memory card must be installed in every node. The check for this card is performed at boot time by the RSA II SL service processor in each node (the RSA II SL card needs firmware 26A or later). If a 40K2450 memory card is not present in every node, then the complex will not merge and the systems will boot in stand-alone mode instead.

Note: The enablement CD-ROM, part 31R1558, that was previously used to unlock 32-way configurations, is no longer being used to fulfill this control.

If a 32-core or greater complex is desired, the client and the IBM salesperson or IBM Business Partner will need to work closely with the IBM Export Regulation Office (ERO) and the in-country Export Regulations Coordinator to initiate the licensing process and acquire the documentation and approvals necessary for IBM to be able to ship the system to the intended client. IBM employees can access the ERO at this internal IBM Web site:

<http://w3.ibm.com/chq/ero>

Be aware that this process can be quite lengthy, lasting up to six months. Therefore, the process should be initiated very early in the procurement cycle.

Once the required approvals have been obtained, the client will be eligible to receive the memory card that will permit 32-core configurations. You can order the memory card before approvals have been received, but it will ship only when approvals are in place.

The export-controlled memory card simply replaces the regular memory card and otherwise provides identical function.

3.4 RSA II SL and BMC initial setup

The Baseboard Management Controller (BMC) is standard on all three X3 Architecture servers. The RSA II SlimLine adapter is standard on the x460 and MXE-460, and optional on the x260 and x366.

Both the BMC and the RSA II SlimLine can be initially configured using the BIOS Configuration/Setup utility. Use it to configure the IP addresses of both service processors.

Tip: Since the RSA II SlimLine is optional on the x260 and the x366, when you do install it, you must flash the adapter with the latest firmware for that particular system. See 3.4.1, “RSA II SlimLine setup” on page 120 for more information.

3.4.1 RSA II SlimLine setup

By default, the RSA II SlimLine is configured to obtain an IP address using DHCP and, if that is not available, then to configure a static IP address of 192.168.70.125.

To configure the RSA II SlimLine to use another static address, you can use one of two methods:

- ▶ Access the adapter from a client computer, for example using a mobile computer computer connected to the RSA with a crossover Ethernet cable. Open a Web browser and point it to the adapter’s IP address (192.168.70.125).

Tip: The default logon credentials are USERID and PASSWORD. All characters are uppercase, and the 0 in PASSWORD is a zero and not the letter O).

- ▶ Configure the adapter from the server’s BIOS Setup/Configuration:
 - a. Press F1 during system startup when prompted to enter the Configuration/Setup utility.
 - b. Select **Advanced Setup** → **RSA II Settings**.
 - c. Select **Use Static IP** under DHCP Control and enter a static IP address, subnet mask, and gateway in the appropriate fields as required, as shown in Figure 3-20 on page 121.
 - d. Select **Save Values and Reboot RSA II** to make the configured network settings active.

RSA II Settings	
RSA II MAC Address	00-0D-60-46-D9-FC
DHCP IP Address	000.000.000.000
DHCP Control	[Use Static IP]
Static IP Settings	
Static IP Address	[009.042.171.230]
Subnet Mask	[255.255.255.000]
Gateway	[009.042.171.003]
OS USB Selection	[Linux OS]
Periodic SMI Generation	[Disabled]
Save Values and Reboot RSA II	
<<<RESTORE RSA II DEFAULTS>>>	

Figure 3-20 RSA setup within Configuration/Setup utility

Alternatively, an IP address also can be assigned using DHCP by selecting **DHCP Enabled** in the DHCP Control field. If DHCP is used, we recommend you use an IP reservation. The IP address assigned by the DHCP server is also displayed in this window in the **DHCP IP Address** field.

The RSA's MAC address is displayed as RSA II MAC Address, which can be helpful for network switch configuration.

In the OS USB Selection field, select the appropriate value for your operating system. This setting determines how the RSA presents an emulated keyboard and mouse in a remote control session to the operating system.

- ▶ For Microsoft Windows and Novell NetWare, select **Other OS**.
- ▶ For Linux, select **Linux OS**.

The Periodic SMI Generation setting is set to **Disabled** by default and should not be changed on the X3 Architecture servers. This feature was intended for support of older operating systems that did not include adequate checking of CPU states. Modern operating systems poll for CPU machine checks without this feature. No function is lost by disabling it.

3.4.2 BMC IP address setup

The Baseboard Management Controller (BMC) is common in all three servers. We strongly recommend configuring an IP address for the BMC, even when an RSA is installed, so you can access and clear the BMC event log if it fills up. For further information about how to access the BMC and perform tasks on it, please refer to 5.3, “Baseboard Management Controller” on page 177.

The IP address configuration is done in the Configuration/Setup utility:

Restriction: The BMC cannot be configured via DHCP.

1. Press F1 during system startup when prompted to enter the Configuration/Setup utility.
2. Select **Advanced Setup** → **Baseboard Management Controller (BMC) Settings** → **BMC Network Configuration**.
3. Enter your static IP address, subnet mask, and a gateway.
4. Select **Save Network Settings in BMC**.

As with RSA the default logon credentials for the BMC are USERID and PASSWORD (with a zero). They can be changed by selecting **Advanced Setup** → **Baseboard Management Controller (BMC) Settings** → **User Account Settings**.

Note: Unlike the RSA II, you cannot administer the BMC with a Web browser. The administrative interface is SMBridge, as described in 5.5, “OSA SMBridge utility” on page 194.

3.5 Storage options

This section discusses the available storage options for the x260, x366, and x460.

3.5.1 Serial Attached SCSI (SAS) subsystem

The onboard disk subsystem is driven by an Adaptec AIC-9410 single chip Serial Attached SCSI (SAS) controller. This controller connects to the internal SAS drive bays via one or two backplanes.

Supported internal disks are described in 3.5.3, “Internal storage” on page 124. For more information about the on-board SAS controller, see 1.11, “Serial Attached SCSI” on page 31.

3.5.2 ServeRAID-8i

The onboard SAS controller does not support any host RAID. To enable RAID support, an optional ServeRAID-8i SAS RAID controller (part 13N2227) can be installed in a dedicated slot (see Figure 3-21).



Figure 3-21 ServeRAID-8i

The ServeRAID-8i supports the following RAID levels for the internal disks:

- ▶ RAID 0
- ▶ RAID 1
- ▶ RAID 1E
- ▶ RAID 5
- ▶ RAID 5EE
- ▶ RAID 6
- ▶ RAID 10
- ▶ RAID 50
- ▶ RAID 60 (the x366 and x460 do not support this RAID level because it requires a minimum of eight drives)

The ServeRAID-8i has the following possible stripe sizes: 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, and 512 KB. The default stripe size is 64 KB. The adapter supports up to 24 logical drives. The ServeRAID-8i does not have any connectors for external SAS devices.

The ServeRAID-8i supports two additional RAID levels compared to other ServeRAID controllers:

- ▶ RAID 6 is a multiparity type that requires a minimum of four drives, of which two drives' worth of capacity are consumed for redundancy. RAID-6 allows the loss of two drives without data loss, but this loss is associated with a substantial performance degradation.
- ▶ RAID 60 is a hierarchical array made up of multiple RAID 6 arrays at the lower level with data striped over these arrays (the 0 in 60) at the upper level. Even though the ServeRAID-8i supports this RAID level, it is not supported on the x460 because it requires a minimum of eight drives.

Note: When an optional ServeRAID-8i SAS RAID adapter is installed in the system, use the existing internal cabling.

For a comparison of features of members of the ServeRAID family, see the Technote *ServeRAID Adapter Quick Reference*, TIPS-0054, available from:

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

If the system is not running applications with high microprocessor requirements, software RAID may be a suitable option. Check with your selected OS vendor to see what the vendor-specific requirements are. We recommend you use hardware RAID for two reasons:

- ▶ Software RAID uses the host CPUs, which may impact the performance of the production applications.
- ▶ Software RAID does not include any form of battery backup in case the server fails in the middle of a write operation.

For hints and tips about using the ServeRAID-8i in the X3 Architecture servers, see the publication *xSeries 260, xSeries 366, xSeries 460, and xSeries MXE-460 FAQ — Hints and Tips*, available at:

<http://www.pc.ibm.com/support?page=MIGR-61395>

3.5.3 Internal storage

Standard removable media are shown in Table 3-7 on page 125.

Table 3-7 Diskette, optical and tape drives

Server	Internal diskette drive	Optical drive	Tape option
x460	None ¹	DVD-ROM	None
x366	None ¹	DVD-ROM	None
x260	Standard	CD-ROM	Two half-high bays
Note 1: The USB diskette drive, part number 05K9276 is a supported option			

All three servers use Serial Attached SCSI (SAS) drives for internal disk storage (see Table 3-8).

Table 3-8 Internal disk storage

Server	Disks	Number of bays	Standard disks ¹
x460	2.5" hot-swap SAS drives	Six standard	None
x366	2.5" hot-swap SAS drives	Six standard	None
x260	3.5" hot-swap SAS drives	Six standard, six optional	None
Note 1: Express models may include hard disks			

The x260 has bays for six 3.5" SAS disk drives hot-swap drives, but supports twelve with an optional 6-drive 3.5" SAS backplane, part number 13M7864.

Using the maximum number of hard drives, the storage maximums shown in Table 3-9 can be achieved for the given disk drive capacities.

Table 3-9 Maximum capacities (no RAID or RAID-0)

Server / Disk drives	36.4 GB	73.4 GB	146.8 GB	300 GB
x260	436 GB	880 GB	1761 GB	3600 GB
x366	218 GB	440 GB	880 GB	1800 GB
x460 (per node)	218 GB	440 GB	880 GB	1800 GB

The x260 also has two 5.25" half-height (1.6" high) drive bays that can also be combined to form a single, full-height (3.2") drive bay. These bays are primarily used to install a tape drive. For a list of supported drives, see the ServerProven page at:

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

As well as the drive, clients will also need to install a SCSI adapter to attach to the tape drive, as none is installed.

3.5.4 External storage options

To expand the storage capacity of each server, external storage devices can be attached. Storage devices are available based on Ultra-320 SCSI, iSCSI, and Fibre Channel. Since no external SCSI or Fibre Channel connections are available as a standard, an optional SCSI adapter or Host Bus Adapter has to be installed in order to connect to these external solutions.

IBM supports the following external storage devices:

- ▶ The ServeRAID-6M SCSI RAID controller can attach up to two EXP400 Storage Expansion Unit. The EXP400 can be equipped with up to 14 Ultra320 146 GB disk drives, such that a maximum of approximately 4 TB disk space can be attached.

Note: The ServeRAID-8i is for internal SAS drives only. It does not offer connection to external storage.

- ▶ IBM TotalStorage® DS300 iSCSI, using an optional iSCSI initiator adapter or a software iSCSI initiator via a regular supported Ethernet adapter.
- ▶ IBM TotalStorage DS400 Fibre Channel attached by the IBM TotalStorage FC2-133 Host Bus Adapter (HBA).
- ▶ Using the FC2-133 HBA, several further storage solutions, such as the IBM TotalStorage DS family, can be attached.
- ▶ Ultra320 SCSI Controller attaches external tape drives respectively, for example, tape libraries.

ServerProven lists all supported adapters and external storage options at:

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

3.6 Power considerations

This section will cover several power related topics, such as available power upgrade options, power redundancy, and recommendations for handling power needs in different environments. All X3 Architecture servers support 220V and 110V power, but IBM recommends you use 220V power for the x366, x460 and the MXE-460.

Power is delivered as follows:

- ▶ The x260 and the x366 can be equipped with additional hot-swap power supplies to provide a redundant power configuration.
- ▶ The x460 and MXE-460 ship with two redundant hot-swap power supplies standard.

Table 3-10 shows the default and the maximum configuration for each server.

Note that unlike the power supplies in the x366 and x460, the power output of the power supply unit in the x260 is independent of the source voltage and is 775W both at 110V and 220V.

Table 3-10 Hot-swap power supplies

Server	x260	x366	x460	MXE-460
Standard power supplies	Two 775W	One 1300W	Two 1300W	Two 1300W
Maximum supported	Three 775W	Two 1300W	Two 1300W	Two 1300W
Redundancy requirements	Three power supplies	Two power supplies	Two power supplies	Two power supplies
Power output from each power supply	775W at 110V 775W at 220V	650W at 110V 1300W at 220V	650W at 110V 1300W at 220V	650W at 110V 1300W at 220V
Part number	24R2656	13M7413	Not applicable	Not applicable

A special note about x260 power: The x260 supports up to three power supplies.

- ▶ To provide sufficient power for all server configurations, two power supplies are required (110V and 220V). This is the standard configuration.
- ▶ If you require a redundant power supply (which we recommend), then you should have three power supplies installed.

3.6.1 Working in 110V environments

At 220V, all of these servers will operate with power supply redundancy when both power supplies are installed in the x366, x460, and MXE-460. You get redundancy in the x260 with three power supplies.

However, at 110V, the power supplies of the x366 and x460 may be insufficient for heavily loaded configurations (that is, servers with many CPUs, DIMMs, drives, and adapters):

- ▶ If one power supply is installed, it may not supply enough power; in this case, the OVERSPEC (over specification) LED will light on the operator panel
- ▶ If two power supplies are installed, the second one may not provide power redundancy; in this case, the NONRED (non-redundant) LED will light on the operator panel.

The two standard power supplies of the x260 are sufficient to power all configurations of the server and three power supplies is always enough to provide power redundancy, both in 110V and 220V installations.

For example, the following x366 configuration can be powered by one power supply and would run redundant with two power supplies:

- ▶ One processor
- ▶ Two PCI-X adapters
- ▶ Three hard disk drives
- ▶ Four 1 GB DIMMs

It is not a trivial exercise to calculate in advance the power requirements of a specific configuration to determine how many power supplies are needed to drive the server and how many to make the power redundant. As a result, there are two options clients can take:

- ▶ Use 220V power, either a supply direct from the utility company or via a step-up transformer (if using the latter, ensure your power circuit has a sufficient current rating).
- ▶ Install the server with all options needed, and test this configuration under real life conditions to see if the NONRED or OVERSPEC LEDs light up.

If neither the NONRED or the OVER SPEC LED on the light path diagnostic are lit at any time, your server can operate properly with the power supplies currently installed.

For simplicity of configuration, we recommend you use 220V power.

Summary: With 110V power, heavily configured x366 and the x460s may require two power supplies for normal operation and may not be redundant. 220V power is recommended.

3.7 Performance tuning and optimization

The X3 Architecture servers are designed to deliver superior performance. This section explains how the performance can be improved further by tuning certain settings in the Configuration/Setup Utility or choosing the optimal way to install hardware options.

3.7.1 Optimal memory module installation

The X3 servers support a maximum number of four memory cards. Up to four memory DIMMs can be installed in each memory card. Each memory card is driven by its own memory interface with a maximum bandwidth of 5.3 GBps.

Aspects of the memory configuration that can affect performance include:

- ▶ Spread the installed DIMMs over as many memory cards as possible, preferably up to the maximum of four memory cards. Because each memory card has a dedicated link to the memory controller, the performance gain in doing this can be significant. For example, with four DIMMs in a single memory card, the x366 is about 50% slower than with four DIMMs across two memory cards (two in each).

From a performance point of view, we do not recommend you operate the server with only one memory card installed.

- ▶ Use as many DIMMs as possible. The optimal configuration is to have all 16 slots in the server populated with DIMMs. The gains are not as significant as with the use of memory cards (perhaps a 3% to 5% improvement) but this method can still make a difference if this memory configuration is suitable for your client.

3.7.2 Memory settings in BIOS

By default the Memory Array Settings is set to Redundant Bit Steering (RBS), as shown in Figure 3-22, which provides better protection from multiple single-bit errors.




Figure 3-22 Memory settings

However, if maximum of performance is required, choose **High Performance Memory Array (HPMA)** in the menu **Advanced Setup** → **Memory Settings**. In this mode, only one single-bit correctable error can be recovered in a chip select group using ECC circuitry. See 3.2.5, “Memory configuration in BIOS” on page 102 for details.

3.7.3 CPU settings in BIOS

By default, the X3 Architecture servers are optimized for database transaction processing. This is achieved by enabling hardware prefetching on the processors, which forces the processors to prefetch extra cache lines for every request.

If you plan to run applications that do not take advantage of prefetch, such as Java, file/print, or a Web server, then you can gain 10% to 20% by disabling prefetch. To disable prefetch, select **Advanced Setup** → **CPU Options** and set Processor Hardware Prefetcher to **Disabled**.

The default is Enabled. This setting affects all processors in the chassis.

For other CPU settings, refer to 3.1.2, “Processor configuration options” on page 92.

3.7.4 PCI adapter placement

In x460 multi-node configurations, performance can also be impacted by the installation of PCI cards as network adapters, Fibre Channel HBAs, and so on. To distribute the load equally, we recommend you spread the placement of the adapters across all the nodes.

3.8 ServerProven

The ServerProven program indicates the IBM and third party hardware components and software that are supported by specific servers. More information about this program can be found at:

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

3.9 Solution Assurance

Solution Assurance is a technical inspection of a completed solution design by uninvolved third parties who are experts in its technology components, with the purpose of answering three questions:

- ▶ Will the solution work? That is, is it technically viable?
- ▶ Can we implement it successfully? Is the implementation plan sound?
- ▶ Will it meet the client's requirements and expectations?

The key aspects of the Solution Assurance process are the following:

- ▶ The inspection examines an entire solution, not merely the configuration of an individual product. Solutions typically are comprised of a mixture of hardware, software, and services. They have a life cycle that transcends that of a sales opportunity. Clients must be able to continue to operate a solution in production for its designated life span. Solution Assurance therefore considers:
 - Hardware components
 - Software components
 - Activities and services necessary to successfully implement the solution in a client's environment
 - Activities and services necessary for the client to successfully operate the solution in a production environment on a continuing basis

A key document in this process is the *Solution Assurance Product Review Guide*.

- ▶ The inspection concerns itself solely with the technical viability of the proposed solution. It does not address aspects such as pricing, contracting, terms and conditions, profitability, or assessing the client's financial stability, except insofar as these aspects can affect technical issues. These elements are dealt with elsewhere in the Quality Assurance process. Neither does the inspection address the availability of necessary resources, although it does verify that all necessary resources have been identified.
- ▶ The Solution Assurance process requires a completed solution. It occurs after the designers deem the design finished.
- ▶ The subject matter experts who act as Solution Assurance reviewers should not be part of the solution design team. It is an important element of Solution Assurance that those doing the technical assessment bring fresh eyes to the solution so that they can examine not only the mechanics of the solution itself, but also double-check all of the assumptions that were made in creating the solution.

- ▶ The technical reviewers must possess expertise in the technology components. Most solutions involve many components, which implies that most reviews will be conducted by a panel of reviewers that collectively cover the necessary technical areas. The word *expert* also implies that the reviewers must have credentials that command the respect of the design team. Otherwise, their judgments will not carry weight.

The purpose of the three questions listed at the beginning of this introduction is to form the terms of reference for the review panel. These questions attempt to ascertain whether or not the solution will work as designed, whether there is a feasible and complete implementation plan to implement the solution in the client's environment, and whether the combination of these two elements will indeed satisfy all known client requirements and expectations.

Solution Assurance is not:

- ▶ A session to design a solution
- ▶ The first time technical details are considered
- ▶ Detailed configuration assistance
- ▶ A business case review
- ▶ A client financial health review
- ▶ A way to get an IBM IT Specialist involved
- ▶ A transfer of responsibility for the solution to the Solution Assurance reviewers
- ▶ A cure-all that will catch all problems or replace the need for good solution design

Not all solutions require a review by independent technical subject matter experts, also known as an *expert review*. The criteria used to determine when that is necessary are subject to periodic change. The current criteria are determined by referring to the Solution Assurance Web site or by using the Trigger Tool, linked to from the Web site. If an expert review is needed, it should be scheduled after the solution design, and before the proposal is due to be delivered to the client.

Even if your proposal does not require a Solution Assurance Review, you should still review the items in the Solution Assurance check list to ensure that you have covered everything.

The Solution Assurance Web sites are as follows:

- ▶ For IBM employees:
<http://w3.ibm.com/support/assure>

- ▶ For Business Partners:

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

Select **Design Solutions** on the left and then select **Solution Assurance** at the bottom from the links provided. The above link requires a PartnerWorld® ID and password. If you do not have this information or have forgotten it, please contact PartnerWorld. For contact details, see:

http://www.ibm.com/partnerworld/pwhome.nsf/weblook/cpw_index.html

3.9.1 Solutions Assurance Reviews

Some level of Solution Assurance Review (SAR) should be performed on all IBM solutions. The level of SAR (self, peer, or expert) should match the complexity of the solution. For example, simpler solutions may need only a self review. However, a combination of the client 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.

The three levels of Solution Assurance are:

- ▶ Self

In a self review, the solution designer checks his or her own work using a general or product-specific checklist, as appropriate.

- ▶ Peer

A peer review is conducted by a peer or colleague of the solution design team using available checklists, general and product-specific. The solution designer explains the solution to the peer reviewers, who assess the technical viability of the solution and record their action items and comments in the checklist.

- ▶ Expert

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

The xSeries SAR Trigger Criteria determine if an expert-level review is required for an x460 solution. See the SA document SA200 for details. IBM employees can use the following link:

<http://w3.ibm.com/support/assure/assur30i.nsf/PubA11Num/SA200?OpenDocument>

For further information about what is required, refer to the following documents in the Solution Assurance Web sites listed on page 132 for more information in your specific geography:

- ▶ For the Americas, reference document SA447 or contact TechXpress at 1-800-426-5525 in North America or 770-835-9700 in Latin America.
- ▶ For EMEA, reference documents SA424 or SA359.
- ▶ For Asia Pacific, reference document SA441 or go to:
<http://w3.ibm.com/support/ap/>
Select **SA Contacts/SA Information**.



Operating system installation

This chapter describes operating system specifics and system hardware settings that you may need to configure prior to installing an operating system. The topics in this chapter are:

- ▶ 4.1, “Supported operating systems” on page 136
- ▶ 4.2, “Updating BIOS and firmware” on page 140
- ▶ 4.3, “Available memory in multi-node configurations” on page 141
- ▶ 4.4, “Microsoft Windows Server 2003” on page 142
- ▶ 4.5, “Datacenter offerings” on page 147
- ▶ 4.6, “Microsoft Windows 2000 Server” on page 152
- ▶ 4.7, “Red Hat and SUSE LINUX” on page 153
- ▶ 4.8, “VMware ESX Server 2.5.1” on page 155

The installation instructions for many of the supported operating systems can be found at the OS installation matrix at:

<http://www.pc.ibm.com/support?page=MIGR-4QLNTQ>

4.1 Supported operating systems

For the latest operating system support information, refer to the ServerProven operating system support matrix at:

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

The operating systems shown in Table 4-1 are planned to be supported or are now supported.

Table 4-1 Operating system support (current and planned)

Operating system	x260	x366	x460	MXE-460
Microsoft Windows Server 2003, Standard Edition	Yes	Yes	Yes	Yes
Microsoft Windows Server 2003, Standard for EM64T	Yes	Yes	Yes	Yes
Microsoft Windows Server 2003, Enterprise Edition	Yes	Yes	Yes	Yes
Microsoft Windows Server 2003, Enterprise for EM64T	Yes	Yes	Yes	Yes
Microsoft Windows Server 2003, Datacenter Edition	No	No	Yes	Yes
Microsoft Windows Server 2003, Datacenter for EM64T	No	No	Yes	Yes
Microsoft Windows 2000 Advanced Server	Yes	Yes	Yes	Yes
SUSE LINUX Enterprise Server 9 for x86	Yes	Yes	Yes	Yes
SUSE LINUX Enterprise Server 9 for EM64T	Yes	Yes	Yes	Yes
Red Hat Enterprise Linux 3 AS for x86	Yes	Yes	Yes	Yes
Red Hat Enterprise Linux 3 AS for EM64T	Yes	Yes	Yes	Yes
Red Hat Enterprise Linux 4 AS for x86	Yes	Yes	Yes	Yes
Red Hat Enterprise Linux 4 AS for EM64T	Yes	Yes	Yes	Yes
VMware ESX Server 2.5.1	Yes	Yes	Yes	Yes
Novell NetWare 6.5	Yes	Yes	Yes	Yes

Specific service packs or updates may be required, and some operating systems may be limited as to the number of processors and nodes supported.

4.1.1 Operating systems scalability

The X3 Architecture servers support (or plan to support) operating systems up to 32-way configurations, as described in Table 4-2.

Table 4-2 Operating systems scalability

Operating system	4-way	8-way	16-way	32-way
Windows Server 2003 SP1 Standard for x86	Yes	No	No	No
Windows Server 2003 SP1 Standard for x86-64	Yes	No	No	No
Windows Server 2003 SP1 Enterprise for x86	Yes	Yes	No	No
Windows Server 2003 SP1 Enterprise for x86-64	Yes	Yes	No	No
Windows Server 2003 Datacenter for x86	Yes	Yes	Yes	Yes
Windows Server 2003 Datacenter for x86-64	Yes	Yes	Yes	Yes
Red Hat Enterprise Linux AS 3 for x86	Yes	Yes	No	No
Red Hat Enterprise Linux AS 3 for x86-64	Yes	Yes	No	No
Red Hat Enterprise Linux AS 4 for x86	Yes	Yes	Yes	No
Red Hat Enterprise Linux AS 4 for x86-64	Yes	Yes	Yes	No
SUSE LINUX Enterprise Server 9 for x86	Yes	Yes	Yes	No
SUSE LINUX Enterprise Server 9 for x86-64	Yes	Yes	Yes	No
VMware ESX Server 2.5.1	Yes	Yes	Yes	No
Windows 2000 Server, Advanced Server	Yes	Yes	No	No
NetWare 6.5	Yes	No	No	No

4.1.2 Hyper-Threading support

As described in 1.8, “Processors” on page 18, Hyper-Threading Technology allows a single physical processor to appear to the operating system and applications as two logical processors. The logical processors share the core processing engine of the physical processor and can execute code streams concurrently.

Operating systems must be Hyper-Threading aware before they can see the additional processors. If they are aware, they will see twice as many CPUs as there really are.

In addition, to fully take advantage of Hyper-Threading, the operating system's scheduler must allocate resources with the full knowledge of which logical processors are part of the one physical CPU. For example, consider a server with two physical processors with Hyper-Threading enabled (that is, four logical processors). If the operating system's scheduler is unaware of Hyper-Threading ("Yes" in the Hyper-Threading column in Table 4-3 on page 139), it would treat all four logical processors the same. As a result, if two processes are eligible to run, the scheduler might put those processes on the two logical processors that are part of the one physical processor. Thus, one physical CPU would be busy while the other CPU is idle, leading to poor overall performance.

Simply enabling Hyper-Threading might not guarantee improved overall system performance. To get a benefit from Hyper-Threading, the operating system and server applications need to be capable of detecting the additional logical processors and spawning multiple threads that can exploit the additional processing power.

In addition to considering whether the operating system you are installing supports Hyper-Threading, there can be licensing implications to weigh before enabling Hyper-Threading technology.

For a more detailed discussion of Hyper-Threading Technology, refer to:

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

Table 4-3 on page 139 lists the level of support for Hyper-Threading technology provided by the operating systems.

In the Hyper-Threading column:

- ▶ *Yes* indicates that the operating system recognizes the logical processors and can execute threads on them, but is not optimized for Hyper-Threading. From a licensing perspective, if Hyper-Threading is enabled, the operating system will need to be licensed for twice the number of physical processors to take full advantage of the processors' capabilities.
- ▶ *Optimized* indicates that the operating system recognizes the logical processors and that the operating system code has been designed to take full advantage of the technology. From a licensing perspective, the logical processors do not count toward the number of processors for which an operating system is licensed.

Table 4-3 Supporting Hyper-Threading technology by operating systems

Operating system	Release	Hyper-Threading
Windows 2000 Advanced Server	SP4	Yes
Windows Server 2003, Standard Edition	SP1	Optimized
Windows Server 2003, Standard for EM64T	SP1	Optimized
Windows Server 2003, Enterprise Edition	SP1	Optimized
Windows Server 2003, Enterprise for EM64T	SP1	Optimized
Windows Server 2003, Datacenter Edition	Initial	Optimized
Windows Server 2003, Datacenter for EM64T	Initial	Optimized
Red Hat Enterprise Linux AS for x86	3 Update 2	Optimized
Red Hat Enterprise Linux AS for EM64T	3 Update 2	Optimized
Red Hat Enterprise Linux ES for x86	3 Update 5	Optimized
Red Hat Enterprise Linux ES for EM64T	3 Update 5	Optimized
SUSE LINUX Enterprise Server for x86	9	Optimized
SUSE LINUX Enterprise Server for EM64T	9	Optimized
VMware ESX Server	2.5.1	Yes
NetWare	6.5	Optimized

Hyper-Threading Technology is enabled by default on x260, x366, and x460. To disable it, if necessary, do the following:

1. Press F1 during system startup to enter the System Configuration Utility.
2. From the main menu, select **Advanced Setup** → **CPU Options**. Figure 4-1 appears.




Figure 4-1 Hyper-Threading setting

3. With Hyper-Threading Technology selected, press the right arrow key to change the value to **Disabled**.
4. Save changes and exit the System Configuration Utility.

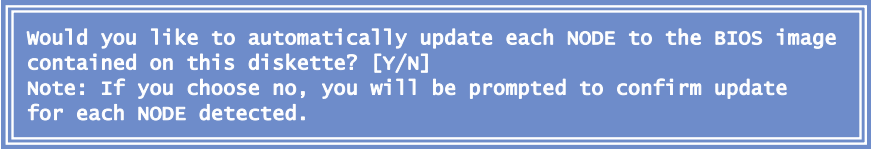
On x460 multi-node configurations, all the nodes in a partition must have the same Hyper-Threading setting. This will need to be set individually on each node.

For descriptions of the other options in this menu, see 3.1.2, “Processor configuration options” on page 92.

4.2 Updating BIOS and firmware

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

- ▶ System BIOS. On x460 multi-node configurations, you can update the BIOS version on all chassis from the primary node. You will be prompted as shown in Figure 4-2.



Would you like to automatically update each NODE to the BIOS image contained on this diskette? [Y/N]
Note: If you choose no, you will be prompted to confirm update for each NODE detected.

Figure 4-2 BIOS update on an x460 multi-node configuration

- ▶ Onboard diagnostics. On x460 multi-node configurations, you will need to update the Diagnostics on each node separately. If you have an RSA II installed, you can use the Remote Console feature with a remote diskette and boot each node from the diagnostics diskette.
- ▶ BMC firmware. On x460 multi-node configuration, you will need to update the firmware separately on each node. If you have an RSA II installed, you can use the Remote Console feature with a remote diskette and boot each node from the BMC firmware diskette.
- ▶ Remote Supervisor Adapter II SlimLine firmware. You can update the RSA firmware from its Web interface. The update involves uploading two separate PKT files before you restart the adapter. In the x460 multi-node configuration, you should update the Diagnostics separately on each node
- ▶ Additional devices if installed, such as ServeRAID adapters and Fibre Channel host bus adapters.

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

<http://www.pc.ibm.com/support?page=MIGR-4JTS2T>

Follow the installation instructions provided with each package.

If you use Microsoft Windows Datacenter, use the driver and firmware matrix located on the following page:

<http://www.pc.ibm.com/support?page=MIGR-4P7RG3>

There are a number of methods you can use to update this code, including:

- ▶ Local diskette drive (x260) or external USB diskette drive, part number 05K9276 (x366 and x460)
- ▶ The Remote Console feature of the RSA II SlimLine adapter (when installed), along with the remote diskette feature. See the redbook *IBM @server xSeries and BladeCenter Server Management*, SG24-6495 for details.
- ▶ Using the bootable UpdateXpress CD-ROM (if supported; see the list of servers), available from:

<http://www.pc.ibm.com/support?page=MIGR-53046>

4.3 Available memory in multi-node configurations

As described in 1.9, “XceL4v cache” on page 27, the X3 Architecture servers use a new technology called XceL4v Dynamic Server Cache as a replacement of the L4 cache used in x440 and x445 servers.

In the x260, x366, and single-node configurations of the x460, the XceL4v cache works as a snoop filter to reduce traffic between the two front-side buses and it uses the memory integrated into the chipset for storage.

In multi-node configurations of x460, 256 MB of main memory in each node is allocated to the XceL4v cache, in addition to building in a chipset for use as an L4 cache for saving memory addresses of data, which are currently located in all the processor’s caches; this reduces traffic between different nodes.

The available memory in multi-node configuration will be less than installed according to Table 4-4.

Table 4-4 Amount of subtracted memory in multi-node configuration

Number of nodes	Total amount of memory allocated to XceL4v
One	0
Two	512 MB
Four	1 GB
Eight	2 GB

For example, when you boot a two-node configuration, you will see the messages shown in Figure 4-3.

Chassis Number	Partition Merge Status	Installed Memory
1	Primary	8GB
2	Merged	8GB

Partition merge successful

16384 MB Memory: Installed
 00512 MB Memory: Consumed by Scalability
 8 Processors Installed

Figure 4-3 Memory allocated to the XceL4v scalability cache

4.4 Microsoft Windows Server 2003

There are three different Microsoft Windows Server 2003 releases:

- ▶ Windows Server 2003: For 32-bit and EM64T/AMD64 systems
- ▶ Windows Server 2003, x64: For EM64T and AMD64-based servers
- ▶ Microsoft Windows Server 2003 64-bit: For Itanium systems

The X3 Architecture servers use Intel Xeon MP processors with EM64T 64-bit extensions. As such, the 32-bit releases and x64 releases of Windows Server 2003 are supported. The 64-bit Itanium release, which Microsoft refers to as

64-bit as opposed to x64, are not supported on the X3 Architecture servers. The Itanium release only runs on servers with the Itanium 2 processor, such as the xSeries 455.

Each of the above three releases come in three or four editions: Standard, Enterprise, Datacenter, and Web. These editions are compared in Table 4-5. Datacenter Edition is discussed in more detail in 4.5, “Datacenter offerings” on page 147.

The latest supported levels of Windows can be found at the ServerProven site: <http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/matrix.shtml>

Table 4-5 Features of the Windows Server 2003 family

Features	Standard Edition	Enterprise Edition	Datacenter Edition	Web Edition
Edition availability				
32-bit release	Yes	Yes	Yes	Yes
x64 release (EM64T & AMD64)	Yes	Yes	Yes	No
64-bit release (Itanium)	Yes	Yes	Yes	No
Scalability				
Processors	1-4	1-8	1-32	1-2
Number of x460 nodes	One	Two	Eight	One
Memory — 32-bit	4 GB	32 GB	64 GB	2 GB
Memory — x64 (64-bit)	32 GB	1 TB ¹	1 TB ¹	N/A
Hyper-Threading	Yes	Yes	Yes	Yes
Hot-add memory	No	Yes	Yes	No
NUMA support	No	Yes	Yes	No
Directory Services				
Active Directory	Yes	Yes	Yes	Limited
Metadirectory support	No	Yes	Yes	No
Security Services				
Internet Connection Firewall	Yes	Yes	No	Yes
PKI, Certify services, Smart Cards	Limited	Yes	Yes	Limited

Features	Standard Edition	Enterprise Edition	Datacenter Edition	Web Edition
Terminal services				
Remote Desktop for Administration	Yes	Yes	Yes	Yes
Terminal Server	Yes	Yes	Yes	N/A
Terminal server Session Directory	No	Yes	Yes	No
Clustering technologies				
Network Load Balancing	Yes	Yes	Yes	Yes
Cluster Service	No	1-8 nodes	1-8 nodes	No
Communication and networking				
VPN Support	Yes	Yes	Yes	Limited
Internet Authentication Service (IAS)	Yes	Yes	Yes	No
Network Bridge	Yes	Yes	Yes	No
Internet Connection Sharing	Yes	Yes	No	No
IP6	Yes	Yes	Yes	Yes
File and print services				
Distributed File System (DFS™)	Yes	Yes	Yes	Yes
Encrypting File System (EFS)	Yes	Yes	Yes	Yes
Shadow Copy Restore	Yes	Yes	Yes	Yes
Removable and Remote Storage	Yes	Yes	Yes	No
Management services				
Remote Installation services (RIS)	Yes	Yes	Yes	No
Remote OS Installation	Yes	Yes	Yes	Yes
Note 1: The x460 is limited to 64 GB per node. This means 128 GB in a two-node configuration and 512 GB in an eight-node configuration.				

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

A custom Hardware Abstraction Layer (HAL) is required for all Windows implementations, when installed on a multi-node x460. This HAL is included in Service Pack 1 for all supported 32-bit versions and also in 64-bit versions of Windows Server 2003.

4.4.1 Installing Windows Server 2003

Instructions on how to install Windows Server 2003 and the necessary drivers are available on the xSeries support site in the support documents listed in Table 4-6.

The latest supported levels of Windows can be found at the ServerProven site:

<http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/matrix.shtml>

Tip: The URLs associated with these support documents are of the following form, where xxxxx is the number listed in the table:

<http://www.pc.ibm.com/support?page=MIGR-xxxxx>

In the PDF version of this redbook, the links in the table below are clickable.

Table 4-6 xSeries support documents for Windows Server 2003

Server	Driver matrix	Windows Server 2003 installation instructions	
		x64 editions	32-bit editions
xSeries 460	MIGR-59923	MIGR-60676	MIGR-61178
xSeries 366	MIGR-59142	MIGR-59608	MIGR-62043
xSeries 260	MIGR-61523	MIGR-62497	MIGR-62498
All servers	MIGR-4JTS2T	MIGR-4QLNTQ	MIGR-4QLNTQ

For installation instructions not listed, see:

<http://www.pc.ibm.com/support?page=MIGR-4QLNTQ>

As well as the normal method of using the Windows installation CD-ROM, you can also install Windows using ServerGuide, available from:

<http://www.pc.ibm.com/support?page=MIGR-4ZKPPT>

Pre-install information

The key points to the installation are as follows:

- ▶ You should update all firmware before starting. Download them from the above driver matrix Web page. This includes:
 - System BIOS
 - Diagnostics
 - BMC firmware
 - RSA II firmware (if installed)
 - ServeRAID firmware (if installed)
- ▶ If you are using a ServeRAID-8i RAID controller, use the supplied CD to configure an array or download the latest configuration CD from the above driver matrix page.
- ▶ In BIOS, ensure the following parameters are set:
 - In CPU Options, ensure that the Clustering Technology parameter is set to **Logical Mode**, as shown in Figure 3-1 on page 92.
 - In **Advanced Settings** → **RSA II Settings**, ensure that the OS USB Selection setting is set to **Other OS**, as shown in Figure 5-3 on page 167.
- ▶ If you plan to install Windows using a regular Windows installation CD, you will need a USB diskette drive to supply the necessary boot device drivers.
- ▶ If you are planning to boot from the internal SAS drives and will be installing using the Windows installation CD-ROM, you will need to press F6 when you see the Setup is inspecting your computer's hardware configuration message, and then insert a driver diskette. The required drivers are one of the following:
 - Adaptec SAS drivers if you are booting from the internal SAS drives and you do not have a ServeRAID-8i installed
 - ServeRAID-8i driver if you are booting from the internal SAS drives and you do have a ServeRAID-8i installed (SAS driver is not needed.)

Download these drivers from the above driver matrix URLs. If you are using ServerGuide to install Windows, you do not need to obtain these drivers separately.

If you do not have a USB diskette drive, but you do have the RSA II adapter installed in the server, you can also install Windows remotely using the

remote console and remote media functions of the RSA II. In this instance, you would put the Windows CD-ROM in the server and the driver diskette in the drive of a remote workstation and control the server remotely via the RSA II Web interface. See 3.5, “Remote console and remote media”, in the IBM Redbook *IBM @server xSeries and BladeCenter Server Management*, SG24-6495 for details.

Post-installation information

The key points to the installation are as follows:

- ▶ After installation, you will need to install additional drivers. Consult the post-install steps in the installation instructions (from the above URLs). In addition, you will need to install:
 - The OSA IPMI driver (for the BMC)
 - The RSA II driver
- ▶ If you are installing a 32-bit version of Windows and you have more than 4 GB of RAM installed, you should add the /PAE switch to the boot.ini file once installation is complete, so that the operating system can access the memory about the 4 GB line (see the last line in Example 4-1).

Example 4-1 boot.ini example for accessing more than 4 GB memory

```
[boot loader]
timeout=3
default=multi(0)disk(0)rdisk(1)partition(1)\WINDOWS
[operating systems]
multi(0)disk(0)rdisk(1)partition(1)\WINDOWS="Windows Server 2003, Enterprise" /fastdetect /PAE
```

4.5 Datacenter offerings

Windows Server 2003, Datacenter Edition is for highly scalable network operating systems designed for mission-critical, enterprise-wide applications. Datacenter Edition is suited for four-node x460 configurations: It enables full use of the large number of processors and large amounts of RAM that can be installed. High-volume online transaction processing, large-scale data warehousing, and scientific simulations are some of the applications for which Datacenter is optimized.

The complete IBM Datacenter Solution Program can be found at:

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

The x460 is fully certified to run Windows Server 2003, Datacenter Edition, up to an 8-node configuration. See the Windows Server Catalog for details:

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

There are two options available for clients who want to run Windows Server 2003, Datacenter Edition:

- ▶ IBM Datacenter High Availability Program Offering
- ▶ IBM Datacenter Scalable Offering

Important: If you plan to implement Datacenter on a 32-way (8-node) x460 configuration, see 3.3.4, “Export regulations for 32-core complexes” on page 118. Gaining approvals can take several months.

Option 1: IBM Datacenter High Availability Program Offering

The Datacenter High Availability Program Offering delivers a fully configured, Datacenter-certified, preinstalled hardware and software operating system solution on 4-way through 32-way server configurations that support up to 8 node Microsoft Cluster Certified solutions for a tightly controlled, end-to-end supported environment for maximum availability.

This end-to-end offering will deliver a fully configured, certified, and preinstalled system for clients who want to maintain a tightly controlled environment for maximum availability. To maintain this high availability, the solution must be maintained as a certified configuration.

Table 4-7 on page 149 shows the models for both 32-bit and 64-bit versions of the operating system.

With the High Availability Program Offering, the client selects memory, disks, and other options and IBM installs them at the factory. The x460 Datacenter models come with four processors, but no memory or disks, in order to provide maximum flexibility in configuration. The operating system is preloaded at the factory. Also shipped with the system is a recovery CD, OS documentation, and a 4-CPU Certificate of Authenticity (COA) to license the system.

The MXE-460 models have no processors, memory, or disks (allowing the client to specify any combination), but have a 4-CPU Certificate of Authenticity (COA) to license the system.

Unlike the x440 and x445 IBM Datacenter Offerings, it is no longer necessary to order the operating system separately. With the x460, the operating system is included when you order the system.

Note: Processors, memory, HDDs, ServeRAID, and so forth, must be ordered as part of the Datacenter solution and will be configured in second-level manufacturing before client shipment. The processors in all nodes must match exactly. The HA models will be configured and tested in manufacturing.

Table 4-7 Certified models for the High Availability Program Offering (Datacenter preloaded)

Server	Model	Standard processors	Cores/socket	L2/L3 cache	Std mem/disk
With 32-bit Windows Server 2003 Datacenter Edition preloaded					
x460	8872-2Ax	4x 3.00 GHz Xeon MP	Single-core	1 MB / 8 MB	None / None
x460	8872-3Ax	4x 3.33 GHz Xeon MP	Single-core	1 MB / 8 MB	None / None
MXE-460	8874-1Ax	(Match the x460)	Single-core	1 MB / 8 MB	None / None
x460	8872-5Ax	4x 2.66 GHz Xeon 7020	Dual-core	1+1 MB / 0	None / None
x460	8872-6Ax	4x 3.00 GHz Xeon 7040	Dual-core	2+2 MB / 0	None / None
MXE-460	8874-2Ax	(Match the x460)	Dual-core	Varies	None / None
With 64-bit Windows Server 2003 Datacenter Edition preloaded					
x460	8872-2Bx	4x 3.00 GHz Xeon MP	Single-core	1 MB / 8 MB	None / None
x460	8872-3Bx	4x 3.33 GHz Xeon MP	Single-core	1 MB / 8 MB	None / None
MXE-460	8874-1Bx	(Match the x460)	Single-core	1 MB / 8 MB	None / None
x460	8872-5Bx	4x 2.66 GHz Xeon 7020	Dual-core	1+1 MB / 0	None / None
x460	8872-6Bx	4x 3.00 GHz Xeon 7040	Dual-core	2+2 MB / 0	None / None
MXE-460	8874-2Bx	(Match the x460)	Dual-core	Varies	None / None

Note: Table 4-7 lists models ending in x. The x changes according to geography. For example, 8872-3Ax will be 8872-3AU in the USA and 8872-3AG in EMEA. For country-specific models, check the announcement letter.

Software Update Subscription

Clients should also purchase the annual Software Update Subscription for the IBM Datacenter High Availability Program Offering. The Software Update Subscription provides periodic updates to the Microsoft Windows Datacenter operating system for maintenance and new versions, which you license for a period of one year. This subscription also includes IBM updates to firmware and device drivers certified by IBM and Microsoft for use with the Datacenter Solution.

Subscriptions are cumulative, meaning that if you have a 16-way solution, you will require four 4-CPU subscriptions (see Table 4-8).

Table 4-8 Software Update Subscriptions for the Datacenter High Availability Program

Software Update Subscriptions	Order number
Windows Server 2003, Datacenter Edition, 32-bit, 1-4 CPUs	4816-GAX
Windows Server 2003, Datacenter Edition, 64-bit, 1-4 CPUs	4816-HAX

IBM builds, tests, and provides the complete certified package of these components. IBM provides program updates as they become available for a period of one year. A program update can contain a new version, release, supplements, or service packs, as IBM determines, announced during the subscription term. IBM does not guarantee that updates will be announced during the annual term.

Tip: The Maintenance Update Subscription is no longer needed and has been withdrawn. Maintenance is now included in the Software Update Subscription.

Option 2: IBM Datacenter Scalable Offering

The Datacenter Scalable Offering is ideal for clients who already have a well-managed IT infrastructure and just want a Windows operating system that scales from 4-way to 32-way and offers maximum performance and scalability in a non-clustered environment.

Table 4-9 on page 151 shows the system models for both 32-bit and 64-bit versions of the operating system. With this option, the x460 models come with four processors, and no memory or disks. The system is shipped with the Datacenter installation CD, OS documentation, recovery CD and a 4-CPU Certificate of Authenticity (COA) to license the system.

The MXE-460 comes with no processors, memory or disks, and a 4-CPU Certificate of Authenticity (COA) to license the system.

Unlike the x440 and x445 IBM Datacenter Offerings, it is no longer necessary to order the operating system separately. With the x460, the operating system is included when you order the system.

Table 4-9 Certified models for the Datacenter Scalable offering without operating system

Server	Model	Standard processors	Cores/socket	L2/L3 cache	Std mem/disk
With 32-bit Windows Server 2003 Datacenter Edition preloaded					
x460	8872-2Dx	4x 3.00 GHz Xeon MP	Single-core	1 MB / 8 MB	4 GB / None
x460	8872-3Dx	4x 3.33 GHz Xeon MP	Single-core	1 MB / 8 MB	4 GB / None
MXE-460	8874-1Dx	(Match the x460)	Single-core	1 MB / 8 MB	None / None
x460	8872-5Dx	4x 2.66 GHz Xeon 7020	Dual-core	1+1 MB / 0	4 GB / None
x460	8872-6Dx	4x 3.00 GHz Xeon 7040	Dual-core	2+2 MB / 0	4 GB / None
MXE-460	8874-2Dx	(Match the x460)	Dual-core	Varies	None / None
With 64-bit Windows Server 2003 Datacenter Edition preloaded					
x460	8872-2Ex	4x 3.00 GHz Xeon MP	Single-core	1 MB / 8 MB	4 GB/None
x460	8872-3Ex	4x 3.33 GHz Xeon MP	Single-core	1 MB / 8 MB	4 GB/None
MXE-460	8874-1Ex	(Match the x460)	Single-core	1 MB / 8 MB	None/None
x460	8872-5Ex	4x 2.66 GHz Xeon 7020	Dual-core	1+1 MB / 0	4 GB / None
x460	8872-6Ex	4x 3.00 GHz Xeon 7040	Dual-core	2+2 MB / 0	4 GB / None
MXE-460	8874-2Ex	(Match the x460)	Dual-core	Varies	None / None

Software Update Subscription

Clients should also purchase the annual Software Update Subscription for the IBM Datacenter Scalable Offering. For this offering, the Software Update Subscription provides the client with operating system upgrade protection, allowing the client to upgrade to the next Datacenter operating system release for no additional charge. The Software Update Subscription is a one-year license.

Subscriptions are cumulative, meaning that if you have a 16-way solution, you will require four CPU subscriptions (see Table 4-10).

Tip: Unlike the Datacenter High Availability Program Offering, maintenance updates are not part of the Software Update Subscription. Instead, they are available directly from Microsoft.

Table 4-10 Software Update Subscription for the Datacenter Scalable Offering

Software Update Subscriptions	Order number
Windows Server 2003, Datacenter Edition, 32-bit, 1-4 CPUs	4818-GAX
Windows Server 2003, Datacenter Edition, 64-bit, 1-4 CPUs	4818-HAX

4.6 Microsoft Windows 2000 Server

Instructions on how to install Windows 2000 Server and the necessary drivers are available on the xSeries support site in the support documents listed in Table 4-11.

Tip: The URLs associated with these support documents are of the following form, where xxxxx is the number listed in the table:

<http://www.pc.ibm.com/support?page=MIGR-xxxxx>

In the PDF version of this redbook, the links in the table below are clickable.

Table 4-11 xSeries support documents for Windows 2000 Server

Server	Driver matrix	OS install instructions
xSeries 460	MIGR-59923 MIGR-62044	
xSeries 366	MIGR-59142	MIGR-59215
xSeries 260	MIGR-61523	
All servers	MIGR-4JTS2T	MIGR-4QLNTQ

For installation instructions not listed, see:

<http://www.pc.ibm.com/support?page=MIGR-4QLNTQ>

As well as the normal method of using the Windows installation CD-ROM, you can also install Windows using ServerGuide, available at:

<http://www.pc.ibm.com/support?page=MIGR-4ZKPPT>

The prerequisites and important notes for installing Windows 2000 Server are similar to those of Windows Server 2003, as listed in 4.4.1, “Installing Windows Server 2003” on page 145.

You will need to install Service Pack 4 for Windows 2000 Server to bring the operating system up to a supported level.

In addition, for multi-node x460 configurations, you will need to a custom hardware abstraction layer (HAL) code, available at:

<http://www.pc.ibm.com/support?page=MIGR-61296>

You will also need an external diskette drive to install the HAL during the operating system installation.

4.7 Red Hat and SUSE LINUX

Red Hat Enterprise Linux Advanced Server 3.0 update 5 and SUSE LINUX Enterprise Server 9.0 SP2 are designed to exploit the capabilities of the X3 Architecture servers, including support for Hyper-Threading technology and the 64-bit extensions.

The latest supported levels of Red Hat and SUSE LINUX products can be found at the ServerProven site:

<http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/matrix.shtml>

Instructions on how to install Linux and the necessary drivers are available on the xSeries support site in the support documents listed in Table 4-12.

Tip: The URLs associated with these support documents are of the following form, where xxxxx is the number listed in the table:

<http://www.pc.ibm.com/support?page=MIGR-xxxxx>

In the PDF version of this redbook, the links in the table below are clickable.

Table 4-12 xSeries support documents for Linux

Server	Driver matrix	OS Installation instructions		
		RHEL 3	RHEL4	SLES 9
xSeries 460	MIGR-59923	MIGR-62630	MIGR-62632	MIGR-62720
xSeries 366	MIGR-59142	MIGR-59187 MIGR-62583	MIGR-62584	MIGR-59157 MIGR-62688
xSeries 260	MIGR-61523	MIGR-62582	MIGR-62616	MIGR-62689
All servers	MIGR-4JTS2T	MIGR-4QLNTQ	MIGR-4QLNTQ	MIGR-4QLNTQ

For installation instructions not listed, see:

<http://www.pc.ibm.com/support?page=MIGR-4QLNTQ>

The key points to installing Linux are as follows:

- ▶ You should update all firmware before starting. Download them from the above driver matrix Web page. This includes:
 - System BIOS
 - Diagnostics
 - BMC firmware
 - RSA II firmware (if installed)
 - ServeRAID firmware (if installed)
- ▶ If you are using a ServeRAID-8i RAID controller, use the supplied CD to configure an array or download the latest configuration CD from the above driver matrix page.

- ▶ In BIOS, ensure the following parameters are set:
 - In CPU Options, ensure that the Clustering Technology parameter is set as follows (see Figure 3-1 on page 92):
 - **Special Mode** for 64-bit RHEL 3.0 when using single-core processors
 - **Logical Mode** for all other Linux distributions (including 32-bit RHEL) and for all distributions on dual-core processors
 - Select **Advanced Settings** → **RSA II Settings** and ensure that the OS USB Selection setting is set to **Linux OS**, as shown in Figure 5-3 on page 167.
- ▶ You may need a USB diskette drive to supply the necessary boot device drivers.
- ▶ At the time of writing, PCI Hot Plug was not supported with RHEL 3 on these systems.

4.7.1 Additional information for SLES 9

The following is of importance when installing SUSE LINUX Enterprise Server 9:

- ▶ If the server contains a Remote Supervisor Adapter II, in the boot screen, press F2 and scroll to chose 1024 x 768 resolution.
- ▶ In the boot screen, choose **Installation**. SLES 9 Update 2 already has all the necessary drivers; simply follow the installation wizard. If you are installing a version of SLES 9 older than Update 2, follow the instructions documented in Table 4-12 on page 154 to obtain the necessary driver diskettes.
- ▶ If you are installing to a server that contains a Remote Supervisor Adapter II, select **Skip** at any graphics configuration screens that are displayed throughout the rest of the installation. After installing, do not reconfigure the monitor or any other graphical settings.

If you must change the graphics settings, type the following command to start the configuration utility:

```
sax2 -m 0=fbdev
```

4.8 VMware ESX Server 2.5.1

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

ESX Server supports two-way SMP virtual machines. This support is made with the add-on component VMware ESX Virtual SMP.

Note the following items if you are using ESX Server:

- ▶ ESX Server must be ordered in conjunction with a supported xSeries server. It cannot be ordered separately. If you want to order ESX Server separately, contact VMware directly.
- ▶ Each VMware ESX Server includes a one-year Software Update Subscription Service from VMware. During this period, product and maintenance updates are made available to registered owners. After the first year, this becomes a fee-based offering.
- ▶ In addition to the Software Update Subscription Service, support services are provided on a fee basis through Support Line offerings which, in certain countries, may be mandatory. VMware ESX Server is a supported product under the Microsoft and Linux Support Line offerings. For additional information, visit:

<http://www.ibm.com/eserver/xseries/vmware>

VMware ESX 2.5.1 supports the following maximum configuration:

- ▶ Storage
 - 16 host bus adapters
- ▶ VMware File System (VMFS)
 - VMFS-2: Up to 64 TB
 - VMFS-1: Up to 2 TB
- ▶ CPU
 - Up to 16 physical processors per system
 - Up to 80 virtual CPUs in virtual machines
 - Up to 200 registered virtual machines
- ▶ Memory
 - 64 GB of RAM
- ▶ Adapters
 - Up to 64 adapters of all types,
 - Up to 8 Gigabit Ethernet or 16 10/100 Ethernet ports per system

Table 4-13 on page 157 represents lists of supported guest operating systems in ESX Server 2.5. SMP-capable means that guest operating system supports VMware Virtual SMP for dual virtual CPU or more.

Table 4-13 Supported guest operating systems in VMware ESX 2.5.1

Guest operating system	SMP-capable
Windows Server 2003 (Enterprise, Standard, and Web Editions)	Yes
Windows XP Professional (Service Pack 1 or 2)	No
Windows 2000 Server (Service Pack 3 or 4)	Yes
Windows 2000 Advanced Server (Service Pack 3 or 4)	Yes
Windows NT® 4.0 (Service Pack 6a)	No
Red Hat Linux 7.2, 7.3, 8.0, 9.0	No
Red Hat Enterprise Linux (AS) 2.1 (Update 5) and 3.0 (Update 3)	Yes
SUSE Linux 8.2, 9.0, and 9.1	No
SUSE Linux Enterprise Server (SLES) 8 and 9.0	Yes
Novell NetWare 6.5 (Support Pack 1 and 2) 5, 6.0 (Support Pack 5) and 5.1 (Support Pack 6 and 7)	No
FreeBSD 4.9	No

VMware ESX Server 2.5.1 supports IBM Director 4.21 or higher.

4.8.1 Support for applications running on ESX Server

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

► Microsoft

See the following Microsoft support Web site for details about their support of applications and operating systems running on ESX Server:

<http://support.microsoft.com/kb/897615/>

► IBM software

If you are running IBM software, such as WebSphere, Lotus®, and Tivoli® products on VMware ESX Server, you must have an IBM Remote Technical Support ServicePac® or IBM VMware Support Line agreement through the IBM Support Line or the IBM equivalent. You must have a current Software Maintenance Agreement in order to receive support for the IBM software products in this environment. Individual IBM software can announce a level of client support beyond that described. If applicable, information about the added support will be included in the specific product announcement letter.

4.8.2 Pre-install information

Before you begin installing ESX Server, do the following:

- ▶ You should update all firmware before starting. Download them from the above driver matrix Web page. This includes:
 - System BIOS
 - Diagnostics
 - BMC firmware
 - RSA II firmware (if installed)
 - ServeRAID firmware (if installed)

Tip: The URLs associated with these support documents are of the following form, where xxxxx is the number listed in the table:

<http://www.pc.ibm.com/support?page=MIGR-xxxxx>

In the PDF version of this redbook, the links in Table 4-14 are clickable.

Table 4-14 xSeries support documents for VMware ESX Server

Server	Driver matrix	Installation instructions
xSeries 460	MIGR-59923	MIGR-60546
xSeries 366	MIGR-59142	
xSeries 260	MIGR-61523	
All servers	MIGR-4JTS2T	MIGR-4QLNTQ

- ▶ If you are using a ServeRAID-8i RAID controller, use the supplied CD to configure an array or download the latest configuration CD from the above driver matrix page.
- ▶ In BIOS, ensure the following parameters are set:
 - In CPU Options, ensure that the Clustering Technology parameter is set to **Logical Mode**, as shown in Figure 3-1 on page 92.
 - Select **Advanced Settings** → **RSA II Settings** and ensure that the OS USB Selection setting is set to **Linux**, as shown in Figure 5-3 on page 167.

ESX Server 2.5.1 includes the necessary drivers for both SAS and ServeRAID-8i, so no additional drivers are typically needed to install the operating system.

4.8.3 Installing ESX Server

Instructions on how to install ESX Server 2.5.1 are available at:

<http://www.pc.ibm.com/support?page=MIGR-60546>

Tip: Be sure to check the Issues section (5.0) in this installation guide. At the time of writing, there was additional information there for multi-node x460 installations.

The upgrade patches and the latest releases of drivers can be found on the VMware Web site:

<http://www.vmware.com/download/>



Management

Like the rest of the xSeries family of servers, the X3 Architecture servers have a number of features that aid in the systems management of the servers. Topics covered in this chapter are:

- ▶ 5.1, “IBM Director” on page 162
- ▶ 5.2, “Remote Supervisor Adapter II SlimLine” on page 164
- ▶ 5.3, “Baseboard Management Controller” on page 177
- ▶ 5.4, “Integrating the service processors with IBM Director” on page 187
- ▶ 5.5, “OSA SMBridge utility” on page 194
- ▶ 5.6, “Predictive Failure Analysis (PFA)” on page 220
- ▶ 5.7, “IBM Dynamic System Analysis” on page 220
- ▶ 5.8, “Partition management” on page 223

5.1 IBM Director

IBM Director is a client/server workgroup manager. With IBM Director, you 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.

All X3 Architecture servers require IBM Director Version 4.22 or later. The latest version of IBM Director can be downloaded from:

<http://www.pc.ibm.com/support?page=SERV-DIRECT>

IBM Director is designed to manage an industry-standard server environment and supports a variety of operating systems, including Windows, Linux, and NetWare. The latest version, IBM Director 5.10, supports the full range of IBM @server platforms, including Linux on zSeries, iSeries, pSeries, xSeries, and BladeCenter. AIX on pSeries systems and i5/OS® on iSeries is also supported.

IBM Director supports many industry standards, such as DMI, CIM, WMI, SNMP, TCP/IP, IPX™, SNA, NetBIOS, SLIP, XML, and HTTP, among others.

The basic management functions included in IBM Director provide the ability to get servers up and running quickly and smoothly, and to simplify the ongoing hardware management. Additional powerful tools are available as extensions to IBM Director to further strengthen its management capabilities.

For more information about IBM Director, see *Implementing Systems Management Solutions Using IBM Director*, SG24-6188, found at:

<http://www.redbooks.ibm.com/abstracts/sg246188.html>

Other IBM Director plug-ins such as Virtual Machine Manager and Scalable Systems Manager can be downloaded at:

http://www.ibm.com/servers/eserver/xseries/systems_management/xseries_sm/dwn1.html

Integration with enterprise managers

Many clients already have an enterprise management system to manage all of their computer systems using centralized alerting, problem determination, and inventory. Because most of these enterprise solutions do not provide the hardware-level information that is available from systems such as the xSeries family, it is important to have tools such as IBM Director to enable integration with these large enterprise solutions.

IBM Director offers a graphical user interface for easy local and remote access, control, and smooth integration into higher levels of workgroup or enterprise management tools, including:

- ▶ Tivoli NetView®
- ▶ Tivoli Enterprise™
- ▶ HP OpenView
- ▶ Microsoft System Management Server (SMS)
- ▶ Computer Associates Unicenter
- ▶ NetIQ
- ▶ BMC Patrol (Morse Systems)

The NetIQ and BMC Patrol integration modules are available from those companies.

By letting IT administrators view the hardware configuration of remote systems in detail and monitor the usage and performance of critical components such as processors, hard disk drives, power supplies, cooling fans, voltage regulator modules (VRMs), and memory, IBM Director can help you manage your server more efficiently and can help you control many of the hidden costs of operation.

For further information, see the IBM Redbook *Integrating IBM Director with Enterprise Management Solutions*, SG24-5388, available at:

<http://www.redbooks.ibm.com/abstracts/sg245388.html>

The latest IBM Director upward integration modules are available at:

<http://www.pc.ibm.com/support?page=SERV-DIRECT>

Restriction: The IBM Director extension Scalable Systems Manger available for xSeries 445 and 455 does *not* work with xSeries 460 and MXE 460.

5.2 Remote Supervisor Adapter II SlimLine

The Remote Supervisor Adaptor II SlimLine (RSA II), shown in Figure 5-1, is a systems management card that ships with every x460 and MXE 460 and is available as an option for the x260 and x366. This section describes how to utilize the RSA II and its features for beneficial systems management.




Figure 5-1 Remote Supervisor Adapter II SlimLine

The most useful functions and features of the RSA II are:

- ▶ Automatic notification and alerts

The RSA II automatically sends different types of alerts and notifications to another server, like IBM Director, SNMP destination, or as e-mail directly to a user by using SMTP.

- ▶ Continuous health monitoring and control

The RSA II monitors all important system parameters like temperature, voltage, and so on, continuously. If a fan fails, for example, the RSA II forces the remaining fans to increase speed to compensate for the failing fan.

- ▶ Web interface

The RSA II is managed via its built-in Web interface. From it, you can control the server (power, remote media, and KVM), check sensors (fans, temperature, and so on) and set monitors.

- ▶ Event log

You can get access to the event logs of the server and the power-on-self-test (POST) log and export them while the server is up and running.

- ▶ Remote control

The RSA II card offers full remote control, including mouse, from power-up and Setup/Diagnostics panels, all the way through to the operating system running as normal. In fact, combined with the remote media function, you boot the server and remotely install an operating system from a remote CD-ROM using this feature.

- ▶ Operating system failure screen capture

When the operating system hangs, for example, with a blue screen, you can do a screen capture for support purposes. Additionally, the RSA II stores the last failure screen in memory so you can refer to it later.

- ▶ Remote media

As a part of the remote control feature, the remote media capability lets you use diskette drives, diskette images, optical drives (such as DVD or CD-ROM), or optical drive images of the system where the Web interface of RSA II is running on the remote PC, and make them appear to be local drives on the server.

Note: At the time of writing, support for optical images (ISO files) was being rolled out to all servers that support RSA II for Windows only. Linux support will come at a later date.

- ▶ Remote power control

The RSA II supports remote power control to power on, power off, or restart the server with or without operating system shutdown over LAN or even WAN connection.

The RSA II SlimLine family does not occupy a PCI slot. It is a small circuit board that plugs into a dedicated socket on the I/O planar.

5.2.1 Installing the RSA II SlimLine

The RSA II SlimLine is installed by default in the x460 and MXE-460. However, if you purchase the card for an x366 or x260 you will need to install it. Follow the instructions in the *Installation Guide* that ships with the server. No additional cabling is required.

5.2.2 Connectivity

Once the RSA II SlimLine is installed, two ports on the rear of the server are enabled. The location of these ports on the x366 is shown in Figure 5-2:

- ▶ Service processor Ethernet port is for TCP/IP-based communication.
- ▶ Service processor serial port is for serial/modem communication.

Tip: If the RSA II SlimLine is not installed, these two ports are non-functional.




Figure 5-2 The RSA II SL ports (x366 shown)

5.2.3 Network settings

After installing the adapter in your server, you have to configure the network settings to connect to the RSA II using the Web interface or telnet.

By default, the RSA II adapter is configured to look for a DHCP server to obtain an IP address, and if none is available, to use the IP address 192.168.70.125.

To view the current DHCP-assigned address, enter Setup (press F1 at boot time when prompted) and select **Advanced Settings** → **RSA II Settings**. The assigned address will be listed on the second line, as shown in Figure 5-3 on page 167.

If possible, we recommend you change this to always use a static IP address or at least set up a reservation in your DHCP server. This is changed in Setup by selecting **Advanced Settings** → **RSA II Settings** (see Figure 5-3 on page 167).

To configure the network settings, do the following:

1. Boot your server and press F1 to go to the BIOS settings.
2. Select **Advanced Setup** → **RSA II Settings**.

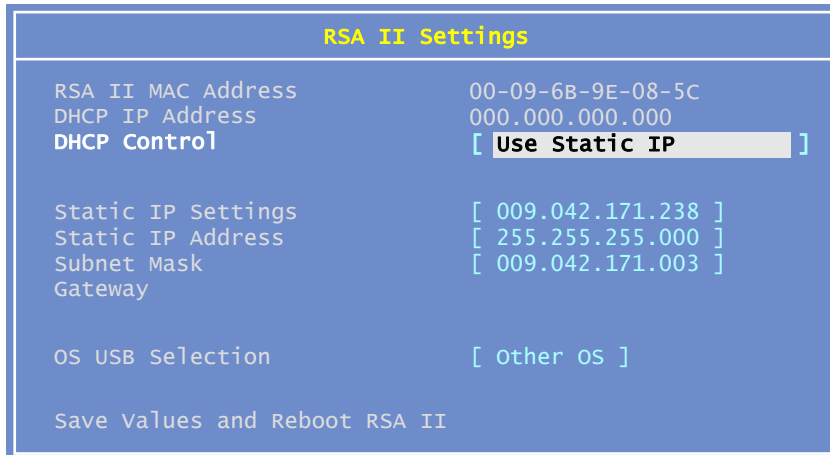


Figure 5-3 RSA II Settings in server BIOS

3. Change the DHCP control to **Use Static IP** by using the right and left arrow keys. We recommend that you use a static IP address for the RSA II so that you can still get access even if DHCP problems occur.
4. Fill in the IP address you want to assign to the RSA II, the network's subnet mask, and the standard gateway. Contact your network administrator for details.
5. Select **Other OS** for Windows operating system or **Linux OS** as OS USB Selection. Use the right and left arrow keys for selection.

The purpose of this selection is to prevent a known problem with Linux and its generic human interface device (HID) driver. Linux cannot establish USB communication with the RSA II using the generic HID (which Windows uses). By selecting **Linux OS** here, it makes the RSA II appear as an OEM HID instead of generic HID, which then functions properly.

Now select **Save the Values and Reboot RSA II**, with the arrow keys and press Enter. Exit the utility.

Tip: To check the network connection of RSA II use the PING command from another system connected to the network.

5.2.4 Web interface

The RSA II can be controlled from a Web browser. As described in the previous section, the adapter is configured by default to look for a DHCP server to obtain an IP address, and if none is available, to use the IP address 192.168.70.125. You can view the assigned address from the BIOS in the **Advanced Settings** menu. Once you have the address, enter it into a Web browser. The default user ID and password are USERID and PASSWORD (where 0 is the number zero).

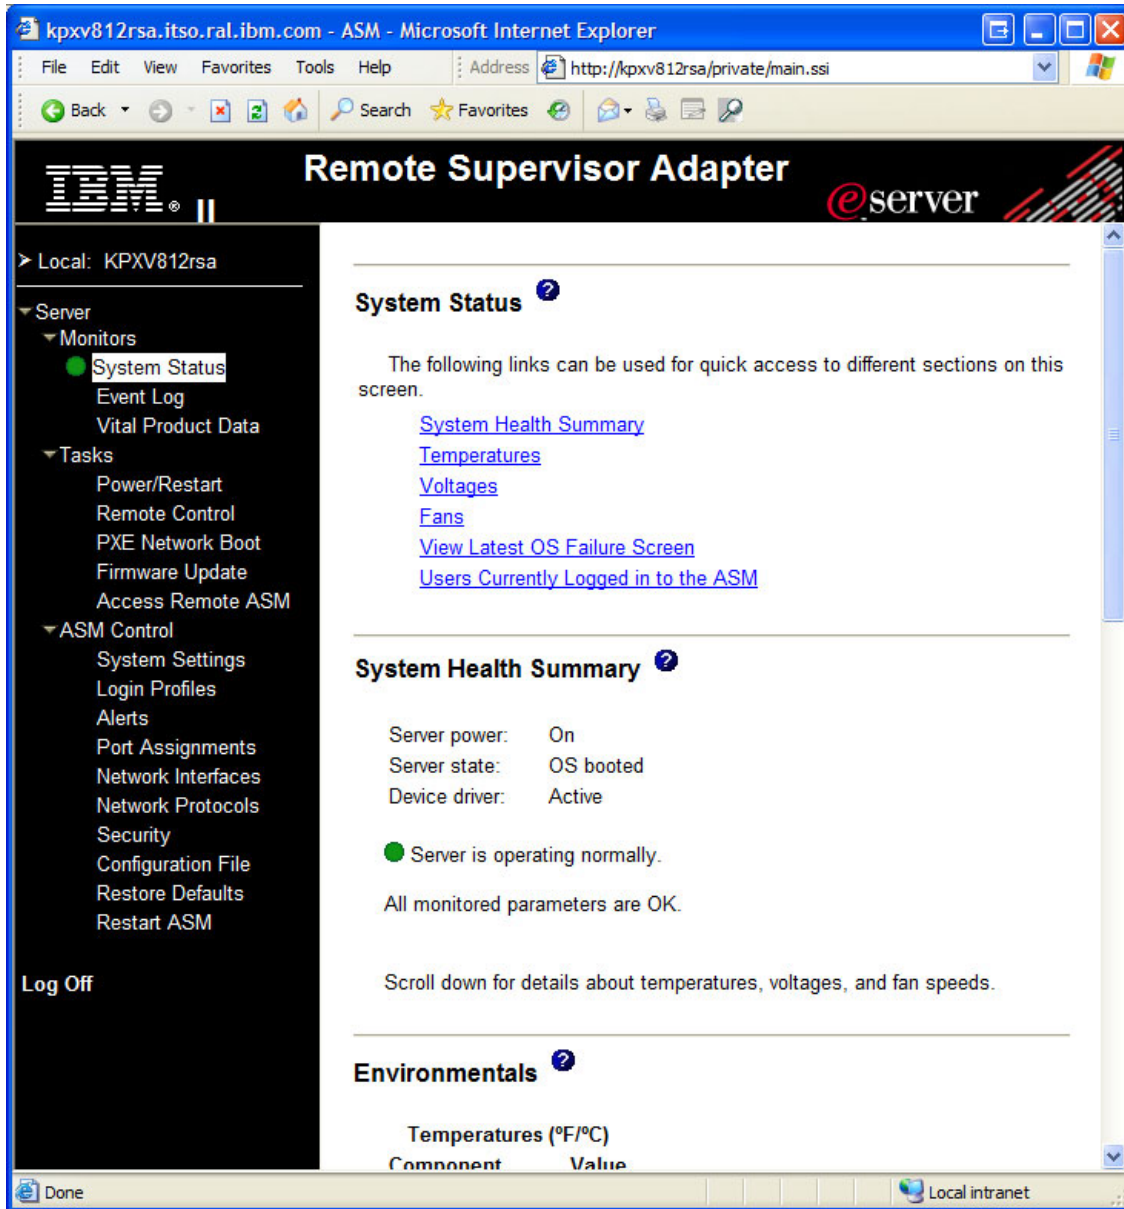


Figure 5-4 RSA II Web interface

5.2.5 Updating firmware

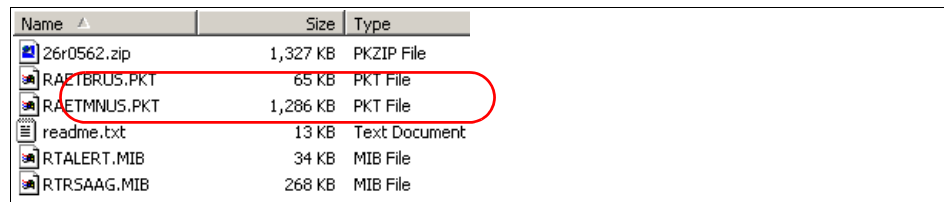
After the card is installed in the server, you will need to flash it with firmware that is customized for your model, the x366 or x260. Download the firmware from the *Remote Supervisor Adapter II Family — Firmware and Drivers Quick Reference*, TIPS0534. In particular, download the Packet files for your server:

<http://www.redbooks.ibm.com/abstracts/tips0534.html>

Tip: IBM provides the packet (PKT) files so you can update the firmware using the Web interface as we are describing here. In addition to that, there are also Windows and Linux tools that let you update the firmware from within the operating system running on the server. All three are listed in the above URL.

The steps to update the firmware using the PKT files and the Web interface are as follows:

1. Click the Packet files link for your server from the URL above.
2. Download the EXE file and save it to a local directory.
3. Run the EXE to extract the files. After extracting the file, take a few minutes to read the readme.txt. The following files should be in your directory. Note that there are two PKT files (Figure 5-5); you will need to perform the firmware update procedure twice, once for each file.



Name	Size	Type
26r0562.zip	1,327 KB	PKZIP File
RAETBRUS.PKT	65 KB	PKT File
RAETMNIUS.PKT	1,286 KB	PKT File
readme.txt	13 KB	Text Document
RTALERT.MIB	34 KB	MIB File
RTRSAAG.MIB	268 KB	MIB File

Figure 5-5 Files of RSA II firmware update package

Tip: The firmware package suitable for update via a Web browser is delivered as a ZIP file, and the ZIP file contains only two PKT files.

4. Connect to the RSA II using a Web browser by simply entering the IP address in the address bar.
5. Log on to the RSA II with the default user USERID and PASSWORD (with a zero, not the letter O) as the password (unless you have changed it). For security reasons, you should change the standard password before you put your server into production.

6. If this is a new adapter (x260 or x366) or new server (x460 or MXE-460), you will be prompted to update the firmware (Figure 5-6). If not, select **Tasks** → **Firmware Update**.

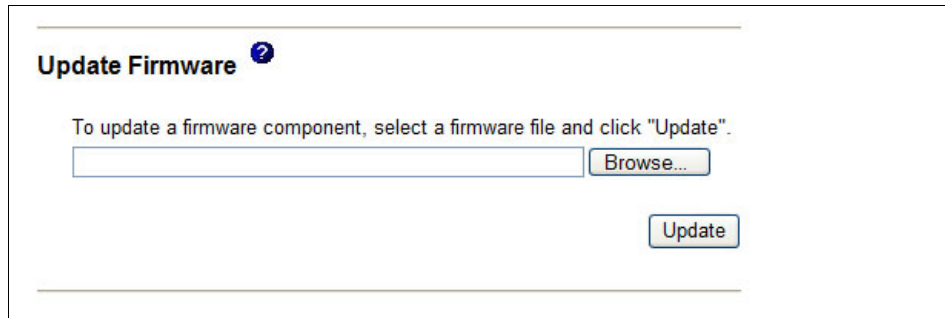


Figure 5-6 RSA II firmware update

7. Click **Browse** to select the first of two files for firmware update.
You should select the files in the correct order for updating the firmware. First, select RAETBRUS.PKT (RSA Boot ROM) and then RAETMNUS.PKT (RSA Main Application). Restart RSA only after applying both files.
8. To update, click **Update**. The file is now transferred to the RSA II.
9. Click **Continue** to flash the RSA II.
10. When prompted, *do not* restart the RSA adapter. You will do this after loading the second firmware PKT file.
11. Repeat steps 6-9 for the second PKT file.
12. Restart the adapter by selecting **ASM Control** → **ASM Restart**.

5.2.6 Installing the device driver

Once you install your operating system, you will need to also install the driver for the RSA II SlimLine adapter.

Tip: If you have an RSA II, you only need to install the RSA II driver. You do not need to also install the BMC drivers (although doing so will not have any negative impact).

Download the driver via the appropriate link from the *Remote Supervisor Adapter II Family — Firmware and Drivers Quick Reference*, TIPS0534, found at:

<http://www.redbooks.ibm.com/abstracts/tips0534.html>

Windows service installation

The installation of the RSA II server software package is unlike the driver installations of older systems management adapters. It is done by executing the downloaded executable file.

Note: If you have not already done so, change the setting OS USB Selection to **Other OS** in the system BIOS, as shown in Figure 5-3 on page 167.

The installation is as follows:

1. Execute the downloaded EXE file on the server with the RSA II.
2. Optionally, click **Change** to specify an alternate temporary folder for the installation files.
3. The installation process starts automatically after the files are copied.
4. Follow the instructions.
5. When the installation finishes, you can delete the files in the temporary folder.

To determine if the installation was successful, check the services for the IBM Remote Supervisor Adapter II by selecting **Start** → **All Programs** → **Administrative Tools** → **Services**. Scroll down to the service IBM Remote Supervisor Adapter II and verify that the status is started (see Figure 5-7).

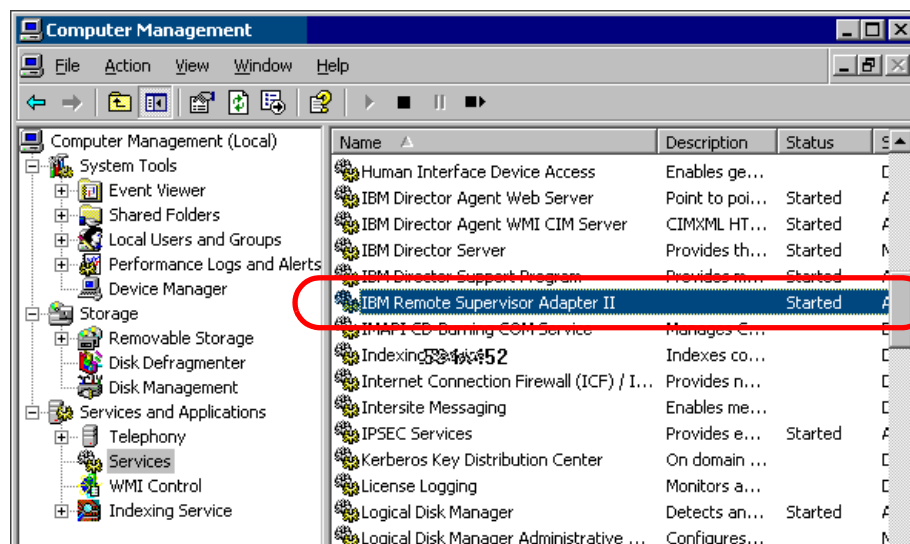


Figure 5-7 RSA II service in Windows Server 2003

Linux daemon installation

To install the Linux daemon for the RSA II, first download it from the *Remote Supervisor Adapter II Family — Firmware and Drivers Quick Reference*, TIPS0534, found at:

<http://www.redbooks.ibm.com/abstracts/tips0534.html>

Note: At the time of writing, drivers were not available for some distributions.

1. Select the correct rpm package for your Linux distribution (Red Hat or SUSE LINUX).
2. Review the appropriate readme file of the rpm package for prerequisites and installation steps.

Note: If you have not already done so, change the setting OS USB Selection to **Linux OS** in the system BIOS, as shown in Figure 5-3 on page 167.

3. Copy the downloaded file to a folder of the Linux server, for example, /tmp/inst.
4. Install the daemon (for example, SUSE, where xx is the version) by running:

```
rpm -ivh ibmusbasm-1.xx.i386.rpm
```

Now you can check to see if the daemon is running. Use the **ps** command, as shown in Example 5-1.

Example 5-1 Command to verify the RSA daemon is running

```
linux:~ # ps -ef | grep ibmasm
root      11056      1  0 10:47 pts/1    00:00:00 /sbin/ibmasm
root      11060 11056   0 10:47 pts/1    00:00:00 /sbin/ibmasm
root      11062 10996   0 10:48 pts/1    00:00:00 grep ibmasm
linux:~ #
```

If /sbin/ibmasm appears in the list, the daemon is running. The ibmusbasm daemon is started automatically during the boot process of the operating system.

To start the daemon manually, use the command **ibmspup**. To stop the daemon, enter **ibmspdwn**.

5.2.7 Remote console and media

To manage servers from a remote location, you often need more than just keyboard-video-mouse (KVM) redirection. For example, for the remote installation of an operating system or patches, you may need remote media to connect a CD-ROM or diskette to the server.

The RSA II SlimLine offers the ability to make available a local diskette, CD-ROM, or image to a remote server and have that server treat it as a local USB-attached device.

Tip: It is possible to mount more than one remote drive concurrently. For example, you could mount a CD-ROM and a diskette or diskette image.

Using remote media requires USB support from the operating system while the OS is up and running or during the installation of the OS. Remote media works with the following operating systems:

- ▶ Windows Server 2003
- ▶ Windows 2000 Server with Service Pack 4 or later
- ▶ Red Hat Enterprise Linux AS 3, but not for OS installation
- ▶ SUSE LINUX Enterprise Server 8, but not for OS installation

A Java runtime is required, which can be installed by going to:

<http://www.java.com/en/download/manual.jsp>

In the remote control window, there is a set of buttons that simulate specific keystrokes and the video speed selector, as shown in Figure 5-8. The slider is used to limit the bandwidth that is devoted to the remote console display on your computer.

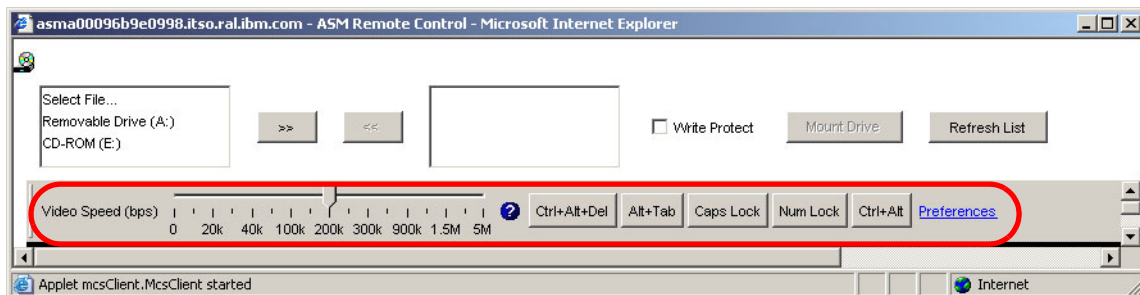


Figure 5-8 RSA II: Remote control buttons

Reducing the video speed can improve the rate at which the remote console display is refreshed by limiting the video data that must be displayed. You can

reduce, or even stop, video data to allow more bandwidth for remote disk, if desired. Move the slider left or right until you find the bandwidth that achieves the best results.

Each of the buttons represents a key or a combination of keys. If you press a button, the corresponding key stroke sequence will be send to the server. If you require additional buttons, click **Preferences**, where you can modify or create new key buttons.

The button bar can be detached by clicking anywhere in the grey background and dragging. Drop the button bar to create a separate window.



Figure 5-9 Detached button bar

The Preferences link also lets you specify your keyboard and enable mouse synchronization (that is, ensure the mouse pointer on the remote system precisely follows the local mouse pointer). The following keyboard types are supported:

- ▶ US 104-key keyboard
- ▶ Belgian 105-key keyboard
- ▶ French 105-key keyboard
- ▶ German 105-key keyboard
- ▶ Italian 105-key keyboard
- ▶ Japanese 109-key keyboard
- ▶ Spanish 105-key keyboard
- ▶ UK 105-key keyboard

For more information about the RSA II SlimLine, see Chapter 3, “Remote Supervisor Adapter II” in *IBM @server xSeries and BladeCenter Server Management*, SG24-6495.

5.2.8 TCP/UDP ports used by the RSA II SlimLine

The RSA II is using several TCP/UDP ports for communication. If the communication with the RSA II passes through firewalls, it is important to know which ports you have to enable on the firewalls to communicate with the RSA. Table 5-1 shows the default ports. Remember when you change the ports in the RSA you have to change them in the firewalls too.

Table 5-1 User configurable TCP/IP ports used by the RSA II

Port name	Port number	Description
http	80 (default)	Web server HTTP connection (TCP)
https	443 (default)	SSL connection (TCP)
telnet	23 (default)	Telnet command-line interface connection (TCP)
SSH	22 (default)	Secure Shell (SSH) command-line interface (TCP)
SNMP Agent	161 (default)	SNMP get/set commands (UDP)
SNMP Traps	162 (default)	SNMP traps (UDP)

Some other ports are fixed and cannot be changed, as shown in Table 5-2.

Table 5-2 Fixed TCP/IP ports used by the RSA II

Port number	Description
427	SLP connection (UDP)
1044	Remote disk function (TCP)
1045	Persistent remote disk (disk on card) (TCP)
2000	Remote Console video redirect (TCP)
6090	IBM Director commands (TCP)
7070-7074	Partition management (TCP)

5.2.9 MIB files

The RSA II supports SNMP from many management tools, including IBM Director. If you require MIB files, these can be found on the RSA II firmware update for your server, in the ZIP file that also includes the PKT files. Links to the appropriate firmware files are listed in technote *Remote Supervisor Adapter II Family — Firmware and Drivers Quick Reference*, TIPS0534, available from:

<http://www.redbooks.ibm.com/abstracts/tips0534.html>

5.3 Baseboard Management Controller

The Baseboard Management Controller (BMC) is standard on all three X3 Architecture servers, and provides basic management features and environmental monitoring. When errors or warnings occur, it controls the light path diagnostics panel to aid in problem determination. It also provides limited text-based remote control using Serial-over-LAN.

The BMC is based on an Hitachi 2166 chip and implements Version 1.5 of the IPMI specification. The spec document is available from:

ftp://download.intel.com/design/servers/ipmi/IPMIv1_5rev1_1-012904markup.pdf

This integrated BMC has the following functions:

- ▶ Monitoring of system voltages
- ▶ Battery voltage monitor
- ▶ System temperature monitors
- ▶ Fan speed control
- ▶ Fan tachometer monitor
- ▶ Power Good signal monitor
- ▶ System ID and planar version detection
- ▶ System power control
- ▶ System reset control
- ▶ NMI detection
- ▶ SMI detection and generation
- ▶ Serial Port text redirection using SOL
- ▶ Remind button detection
- ▶ System LEDs control (power, HDD activity, alert, and so on)
- ▶ Control of Lightpath LED

5.3.1 Connectivity

The BMC communicates via port 1 of the integrated Ethernet controller of each of the X3 Architecture servers. To communicate with the BMC, you would attach a standard Ethernet cable.

Figure 5-10 on page 178 shows the location of the port on the x366.

By default, the BMC is configured with an IP address of 10.1.1.97 (some servers may have a default BMC address of 0.0.0.0). You will need to set a valid static IP address using Setup in the BIOS or IBM Director. DHCP is currently not supported.




Figure 5-10 The BMC is accessed using the system Gigabit Ethernet port 1 (x366 shown)

Tip: Even if an RSA II SlimLine is installed in the server, you still access the BMC using system Ethernet port 1.

The servers have an RJ-45 port labeled “Systems Management”. This port is the Remote Supervisor Adapter II SlimLine Ethernet connector, and for the x260 and x366, is only active when the RSA II SL is installed.

In addition, there is a port labeled RS-485 and this port is only for use to connect the server to an Integrated xSeries Adapter (IXA) if one is installed in the server. They are *not* used to form an ASM interconnect network; the interconnect network is not supported.

5.3.2 Updating the BMC firmware

It is a best practice to ensure that the BMC firmware is at the latest level to avoid any issues.

Note: Updating the firmware of the BMC does not change any user settings.

The steps are as follows:

1. Download the firmware appropriate for your server. Refer to the IBM Technote *IBM @server xSeries BMC — Firmware and Drivers Quick Reference*, TIPS0532, and click the link for the firmware for the appropriate server:

<http://www.redbooks.ibm.com/abstracts/tips0532.html>

Alternatively, you can navigate to the appropriate download page from the software matrix:

<http://www.pc.ibm.com/support?page=MIGR-4JTS2T>

2. The BMC firmware update is available in bootable diskette form (an EXE file that you run to extract the files to diskette) or as a ISO image file for CD creation.
3. Insert the media into the server and boot the server. You may need to use POST/BIOS setup to configure the correct boot devices.
4. Once your system starts, a Ramdrive will be created, the appropriate files will be copied to this Ramdrive, and the BMC firmware update will execute automatically.

5.3.3 Configuring the BMC in BIOS

Within the system BIOS, you are able to configure the following settings:

- ▶ IP address
- ▶ Subnet mask
- ▶ Default gateway

The default IP address is 10.1.1.97 (some servers may have a default BMC address of 0.0.0.0).

Tip: If you have an RSA II also installed in your server, then you do not need to enable the BMC network interface by supplying a static IP address. You can perform all management functions from the RSA II Web interface or IBM Director.

If you want to modify the default user ID and password, either use IBM Director (the preferred method) or the `bmc_cfg` command-line tool supplied on the firmware diskette.

The `bmc_cfg` command-line tool is described in detail in Chapter 2, “Baseboard Management Controller” in *IBM @server xSeries and BladeCenter Server Management*, SG24-6495. To get to a command prompt where you can run `bmc_cfg`, boot to the firmware diskette/CD-ROM and exit from the firmware update when prompted.

To configure the BMC using the BIOS, follow these steps:

1. Reboot the server if currently running. During post, press F1 to enter the Configuration and Setup utility.
2. Select **Advanced Setup**.

3. Select **Baseboard Management Controller (BMC) Settings**. The menu is as shown in Figure 5-11.

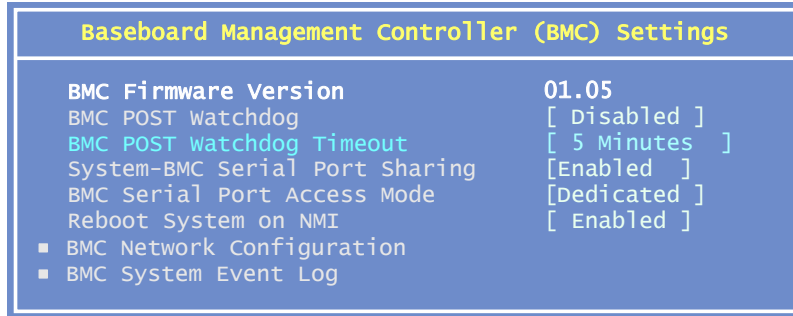


Figure 5-11 BMC Settings panel in BIOS

4. Select **BMC Network Configuration**.

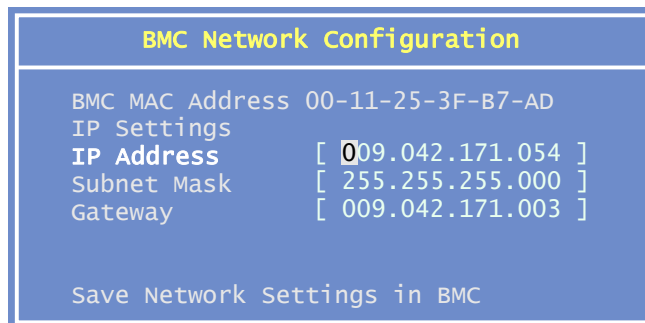


Figure 5-12 BMC network settings in BIOS

5. Enter the appropriate IP address, subnet mask, and gateway addresses, and then select **Save Network Settings in BMC**.

Further parameters will need to be adjusted if you plan to use console redirection and Serial over LAN. This can be used with a tool such as OSA SMBridge (see 5.5, “OSA SMBridge utility” on page 194).

5.3.4 Event Log

The BMC maintains a separate event log that contains, for example, entries like power events, events from environmental sensors, and chipset specific entries.

You can access the BMC System Event Log (SEL) via the menu shown in Figure 5-11 or using tools such as OSA SMBridge. This event log records all the hardware alerts for the server. The event log displays one event per screen. Use

the **Get Next Entry** and **Get Previous Entry** links to page through the events, as shown in Figure 5-13.

Tip: You may notice that some events have a date stamp of 2070. When power is removed from the server and later restored, the BMC clock resets to 1970 (which the BIOS displays as 2070). As soon POST completes and the operating system is passed control, the BMC clock is updated with the correct time from BIOS.

The incorrect time does not affect the timestamps of the events as received by IBM Director. The events as seen in the IBM Director Event Log have the correct time (the time the events were received) as long as the IBM Director Management Server has the correct time.

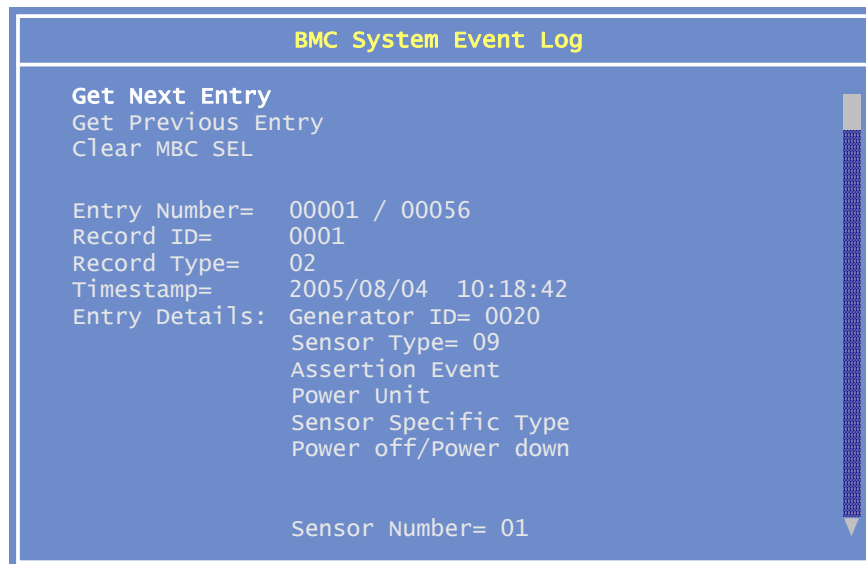


Figure 5-13 BMC Event Log

The system event log has room for 512 entries. You will be alerted if the log reaches 75 percent or 90 percent full. However, unlike the RSA II or BladeCenter management module, once the log is full, new entries are not saved and the LOG LED on the light path diagnostics panel and the Information LED on the front panel will be lit. You will need to clear the log in this instance using tools such as SMBridge (command-line interface or SOL interface) or IBM Director.

When an RSA II SlimLine is installed in the server, all events in the BMC System Event Log are also made available to the RSA II. The RSA II also maintains a separate log, and when viewed by accessing the RSA II, you will see both RSA II and BMC-based events. However, the reverse is not true. You cannot view the RSA II event log by viewing the event log in BIOS. The BIOS event log only shows BMC-based events.

5.3.5 Remote control

The BMC supports remote control using the OSA SMBridge utility and Serial over LAN. This provides a text-only console interface that lets you control BIOS screens and specific operating system consoles. Both Linux and Windows provide such text-only consoles (see 5.5, “OSA SMBridge utility” on page 194).

5.3.6 Installing the BMC device drivers

The device drivers are required to provide operating support and inband communication with IBM Director. This section describes how to install the IPMI device drivers on Windows and Linux platforms. The required device drivers are listed in Table 5-3.

Tip: If you have an RSA II SL installed in your server, you do not need to install the BMC drivers as described in this section. However, doing so will not cause problems.

Table 5-3 IPMI required device drivers

Device driver	Additional comments
IPMI device drivers	Required for in-band communication with IBM Director.
IPMI Library	This is the OSA BMC IPMI mapping layer. Includes the BMC Mapping Layer, which maps the dot.commands to IPMI commands. Required for in-band communication with IBM Director.
ASR Server Restart software	Required for ASR functionality.

The device drivers must be installed in a specific order or the installation will fail. The order is as follows:

1. IPMI device driver
2. IPMI library (mapping layer)
3. IPMI ASR service

To download the drivers appropriate for your server, refer to the IBM Technote *IBM @server xSeries BMC — Firmware and Drivers Quick Reference*, TIPS0532, and click the link for the firmware for the appropriate server:

<http://www.redbooks.ibm.com/abstracts/tips0532.html>

Alternatively, you can navigate to the appropriate download page from the software matrix:

<http://www.pc.ibm.com/support?page=MIGR-4JTS2T>

Installing the device drivers on Windows

This section describes how to install the drivers under Windows.

IPMI device driver

To install the OSA IPMI device driver, follow these steps:

1. Run Setup.exe. After the usual initial windows, you will be prompted to select a driver parameter, as in Figure 5-14.




Figure 5-14 Driver parameters

2. Click **No**.

Clicking **No** means that you want the driver to query SMBIOS record 38 for the necessary device driver parameters. Clicking **Yes** means that you will manually set these parameters, and this should only be used in unusual situations as directed by IBM support. If you manually set an invalid parameter, you may cause Windows to crash or reboot.

3. Click **Next** to begin the installation. When it completes, you will be prompted to reboot the server, but the installer will not do this automatically.

IPMI mapping layer (library) files

To install the IPMI mapping layer (library) files, do the following:

1. Ensure that the IPMI device driver is installed before installing this software.
2. Download the EXE from the above Web site and run it.
3. Follow the on-screen instructions.
4. Reboot the server if the installation procedure prompts you to do so.

IPMI ASR service

To install the ASR service, do the following:

1. Ensure that the IPMI device driver and IPMI library files are installed before installing this software.
2. Download the EXE from the above Web site and run it.
3. Follow the on-screen instructions.
4. Reboot the server if the installation procedure prompts you to do so.

Installing the device drivers on Linux

This section describes how to install the drivers under Linux.

IPMI device driver

To install the OSA IPMI device driver, launch a UNIX shell and enter the following command to build and install the driver module on your system:

```
rpm -i osa_ipmi-x.x.x-x.i386.rpm
```

If you upgrade your Linux kernel, you should uninstall and then recompile/re-install the OSA IPMI device driver. To rebuild the driver, change to directory `/usr/osa/osa_ipmi-x.x.x-x` and enter:

```
sh build_osadrv
```

Notes:

- ▶ You should install the necessary source code packages on your system. We recommend that you install them in `/usr/src`.
- ▶ Make sure you have gcc 3.2 available. You may also need to upgrade the `binutils` package.
- ▶ When installing on 32-bit SUSE LINUX 8.2, which has gcc Version 3.3 20030226 (pre-release) installed, `insmod` does not work without the `-f` option. You should manually add the `-f` option to the `insmod` command in the `/sbin/ipmi_load` script for the force loading. You may, however, receive a warning that the kernel is tainted. Other gcc versions may cause the same problem.

To uninstall the OSA IPMI device driver, enter one of the following:

```
rpm -e osa_ipmi-x.x.x-x
rpm -e osa_ipmi
```

See the README.TXT file available with the driver for more information.

IPMI mapping layer (library) files

The IBM mapping layer software is installed and removed via the Linux RPM package management tool. Ensure that you have first installed the IPMI driver.

If this is an upgrade to an existing software package, remove the old version first with the command:

```
rpm -e ibmsp6a
```

Depending upon your system's configuration, you may see messages about missing files. These may be ignored.

To install the IPMI mapping layer (library) files, issue the following commands:

EM64T and AMD64 note: On x86_64 kernels, this RPM will build a 64-bit shared object and a 32-bit compatibility shared object. Before installing the RPM on an x86_64 kernel, make sure the 32-bit compatibility development packages are installed.

```
rpmbuild --rebuild ibmsp6a-x.xx-y.src.rpm
```

Followed by:

```
cd /usr/src/package-dir/RPMS/architecture
rpm -ivh ibmsp6a-x.xx-y.architecture.rpm
```

Where:

- ▶ `package-dir` is the distribution-specific name of the RPM build directory (usually "redhat" or "packages").
- ▶ `architecture` is the architecture of the kernel in use (i386™, i586, or x86_64).

For example, to install the rpm on an x86_64 SUSE LINUX, the commands are:

```
rpmbuild --rebuild ibmsp6a-x.xx-y.src.rpm
cd /usr/src/packages/RPMS/x86_64
rpm -ivh ibmsp6a-x.xx-y.x86_64.rpm
```

IPMI ASR service

This section describes how to install the ASR (ibmipmiasr) RPM.

Before installing, make sure your server has both the IPMI device driver and the IBM Mapping Layer Software installed.

The system that the source rpm file is to be run on must have Linux development/build capability.

If this is an upgrade to an existing software package, remove the old version first with the command:

```
rpm -e ibmipmiasr
```

Depending upon your system's configuration, you may see messages about missing files; these may be ignored.

To install the source rpm, execute the following command:

```
rpm -ivh ibmipmiasr-x.xx-y.i386.rpm
```

Once the installation is complete, check the log file /var/log/message. A successful installation will write the following message to the log:

```
IBM IPMI ASR application loaded
```

To uninstall the binary rpm, execute the following command:

```
rpm -e ibmipmiasr
```

RPM will unload the ASR application, and remove all ibmipmiasr-related files from your system.

5.3.7 Ports used by the BMC

The BMC uses several TCP/UDP ports for communication, as shown in Table 5-4. If the communication with the BMC passes firewalls, it is important to know which ports you have to enable on the firewalls to communicate properly.

Table 5-4 TCP/IP ports used by the BMC

Port number	Description
623	IPMI communications to SMBridge and Director
664	IPMI communications (secondary)
161	SNMP get/set commands
162	SNMP traps and PET alerts to Director

5.4 Integrating the service processors with IBM Director

If you have the IBM Director Server installed in your environment, you will most likely want to manage the service processors (BMC and RSA II) with IBM Director. In this section, we describe how to configure the user ID, passwords, and alert-forwarding settings.

Tip: If you have an RSA II also installed in your server, then you do not need to enable the BMC network interface by supplying a static IP address. You can perform all management functions from the RSA II Web interface or IBM Director.

5.4.1 Adding the service processor to IBM Director

If you have the service processor device driver and the IBM Director Agent installed on the server, then IBM Director can automatically discover the service processor as a Physical Platform. Consequently, you do not need to manually add it as described in these steps.

To manually add the BMC or RSA to Director as a managed object, do the following:

Tip: If your server has both an RSA and a BMC installed, only the RSA will be added to Director. You cannot add both service processors in the same server to Director.

1. From the Director console, right-click the middle pane in a blank area.
2. Select **New** → **Physical Platform**.
3. Enter the appropriate details for the service processor you want to add and select **OK**, as shown in Figure 5-15.

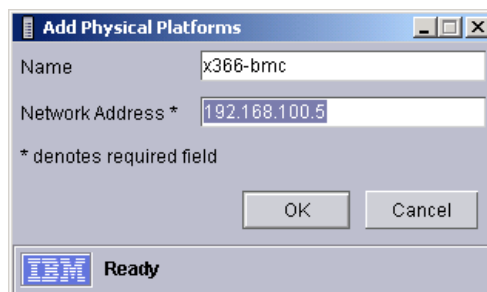


Figure 5-15 Add physical platforms window

4. Once the service processor is discovered, it will appear as a Physical Platform object on the Director console.
5. IBM Director attempts to access the BMC using the default USERID/PASSWORD combination.

If you have deleted or changed the default USERID/PASSWORD combination, then a small padlock icon will appear next to the device. Right-click the device and click **Request Access**, and enter a valid user ID and password.

5.4.2 Adding users

In Director 5.10, the recommended way to add users to the service processors is to use the Server Configuration Manager. This is a task new to Director 5.10. With it, you create a profile and then apply it to multiple systems at once. It does have some limitations, however:

- ▶ You can only add new users. You cannot replace existing users, including the default USERID/PASSWORD combination.
- ▶ Since the RSA II supports 12 users at most and the BMC supports four users at most (one of which is reserved), then this means you can only add 10 and two, respectively.
- ▶ You cannot delete users. This function, however, is planned to be added in a future release of IBM Director.

To add a user to one or more service processors, do the following:

1. Double-click **Server Configuration Manager**.
2. Click **Create a new profile**, as shown in Figure 5-16 on page 189. You can also edit an existing profile or import a profile from an XML file (which you probably based on an XML file you previously exported).

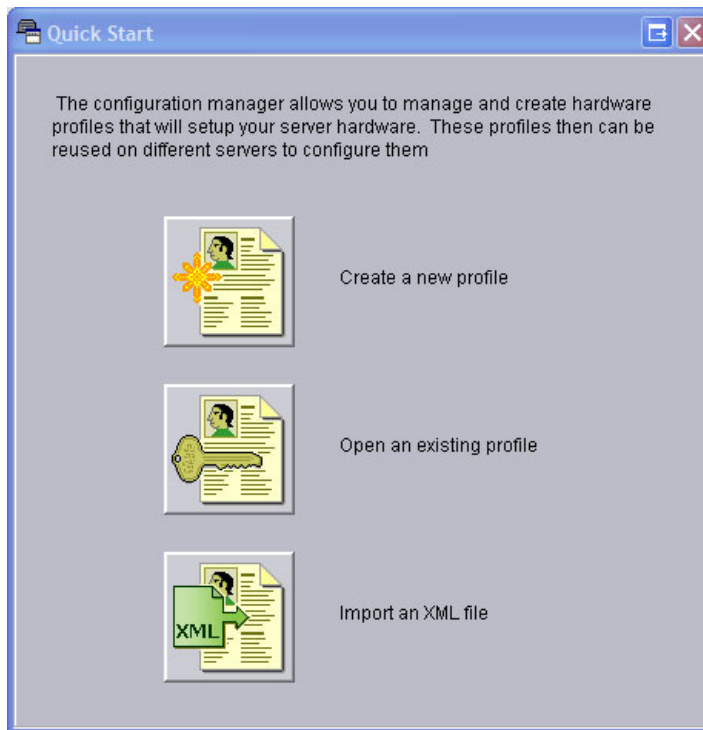


Figure 5-16 Server Configuration Manager: Quick Start menu

3. Enter a name for the profile (for example, **spadmin userid**) and click **OK**.
You can select any of three components, as shown in Figure 5-17:
 - IP Address Pool: Only applies to BladeCenter Management Modules
 - Service Processor Login Profile: To add a new service processor user ID
 - Service Processor Network Configuration: To configure DHCP and SNMP properties

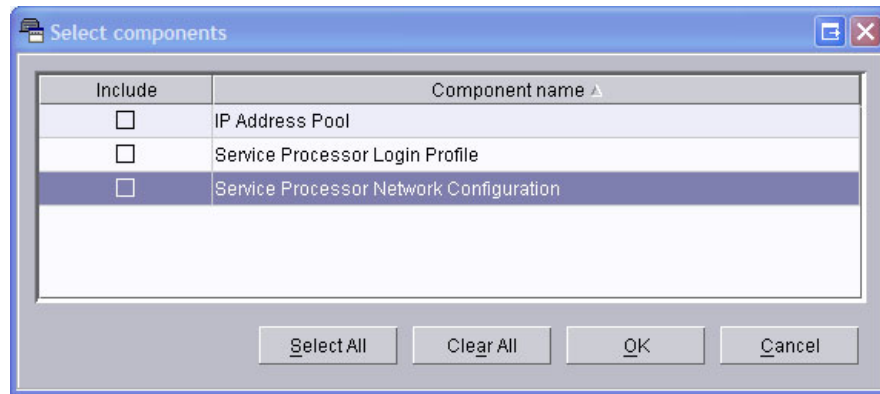


Figure 5-17 Select components window

We select only the Service Processor Login Profile. Click **OK** to continue.

4. Enter a suitable user name, password, and authority level, as shown in Figure 5-18 on page 191. If you wish to replace the default USERID/PASSWORD combination, check the check box **Replace default user ID and password**.

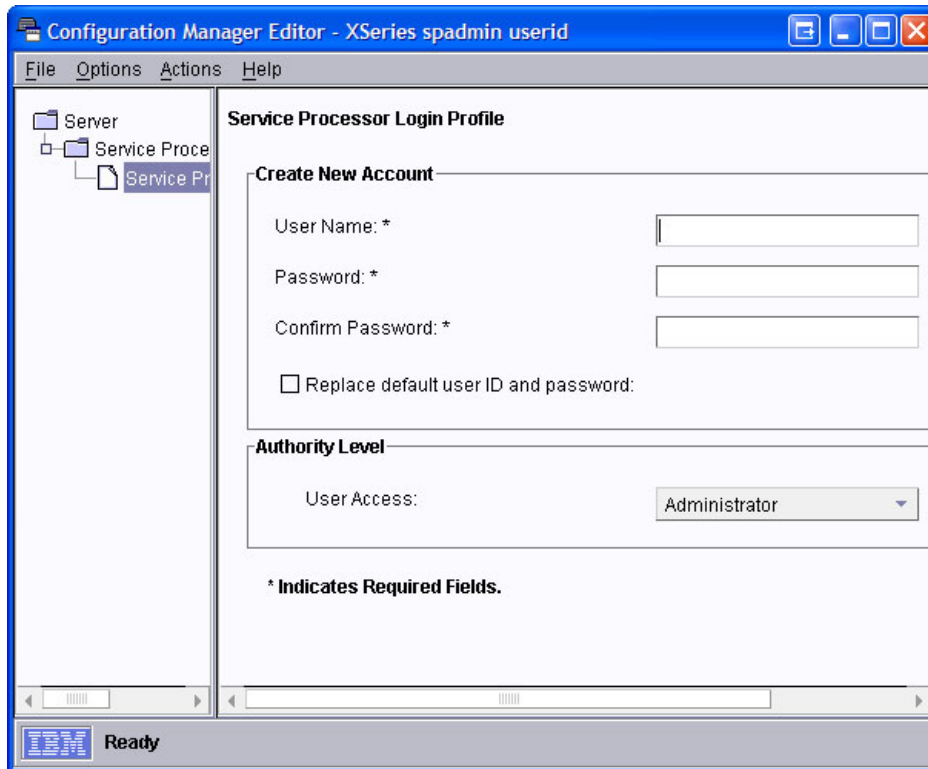


Figure 5-18 Server Configuration Manager: adding a login profile

5. Select **File** → **Save** to save the profile, and then close the editor.

6. A profile will now appear in the tasks pane in the Director console (see Figure 5-19). You can now drag and drop that profile to one or more service processors (Physical Platforms) or groups to either execute the task now or schedule it to run later.
7. You will see the status of each system as it is processed.

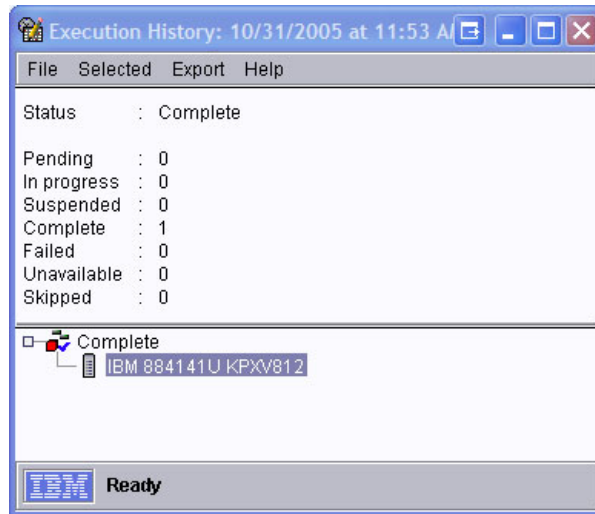


Figure 5-19 Results of applying the task against a service processor

8. Once complete, you can view the detailed results by double-clicking any of the systems listed in the bottom half of the window. To see the status of an individual system, double-click it in the lower pane of the Execution History window.

If successful you will see:

```
-VENDOR- sponfig: Creating a new account for user spadmin
-VENDOR- sponfig: Setting the user's password
-VENDOR- sponfig: Setting the user's authority level
```

Note that if you already have the maximum number of users already defined, then you will get an error:

```
-VENDOR- sponfig: User does not already exist and maximum number of
users has already been reached
```

Note: Server Configuration Manager adds users to the highest numbered available slot first. For example, with the RSA II, the first user it adds will be user ID 12 if that slot is available.

For the RSA II (with a limit of 12 users, including USERID), you will need to go to the Web interface to delete an existing users. For the BMC (with a limit of three users, including USERID), you will need to delete the user using the MPCLI (management processor command-line interface). You can also use the MPCLI with the RSA II as well.

Using the MPCLI, you can view the user IDs already defined using the **getdialinentry** command and you can delete existing user IDs using the **setdialinentry** command (see Figure 5-20 for an example).

Tip: To delete a user ID, set the user ID to a blank with the parameter `-id ""`

```
logged on ip=9.42.171.192>getdialinentry -index 11
Dialin Configuration:
  Entry Number: 11
  Id: blah
  Last login: Tue Nov 30 00:00:00 EST 1999
  Dialback enabled: false
  Dialback number: null
  Authority: supervisor
logged on ip=9.42.171.192>setdialinentry -index 11 -id ""
true
logged on ip=9.42.171.192>getdialinentry -index 11
Dialin Configuration:
  Entry Number: 11
  Id:
  Last login: Tue Nov 30 00:00:00 EST 1999
  Dialback enabled: false
  Dialback number: null
  Authority: supervisor
```

Figure 5-20 MPCLI commands to read and delete a service processor user ID

Configuring alert forwarding

To configure the service processors to send all alerts to IBM Director 5.10, use the following MPCLI commands (where, in our example, 9.42.171.229 is the address of our Director 5.10 server):

```
setalertentry -index 1 -type director.all
setalertentry -index 1 -ipaddress 9.42.171.229
```

You can use the following command to verify the setting

```
getalertentry -index 1
```

For the RSA II adapter, you can also use the Web interface to configure alert forwarding by setting it to type “Director Comprehensive”.

5.5 OSA SMBridge utility

The OSA System Management Bridge (SMBridge) is a utility that lets you perform certain remote management functions on servers with a BMC service processor, such as the X3 Architecture servers. It allows the administration of servers using IPMI1.5 protocol and the Serial Over LAN (SOL) protocol via either the server's Ethernet or serial interfaces. The primary function of SMBridge is to provide remote control of the text-mode console via Ethernet.

There are two ways to use the SMBridge utility: as a telnet server and as a direct command-line interface to the BMC. These are shown in Figure 5-21.

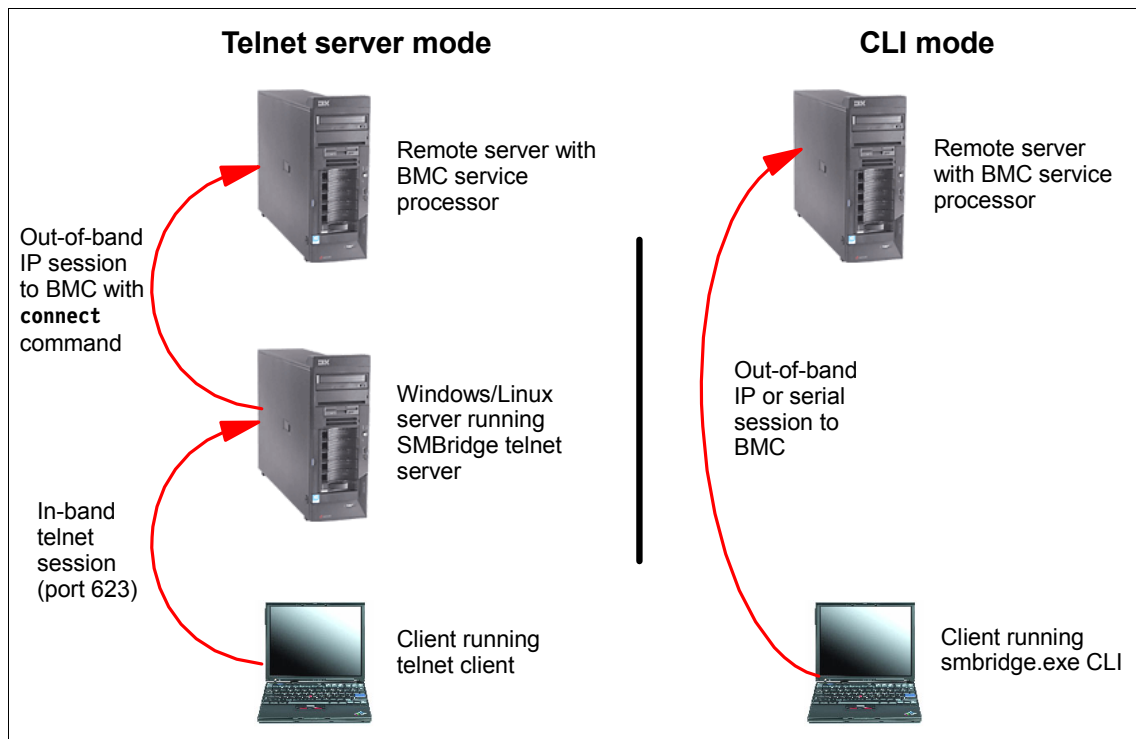


Figure 5-21 The two modes of the SMBridge utility

- Telnet server connection

Used as a telnet server, SMBridge is started as a background service or daemon on a system on your network (typically not the server with the BMC).

You connect to the telnet server, then from there, connect to the BMC via the server Ethernet port.

SMBridge uses the Serial Over LAN protocol to let the administrator remotely control text-mode tasks, such as POST messages, BIOS setup, and text-mode tasks with operating systems. Tasks you can perform are:

- Establish a text-mode console session with the remote server.
- Power on, power off (immediate and graceful), or reboot a server.
- Turn on/off the blinking system identifier.
- Display the current power status.
- Display the event log.

Any standard telnet client application, such as HyperTerminal on Microsoft Windows or telnet on Linux, can be used to access the server's features.

The SOL protocol coupled with the remote system's BIOS console redirection allows administrators to view and change the BIOS settings over LAN. Linux serial console and Microsoft's Emergency Messaging Service (EMS)/Special Administration Console (SAC) interfaces can also be accessed over LAN using SOL.

This is discussed further in 5.5.3, "Connecting via the telnet server" on page 201.

► Command-line interface

Used in this way, SMBridge lets an administrator perform the following tasks on a remote BMC service processor to:

- Power on, power off (immediate and graceful), or reboot a server.
- Turn on/off the blinking system identifier.
- Display the current power status.
- Display or clear the event log.

The CLI lets you do all but the remote console feature offered by the telnet server.

This is discussed further in 5.5.7, "Connecting via the command-line interface" on page 217.

SMBridge can be downloaded from:

<http://www.ibm.com/pc/support/site.wss/MIGR-57729.html>

The current version supports the following operating systems:

- ▶ Red Hat Linux 7.2
- ▶ Red Hat Linux 8.0
- ▶ Red Hat Linux 9.0
- ▶ Red Hat Enterprise Linux 3.0
- ▶ Microsoft Windows XP
- ▶ Microsoft Windows 2000 Professional
- ▶ Microsoft Windows 2000 Server
- ▶ Microsoft Windows Server 2003

The *OSA System Management Bridge User's Guide* is available at:

<http://www.ibm.com/pc/support/site.wss/MIGR-57816.html>

5.5.1 Configuring BIOS

Before SMBridge can be used to manage a remote server via SOL, the BMC and BIOS of the remote server must have the following settings configured:

Note: This procedure disables PXE boot on Gigabit port 1 on the server. If you plan to use PXE, you will need to connect Gigabit port 2 to your network and ensure that your remote install procedure is configured to use that port.

1. Enter the BIOS Setup by pressing F1 when prompted during boot.
2. If you have not done so already, configure the static IP address, subnet mask, and gateway of the BMC, as described in 5.3.3, “Configuring the BMC in BIOS” on page 179.
3. From the main menu, select **Devices and I/O Ports**. Set the following:
 - Set field Serial Port A to Auto-configure.
 - Set field Serial Port B to Auto-configure.
4. Select **Remote Console Redirection**. Set the following:
 - Remote Console Active to Active
 - Remote Console Text Emulation to VT100/VT220
 - Remote Console Keyboard Emulation to VT100/VT220
 - Remote Console Active After Boot to Enabled
 - Remote Console Flow Control to Hardware

The result is shown in Figure 5-22 on page 197.

```

*****
*                               *
*           Remote Console Redirection           *
*                               *
*****
* Remote Console Active           [ Enabled ] *
* Remote Console COM Port        [ COM 1 ] *
* Remote Console Baud Rate       [ 19200 ] *
* Remote Console Data Bits       [ 8 ] *
* Remote Console Parity          [ None ] *
* Remote Console Stop Bits       [ 1 ] *
* Remote Console Text Emulation  [ VT100/VT220 ] *
* Remote Console Keyboard Emulation [ VT100/VT220 ] *
* Remote Console Active After Boot [ Enabled ] *
* Remote Console Flow Control    [ Hardware ] *
*****

```

Figure 5-22 Remote Console Redirection settings to enable SOL

5. Press Esc twice to return to the main menu, then select **Start Options**. Set the following:
 - Planar Ethernet 1 PXE to Disabled
 - Planar Ethernet 2 PXE to Enabled
 - Planar Ethernet PXE/DHCP to Planar Ethernet 2
 - Run PXE only on Selected Planar NIC to Enabled

Note that you will most likely only have some of these options on your server.
6. Press Esc to return to the main menu, select **Advanced Options**, and then select **Baseboard Management Controller (BMC) Settings**. Set the following:
 - System-BMC Serial Port Sharing to Enabled
 - BMC Serial Port Access Mode to Dedicated
7. Save the BIOS settings and reboot the server.

5.5.2 Installation

This section describes how to install the SMBridge utility on both Windows and Linux platforms. The latest version of the utility is available from:

<http://www.ibm.com/pc/support/site.wss/MIGR-57729.html>

Microsoft Windows

The procedure both installs the CLI, and it installs and enables the telnet server.

Tip: This installation is normally run on the server you plan to have act as the telnet server, as shown in Figure 5-21 on page 194. If you plan to use the CLI, you do not actually need to install the tool, since the only files you need to run the CLI are smbridge.exe and smbridge.cfg.

Consequently, if you plan to use the CLI on other systems, you can either:

- ▶ Copy the files smbridge.exe and smbridge.cfg from the telnet server.
- ▶ Install SMBridge and then disable the service.

1. Run Setup, agree to the license, and specify an installation directory.
2. You will now be prompted to enter an IP address and TCP/IP port number, as shown in Figure 5-23.

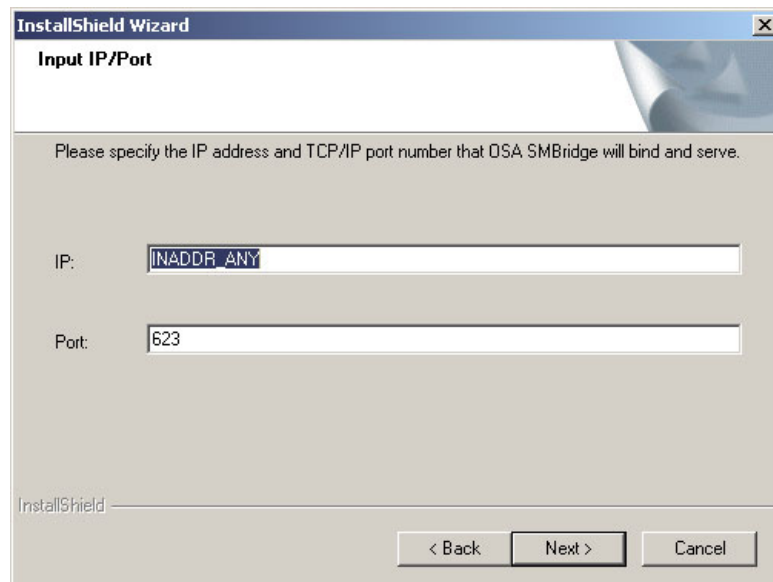


Figure 5-23 IP address and port number window

These values are as follows:

- IP specifies the server IP address that SMBridge will bind to.

Since a server may have multiple valid IP addresses, SMBridge allows you to restrict access to it via a single IP address. Specify INADDR_ANY as the IP address if any of the multiple IP addresses can be bound to SMBridge. Specify 127.0.0.1 or localhost as the IP address if SMBridge

should only accept local connections. Specify a specific IP address if only this IP address should be bound to SMBridge.

- Port specifies the server port number that SMBridge will listen on.

Note: These two values will be recorded in the `smbridge.cfg` file for automatic startup of SMBridge as a service daemon.

3. Next, you are asked to specify timeout values for telnet sessions (in minutes) and the power-off command (in seconds) (Figure 5-24).

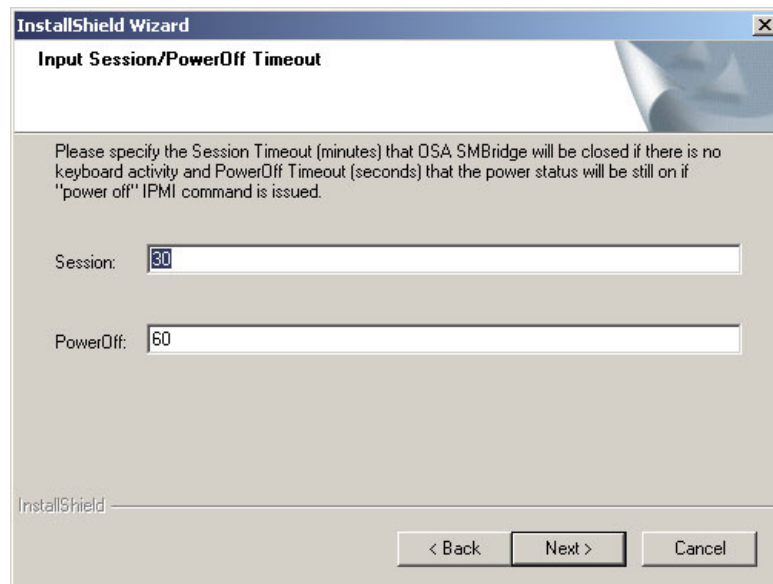


Figure 5-24 Session and power off timeout settings window

These values are as follows:

- Session specifies the number of minutes without any keyboard activity before an established telnet session is ended.
- PowerOff specifies the number of seconds to wait for an IPMI power off command (graceful or forced) to complete. If time has exceeded the timeout value and the power status is still on, an error code will be returned to indicate that the power off command may have failed.

4. Click **Next** to confirm your choices, and then begin the installation. Once the installation is complete, click **Finish** to end the installer.

The OSA SMBridge service is started automatically and is configured to start every time the server starts. You can change this by selecting **Control Panel** → **Administrative Tools** → **Services**.

Installation on a Linux platform

To install SMBridge on Linux, follow these steps:

1. Log in as root.
2. If the SMBridge RPM file is on CD, then insert the CD into the drive and enter the following commands to mount the drive and change to the root directory of the CD:

```
mount /mnt/cdrom  
cd /mnt/cdrom
```

3. Run the installation with the following command (substitute the file name of the rpm file you have if it is different):

```
rpm -i osasmbridge-1.0.3-1.i386.rpm
```

When the installation process has finished successfully, files are copied to the following directories:

- ▶ /etc/init.d/smbridge
- ▶ /etc/smbridge.cfg
- ▶ /usr/bin/smbridge
- ▶ /var/log/smbridge
- ▶ /var/log/smbridge/LICENSE

Additionally, the symbolic link /usr/sbin/smbridge is created.

The text file /etc/smbridge.cfg contains a number of SMBridge runtime parameters that you should review and modify if necessary.

To start the daemon, navigate to directory /etc/int.d and use the following command to start or stop the OSA SMBridge daemon service:

```
smbridge start  
smbridge stop
```

You can also start the daemon using the command:

```
smbridge -d -c config-file
```

Where config-file is the name of the file containing the telnet server configuration. By default, it is /etc/smbridge.cfg. See Appendix C, “Configure OSA SMBridge”, in the *OSA System Management Bridge User’s Guide* for more details about this file, which can be found at:

<http://www.pc.ibm.com/support?page=MIGR-57816>

5.5.3 Connecting via the telnet server

As shown in Figure 5-21 on page 194, you can use SMBridge as a telnet interface (or a “bridge”) to the BMC. You connect to a telnet server (where you installed SMBridge), and from there you connect to the BMC using a Serial Over LAN (SOL) connection.

Using a SOL connection, you can perform over a LAN connection all the tasks that you would normally only be able to do while connected directly to the server’s serial port:

- ▶ Change the BIOS settings.
- ▶ Linux serial console.
- ▶ Emergency Messaging Service (EMS) from Microsoft.
- ▶ Special Administration Console (SAC) from Microsoft.

Information about Microsoft EMS is available from:

http://www.microsoft.com/resources/documentation/WindowsServ/2003/standard/proddocs/en-us/EMS_topnode.asp

The SAC commands you can perform are described here:

http://www.microsoft.com/resources/documentation/WindowsServ/2003/standard/proddocs/en-us/EMS_SAC_commands.asp

With BIOS console redirection to serial port enabled on the remote server, applications that use the BIOS to read and write to the system console will have their I/O redirected to the serial port. With SOL, the BMC firmware then reads the data written to the serial port and transmits it to the SMBridge as LAN packets. SMBridge then forwards the data to the telnet client as TCP/IP packets.

One SMBridge session supports one SOL session with one BMC at a time.

Telnet clients

To access the SMBridge telnet server, connect using port 623 (or the port you selected during installation if you changed it from this default). For example:

```
telnet smbridge-server 623
```

Note: By default, the SMBridge telnet server listens on port 623.

Telnet clients that support VT100 terminal emulation can be used to access the BMC via SMBridge, including the following:

► The **telnet** command-line utility in Windows

Additional information:

- Our testing showed that on a Windows XP system, pressing F1 on the keyboard correctly sent us to the remote server (for example, to enter the BIOS Setup). If your telnet client does not work this way (such as the **telnet** command in Windows Server 2000), you can simulate the F1 and F2 keys in two ways. To simulate F1, either press:

- Esc, then Numeric+1 (that is, the 1 key on the numeric keypad)
- Esc, then Shift+O, then Shift+p

To simulate F2, either press:

- Esc, then Numeric+2 (that is, the 2 key on the numeric keypad)
- Esc, then Shift+O, then Shift+q

Tip: We recommend that you use this telnet client when running Windows. It is a standard implementation of telnet and also supports colors.

► HyperTerminal in Windows

Additional information:

- HyperTerminal supplied with Windows Server 2003, Enterprise Edition may not work correctly. You should upgrade to the Private version.
- HyperTerminal supplied with Windows 2000 Server displays random characters and loses some text. We recommend that you do not use this client.
- When creating a new connection, select **TCP/IP (WinSock)** in the Connect window using the drop-down menu. Enter the IP address of the telnet server and specify port 623 (or the port you specified when you installed SMBridge).
- Turn off automatic line wrapping by selecting **File** → **Properties** → **Setting** → **ASCII** → **Setup** and unchecking the **Wrap lines that exceed terminal width** check box.
- Configure the connection to emulate a VT100 terminal in the Properties window.

► The **telnet** command as part of csh or ksh in Linux.

The F1 and F2 keys may not work correctly, especially outside of an X-Windows environment. You may be able to reconfigure X-Windows to generate VT100 keystrokes. With KDE, for example, use Settings to reconfigure the keyboard.

Connecting

In this section, we describe the process by referring to servers in our lab. We have installed the SMBridge telnet server on the system at address 9.42.171.121, and the remote server has a BMC configured to use address 9.24.171.237. Our BMC uses the default USERID/PASSWORD authentication.

To connect to the BMC, do the following (using our example addresses):

1. Issue the following to connect to the telnet server on port 623:

```
telnet 9.42.171.121 623
```

2. You will be given the prompt shown in Figure 5-25.

```
Username:
```

Figure 5-25 Username prompt

3. Enter an administrator user ID and password that is valid for the telnet server (for example, Administrator or root).

Note: If the SMBridge telnet server is running on the local system, you will not see this prompt, as SMBridge uses the authority of the current user logged on.

4. You will then see the welcome message shown in Figure 5-26.

```
Username:Administrator
Password:

Administrator login successful.

OSA System Management Bridge (SMBridge), Version 1.0.3.1
Copyright (c) 2004 - OSA Technologies, an Avocent Company. All
Rights Reserved.

SMBridge>
```

Figure 5-26 Welcome message

5. Connect to the BMC on the server you wish to manage using the **connect** command, as shown in Figure 5-27.

```
SMBridge>connect -ip 9.42.171.237 -u USERID -p PASSWORD
SMBridge>
```

Figure 5-27 Connect to the BMC

The user ID and password here are ones that have previously been configured as users able to log into the BMC.

If the command was successful, you will be returned to the **SMBridge** command prompt.

6. You can now issue commands against the remote BMC, as described below.
7. To exit, enter the **exit** command.

Available commands

These commands listed in Table 5-5 are a super-set of those of the command-line interface. To get detailed help about a command, issue the **help** command. For example:

```
help power
```

Tip: Most of the commands available to the telnet interface are the same as those used in the CLI. The additional telnet commands are **console**, **sol**, and **reboot**.

Table 5-5 *SMBridge telnet subcommands*

Subcommand	Description and syntax
console	<p>Starts a Serial Over LAN (SOL) session with the BMC, displaying the text that has been redirected from the console to the serial port. There are no parameters.</p> <p>When you enter the console command, you will see:</p> <pre>Activating remote console now. Remote console is now active and ready for user input.</pre> <p>To return to the telnet session, press the tilde key followed by the period key, as in:</p> <pre>~.</pre>

Subcommand	Description and syntax
sol	<p>Used to enable or disable Serial Over LAN and to configure serial parameters to match the Console Redirection parameters of the remote server's BIOS. The options are:</p> <pre>sol enable sol disable sol config [-baud <i>baud_rate</i>] [-priv <i>privilege_level</i>] [-retry count <i>retry_count</i>] [-retry interval <i>retry_interval</i>]</pre>
reboot	<p>Performs the equivalent of a power off (graceful shutdown), power on, then starts the remote console. The options are:</p> <pre>reboot reboot -force</pre> <p>Note that the x236, x336, and x346 do not support the graceful shutdown option. The <code>-force</code> parameter is required on these servers.</p>
sysinfo	<p>Displays general system information related to the server and BMC. The options are:</p> <pre>sysinfo fru sysinfo id</pre> <p>id is the default if no parameter is specified.</p>
identify	<p>Controls the blue identification LED on the front panel of the server. The options are:</p> <pre>identify on [-t <seconds>] identify off</pre> <p>on is the default if no parameter is specified.</p>
power	<p>Controls the power options of the server. The options are:</p> <pre>power status power on power cycle power reset power off [-force]</pre> <p>status is the default if no parameter is specified.</p> <p>Note that the x236, x336, and x346 do not support the graceful shutdown option. The <code>-force</code> parameter is required on the server.</p>

Subcommand	Description and syntax
sel	<p>Performs operations with the System Event Log (SEL). The options are:</p> <pre> sel status sel get set get -last <n> sel get -begin <index1> -end <index2> sel get -begin <index1> -max <count> sel clear sel set -time <YYYY/MM/DD hh:mm:ss> </pre> <p>status is the default if no parameter is specified.</p>
help	Displays general help about all commands or help about a specific command.

5.5.4 Configuring Windows Server 2003 to support SOL

When you connect to the BMC using the SMBridge telnet server, you can remotely control the text console. With SOL, this also includes operating systems such as Windows Server 2003 and Linux.

Windows Server 2003 has two components that work with SMBridge and the BMC to provide out-of-band access to the operating system:

- ▶ Microsoft Emergency Messaging Service (EMS)
- ▶ Microsoft Special Administration Console (SAC)

Information about Microsoft EMS is available from:

http://www.microsoft.com/resources/documentation/WindowsServ/2003/standard/proddocs/en-us/EMS_topnode.asp

The SAC commands you can perform are summarized in Table 5-6.

Table 5-6 Windows Server 2003 Special Administration Console (SAC) commands

Command	Description
ch	Lists all channels.
cmd	Creates a command-prompt channel. You will be asked to log on.
crashdump	Manually generates a Stop error message and forces a memory dump file to be created.
d	Dumps the current kernel log.

Command	Description
f	Toggles the information output of the t-list command, which shows processes only, or shows processes and threads.
i	Lists the TCP/IP details of all network interfaces, and lets you configure the IP address, subnet mask, and gateway of a given network interface. To change the parameters, specify them as: network# IPaddress subnet gateway
id	Displays identification information about the server.
k pid	Ends the given process. pid is the process identification number you specify.
L pid	Lowers the priority of a process (and any associated child processes) to the lowest possible level.
lock	Restricts access to Emergency Management Services command-prompt channels. You must provide valid logon credentials to unlock a channel.
m pid mb-allow	Limits the memory usage of a process (and any associated child processes) to a specified number of megabytes. mb is the number of megabytes you specify.
p	Causes t-list command output to pause after displaying one full screen of information.
r pid	Raises the priority of a process and any associated child processes by one level.
restart	Restarts the server.
s	Displays or sets the system time. To set the time, use the format: mm/dd/yyyy hh:mm
shutdown	Shuts down and powers off the server. Terminates the console session and returns you to the SMBridge prompt.
t	Lists the processes and threads that are currently running.
? or help	Lists the available commands.

To exit SOL and return to the SMBridge prompt, press the tilde key and the period key (that is, ~.).

For more information, see:

http://www.microsoft.com/resources/documentation/WindowsServ/2003/standard/proddocs/en-us/EMS_SAC_commands.asp

To enable EMS on a Windows Server 2003, do the following.

1. Log in to Windows as an administrator.
2. Launch a command prompt and enter the command **bootcfg**. This command should produce the output shown in Figure 5-28.

```
C:\>bootcfg

Boot Loader Settings
-----
timeout:30
default:multi(0)disk(0)rdisk(0)partition(1)\WINDOWS

Boot Entries
-----
Boot entry ID: 1
OS Friendly Name: Windows Server 2003, Enterprise
Path: multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
OS Load Options: /fastdetect
```

Figure 5-28 Output from the bootcfg command

3. Examine the output. If there is more than one boot entry, then you will need to determine the default entry by looking at the default line under Boot Loader Settings and determining whether Boot Entry has a matching Path value. In our case, there is only one boot entry, 1.

4. Issue the following command, substituting your boot entry number in the /id parameter if it is not 1, as in our example:

```
bootcfg /ems on /port com1 /baud 19200 /id 1
```

Figure 5-29 shows the output of this command.

```
C:\>bootcfg /ems on /port com1 /baud 19200 /id 1
SUCCESS: Changed the redirection port in boot loader section.
SUCCESS: Changed the redirection baudrate in boot loader section.
SUCCESS: Changed the OS entry switches for line "1" in the BOOT.INI
file.
```

Figure 5-29 Changing the boot configuration

5. Reissue the **bootcfg** command to see the result. The changes in our example are highlighted, as shown in Figure 5-30 on page 209.

```
C:\>bootcfg /ems on /port com1 /baud 19200 /id 1
SUCCESS: Changed the redirection port in boot loader section.
SUCCESS: Changed the redirection baudrate in boot loader section.
SUCCESS: Changed the OS entry switches for line "1" in the BOOT.INI
file.

C:\>bootcfg

Boot Loader Settings
-----
timeout:          30
default:          multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
redirect:         COM1
redirectbaudrate:19200

Boot Entries
-----
Boot entry ID:    1
OS Friendly Name: Windows Server 2003, Enterprise
Path:            multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
```

Figure 5-30 The bootcfg command after enabling EMS

6. Reboot the server to have the changes take effect.

Note: To turn EMS off again, issue the following command:

```
bootcfg /ems off /id 1
```

Where 1 is the boot entry you have modified in the above steps. Reboot to bring the changes online.

Once you have rebooted and engaged the SMBridge console (see the **console** command in Table 5-5 on page 204), you will see the EMS console, as shown in Figure 5-31.

```
<?xml version="1.0"?>
<machine-info>
Computer is booting, SAC started and initialized.
<processor-architecture>x86</processor-architecture>
<os-version>5.2</os-version>
<os-build-number>3790</os-build-number>

SAC>
```

Figure 5-31 Microsoft Emergency Messaging Service console

You can now issue the various SAC commands described in Table 5-6 on page 206. For example, to start a command prompt, the commands are as follows (see Figure 5-32 on page 211).

Tip: After you start the SMBridge console, if you only get a blank screen, press Enter a few times to get the SAC> prompt.


```
SAC>cmd
The Command Prompt session was successfully launched.
SAC>
EVENT: A new channel has been created. Use "ch -?" for channel help.
Channel: Cmd0002
SAC>ch
Channel List

(Use "ch -?" for information About using channels)

# Status Channel Name
0 (AV) SAC
1 (AV) Cmd0002
SAC>ch -si 1
Name: Cmd0002
Description: Command Prompt
Type: <Esc><tab>0 to return to the SAC channel.
Use any other key to view this channel.

Please enter login credentials.
Username:
Domain:
Password:

Attempting to authenticate...

Microsoft Windows [Version 5.2.3790]
(C) Copyright 1985-2003

C:\WINDOWS\system32>
```

Figure 5-32 SAC commands to launch a command prompt channel

To close the command prompt channel, enter `exit`. To leave the channel open and return to the SAC prompt, press `Esc+Tab+0` (the number zero key) (three keys in sequence). To leave the remote console and return to SMBridge, press `tilde+period` (that is, `~.`)

5.5.5 Configuring Red Hat Linux to support SOL

You must configure Red Hat Linux to expose the Linux initialization (booting) process. This enables users to log in to the Linux console through an SOL session and directs output to the serial console. The following instructions are for Red Hat Enterprise Linux ES 2.1 or 3.0 to enable SOL:

1. Log in as root.
2. Modify the `/etc/inittab` file by adding the following line to the end of the `# Run gettys in standard runlevels` section to enable users to log in at the SOL console:

```
7:2345:respawn:/sbin/agetty -h ttyS1 19200 vt102
```

3. Modify the `/etc/securetty` file by adding the following line to enable users to log in as root at the SOL console:

```
ttyS1
```

For LILO users (GRUB users, please go to step 1 on page 214):

1. Modify the `/etc/lilo.conf` file:
 - a. Add `-Monitor` to the first default line.
 - b. Comment out the map line.
 - c. Comment out the message line.
 - d. In the first Image section, append `-Monitor` to the label line, and append the following line:

```
append="console=ttyS1,19200n8 console=tty1"
```

- e. Add the following lines between the two Image sections

```
# This will allow you to Interact with the OS boot via SOL
image=/boot/vmlinuz-2.4.9-e.12smp
label=linux-Interact
initrd=/boot/initrd-2.4.9-e.12smp.img
read-only
root=/dev/hda6
append="console=tty1 console=ttyS1,19200n8"
```

The result is shown in Figure 5-33 on page 213. The changes are highlighted.

```

prompt
timeout=50
default=linux-Monitor
boot=/dev/hda
#map=/boot/map
install=/boot/boot.b
#message=/boot/message
linear

# This will allow you to only Monitor the OS boot via SOL
image=/boot/vmlinuz-2.4.9-e.12smp
label=linux-Monitor
initrd=/boot/initrd-2.4.9-e.12smp.img
read-only
root=/dev/hda6
append="console=ttyS1,19200n8 console=tty1"

# This will allow you to Interact with the OS boot via SOL
image=/boot/vmlinuz-2.4.9-e.12smp
label=linux-Interact
initrd=/boot/initrd-2.4.9-e.12smp.img
read-only
root=/dev/hda6
append="console=tty1 console=ttyS1,19200n8"

image=/boot/vmlinuz-2.4.9-e.12
label=linux-up
initrd=/boot/initrd-2.4.9-e.12.img

```

Figure 5-33 Changes to the lilo.conf file

2. Enter lilo to store and activate the new LILO configuration.
3. Restart Linux.

When the operating system starts to boot, you will now see a LILO boot: prompt instead of the usual GUI interface. Pressing the Tab key while at this prompt will display the boot options. To load the operating system in interactive mode, you would enter:

```
linux-Interact
```

For GRUB users:

1. Modify the `/boot/grub/grub.conf` file as follows:

a. Comment out the splashimage line.

b. Add the following comment before the first title line:

```
# This will allow you to only Monitor the OS boot via SOL
```

c. Append `SOL Monitor` to the first title line.

d. Append the following text to the end of the kernel line of the first title section:

```
console=ttyS1,19200 console=tty1
```

e. Add the following lines between the two title sections:

```
# This will allow you to Interact with the OS boot via SOL
```

```
title Red Hat Linux (2.4.9-e.12smp) SOL Interactive
```

```
root (hd0,0)
```

```
kernel /vmlinuz-2.4.9-e.12smp ro root=/dev/hda6 console=tty1
```

```
console=ttyS1,19200
```

```
initrd /initrd-2.4.9-e.12smp.img
```

The result is shown in Figure 5-34 on page 215. The changes are highlighted.

2. Restart Linux.

```

#grub.conf generated by anaconda
#
# Note that you do not have to rerun grub after making changes to this file
# NOTICE: You have a /boot partition. This means that
#           all kernel and initrd paths are relative to /boot/, eg.
#           root (hd0,0)
#           kernel /vmlinuz-version ro root=/dev/hda6
#           initrd /initrd-version.img
#boot=/dev/hda
default=0
timeout=10
#splashimage=(hd0,0)/grub/splash.xpm.gz
# This will allow you to only Monitor the OS boot via SOL
title Red Hat Enterprise Linux ES (2.4.9-e.12smp) SOL Monitor
    root (hd0,0)
    kernel /vmlinuz-2.4.9-e.12smp ro root=/dev/hda6 console=ttyS1,19200 console=tty1
    initrd /initrd-2.4.9-e.12smp.img

# This will allow you to Interact with the OS boot via SOL
title Red Hat Linux (2.4.9-e.12smp) SOL Interactive
    root (hd0,0)
    kernel /vmlinuz-2.4.9-e.12smp ro root=/dev/hda6 console=tty1 console=ttyS1,19200
    initrd /initrd-2.4.9-e.12smp.img

title Red Hat Enterprise Linux ES-up (2.4.9-e.12)
    root (hd0,0)

```

Figure 5-34 Changes to the grub.conf file

5.5.6 Configuring SUSE LINUX to support SOL

You must configure SUSE LINUX to expose the Linux initialization (booting) process. This enables users to log in to the Linux console through an SOL session and directs output to the serial console. The following instructions are for SUSE LINUX Enterprise Server 8.0 to enable SOL:

1. Log in as root.
2. Modify the /etc/inittab file by adding the following line to the end of the #getty-programs for the normal runlevels section to enable users to log in at the SOL console:

```
7:2345:respawn:/sbin/agetty -h ttyS1 19200 vt102
```

3. Modify the `/etc/securetty` file by adding the following line after the `tty6` line to enable users to log in as root at the SOL console:

```
ttyS1
```

4. Modify the `/boot/grub/menu.lst` file as follows:

- a. Comment out the `gfxmenu` line.

- b. Add the following comment line before the first title line:

```
# This will allow you to only Monitor the OS boot via SOL
```

- c. Append `SOL Monitor` to the first title line.

- d. Append the following text to the kernel line of the first title section:

```
console=ttyS1,19200 console=tty1
```

- e. Add the following lines between the first two title sections:

```
# This will allow you to Interact with the OS boot via SOL
```

```
title linux SOL Interactive
```

```
kernel (hd0,1)/boot/vmlinuz root=/dev/hda2 acpi=oldboot vga=791
```

```
    console=tty1 console=ttyS1,19200
```

```
initrd (hd0,1)/boot/initrd
```

The result is shown in Figure 5-35 on page 217. The changes are highlighted.

5. Restart Linux.

```

#gfxmenu (hd0,1)/boot/message
color white/blue black/light-gray
default 0
timeout 8

# This will allow you to only Monitor the OS boot via SOL
title linux SOL Monitor
# Note: The following "kernel" line is all one line, not two separate lines
# The text has wrapped in this example
kernel (hd0,1)/boot/vmlinuz root=/dev/hda2 acpi=oldboot vga=791
console=ttyS1,19200 console=tty1
initrd (hd0,1)/boot/initrd

# This will allow you to Interact with the OS boot via SOL
title linux SOL Interactive
# Note: The following "kernel" line is all one line, not two separate lines
# The text has wrapped in this example
kernel (hd0,1)/boot/vmlinuz root=/dev/hda2 acpi=oldboot vga=791 console=tty1
console=ttyS1,19200
initrd (hd0,1)/boot/initrd

title floppy
root
chainloader +1
title failsafe
kernel (hd0,1)/boot/vmlinuz.shipped root=/dev/hda2 ide=nodma apm=off vga=normal

```

Figure 5-35 Changes to the menu.lst file

5.5.7 Connecting via the command-line interface

OSA SMBridge also supplies a command-line tool, `smbridge.exe`, which lets you perform a subset of the functions that you can perform using the telnet server. Specifically, the tasks that are missing are the ability to remotely control the text console of the server via SOL. The CLI does, however, let you connect to the server via a serial connection.

In CLI mode, SMBridge supports out-of-band access through the LAN or serial port to one server at a time. However, multiple IPMI sessions can run simultaneously on the same remote server. LAN connections are via Ethernet and serial connections are typically via a null modem.

To run SMBridge in CLI mode, simply open a command prompt/shell prompt at the directory where SMBridge is installed and issue the **smbridge** command.

- ▶ With Windows, SMBridge is installed by default in c:\Program Files\OSA.
- ▶ With Linux, it is installed by default in /usr/sbin.

The syntax is as follows for Ethernet or Serial connectivity.

For Ethernet connections:

```
smbridge -ip address -u user -p password subcommand
```

Where

- ▶ -ip address is the IP address or host name of the remote server.
- ▶ -u user -p password is a valid service processor user ID and password (default USERID/PASSWORD).

For Serial connections:

```
smbridge -com serialport [-baud baudrate] [-flow flowcontrol] -u user  
-p password subcommand
```

Where:

- ▶ -com serialport specifies the serial port on remote server. In Windows systems, it can be 1 for COM1, 2 for COM2, and so on. In Linux systems, it can be ttyS0, ttyS1, and so on.
- ▶ -baud baudrate specifies the baud rate you wish to communicate at, such as 9600 and 19200. It should match the one set in BIOS of the remote server (in the Remote Console Redirection window). If not specified, it defaults to 19200.
- ▶ -flow flowcontrol specifies the flow control. If not specified, it defaults to CTS (hardware flow control). The options are:
 - CTS: Hardware flow control
 - XON: Software flow control
 - NONE: No flow control

The valid subcommands and the syntax of those commands is listed in Table 5-7 on page 219. For more information about the syntax, issue the **-help** command.

For example:

```
smbridge -help power
```


Table 5-7 *SMBridge CLI subcommands*

Subcommand	Description and syntax
<p>sysinfo</p>	<p>Displays general system information related to the server and BMC. The options are:</p> <pre>sysinfo fru sysinfo id</pre> <p>id is the default if no parameter is specified.</p>
<p>identify</p>	<p>Controls the blue identification LED on the front panel of the server. The options are:</p> <pre>identify on [-t <seconds>] identify off</pre> <p>on is the default if no parameter is specified.</p>
<p>power</p>	<p>Controls the power options of the server. The options are:</p> <pre>power status power on power cycle power reset power off [-force]</pre> <p>status is the default if no parameter is specified.</p> <p>Note: The x236, x336, and x346 do not support the graceful shutdown option. The -force parameter is required on these servers.</p>
<p>sel</p>	<p>Performs operations with the System Event Log (SEL). The options are:</p> <pre>sel status sel get set get -last <n> sel get -begin <index1> -end <index2> sel get -begin <index1> -max <count> sel clear sel set -time <YYYY/MM/DD hh:mm:ss></pre> <p>status is the default if no parameter is specified.</p>

5.6 Predictive Failure Analysis (PFA)

PFA is a mechanism developed by IBM, that periodically measures critical components such as hard drives, power supplies, fans, processors, and memory to reduce unscheduled system downtime.

If a predefined threshold of one of these components is exceeded, for example, the number of single bit errors on a memory module, a PFA alert will be issued. This enables you to replace the component before a failure actually occurs.

If a PFA alert is raised, the Error LED on the light path diagnostic panel will be lit and the BMC logs an event that can be forwarded to IBM Director. The RSA II can also be configured to forward alerts via e-mail or to an IBM Director server (see also 5.2, “Remote Supervisor Adapter II SlimLine” on page 164).

5.7 IBM Dynamic System Analysis

IBM Dynamic System Analysis (DSA) is an information gathering tool, which collects and analyzes system information in order to provide the appropriate data to the IBM support in case of problems.

5.7.1 Introduction

DSA can be downloaded from the following Web site:

http://www.ibm.com/servers/eserver/xseries/systems_management/dsa.html

DSA collects information about the following aspects of a system:

- ▶ System configuration
- ▶ Installed applications and hot fixes
- ▶ Device drivers and system services
- ▶ Network interfaces and settings
- ▶ Performance data and running process details
- ▶ Hardware inventory, including PCI information
- ▶ Vital product data and firmware information
- ▶ SCSI device sense data
- ▶ EXA chipset uncorrectable error register information
- ▶ ServeRAID configuration

- ▶ Application, system, security, ServeRAID, and service processor system event logs

As the time of writing this redbook, the following operating systems are supported by DSA: Microsoft Windows 2000 Server, Advanced Server and Datacenter Server with Service Pack 3 or later, Microsoft Windows 2003 Standard, Enterprise, and Web Editions. DSA versions that support 64-bit Windows editions and Linux will be available in the future.

Two versions of DSA are available:

- ▶ DSA Portable Edition is meant to be run from portable media, such as a memory key or CD-ROM. It runs from the command prompt, unpacks to and runs from a temporary directory, but does not actually install on the system.
- ▶ DSA Installable Edition provides a permanent installation of DSA onto a system. You can run DSA Installable Edition from a command line, from the Start menu, as an IBM Director task, or in conjunction with UpdateXpress.

DSA can be downloaded here:

http://www.ibm.com/servers/eserver/xseries/systems_management/dsa.html

Tip: The current and previous versions of DSA can be downloaded from:

<http://www.ibm.com/pc/support/site.wss/document.do?lndocid=SERV-DSA>

Figure 5-36 shows an example of the DSA output in HTML format.

Installed Hotfixes Device Drivers System Services Network Settings Resource Utilization Processes OS Configuration	Computer System	
	Manufacturer	IBM
	Version	
	Product Name	eserver xSeries 460-[88721RZ]-
	Serial Number	23A0029
	System UUID	40de33a3-ea12-b601-9613-000d609812a2
	Operating System	
	Computer Name	X460
	Product Name	Windows Server 2003
	Suite Type	Enterprise Edition
	Version	5.2
	Service Pack	Service Pack 1
Build Number	3790	
Vendor	Microsoft Corporation	
Registered Owner	Sergey&Christian	
Product ID	69713-640-0394874-45654	
Installation Date	06/15/2005 10:04:42	
Uptime	5 minutes 29 seconds	
Time of Last Boot	06/29/2005 19:37:54	
Local Date and Time	06/29/2005 19:43:23	
Hardware		
Hardware Inventory		
PCI Information		
Firmware/VPD		
SP Configurations		
Environmentals		
Drive Health		
LSI Controller		
ServeRAID		
ServeRAID Logs		
Windows Logs		
Application Event		
System Event		
Security Event		
SP Logs		
ASM Event		
ISMP Event		
IPMI Event		

Figure 5-36 DSA results Web page

5.7.2 Usage

DSA can be used in several ways, depending on the desired result. In the following examples, we refer to the DSA Portable Edition, which is simply launched by executing the downloaded program from a command-line prompt.

The Installable Edition is installed by default in C:\Program Files\IBM\DSA\. Once installed, you can launch it in a variety of ways:

- ▶ From the Start menu
- ▶ From a command prompt
- ▶ As an IBM Director task
- ▶ In conjunction with UpdateXpress

The option switches listed for the DSA Portable Edition apply to the installable edition as well.

When DSA is started without any additional command-line parameters, a compressed XML file is created, which can be sent via e-mail to an IBM support agent:

```
C:\temp>dsa101p.exe
```

To create a human readable output in HTML format instead of the XML file, use the following parameters:

```
C:\temp>dsa101p.exe /s /a /x /v
```

Servers that have an Internet connection can create and upload the compressed XML file via FTP to a secured server, so that is available for the IBM support agent for analysis. This is done using the following command:

```
C:\temp>dsa101p.exe /t
```

It is also possible to parse a compressed XML created on an other machine file to inspect the output using the following command:

```
C:\temp>dsa101p.exe /s /a /x /i C:\temp\ 88721RC_KK00011_20050704-081500
```

For further details about the command options, see the DSA readme.txt that accompanies the downloadable files in the above URLs.

5.8 Partition management

Scalable partitions can be managed using the RSA Web GUI. This is described in 3.3.2, “Scalable system setup” on page 110 and 3.3.3, “Partitioning” on page 116.

Abbreviations and acronyms

AC	alternating current	CTP	composite theoretical performance
ACPI	advanced control and power interface	CTS	clear to send
ALU	arithmetic logic unit	DB	database
AMD	Advanced Micro Devices	DBA	database administrator
APIC	Advanced Programmable Interrupt Controller	DHCP	Dynamic Host Configuration Protocol
AS	Australian Standards	DIMM	dual inline memory module
ASCII	American National Standard Code for Information Interchange	DMI	Desktop Management Interface
ASIC	application-specific integrated circuit	DNS	Domain Name System
ASM	Advanced System Management	DP	dual processor
ASR	automatic server restart	DPF	Database Partitioning Feature
AWE	Address Windowing Extensions	DRAM	dynamic random access memory
BI	business intelligence	DSA	digital signature algorithm
BIOS	basic input output system	EB	exabyte
BMC	baseboard management controller	ECC	error checking and correcting
CD	compact disk	EIDE	enhanced IDE
CD-ROM	compact disk read only memory	EMEA	Europe, Middle East, Africa
CE	Conformité Européene	EMS	Emergency Messaging Service
CIM	Common Information Model	ERO	Export Regulation Office
CLI	command-line interface	ERP	enterprise resource planning
COA	Certificate of Authenticity	EXA	Enterprise X-Architecture
COM	Component Object Model	FAMM	full array memory mirroring
CPU	central processing unit	FQDN	fully qualified domain names
CRM	Customer Relationships Management	GB	gigabyte
CSG	Chip Select Group	GPR	general purpose register
		GRUB	grand unified bootloader
		GUI	graphical user interface
		HAL	hardware abstraction layer
		HAM	hot-add memory

HBA	host bus adapter	MIB	management information base
HDD	hard disk drive	MIOC	Memory and I/O Controller
HID	human interface device	MP	multiprocessor
HPMA	high performance memory array	MPA	Management Processor Assistant
HTML	Hypertext Markup Language	MTOPS	millions composite theoretical operations per second
HTTP	hypertext transmission protocol	MXE	modular expansion enclosure
I/O	input/output	NIC	network interface card
IBM	International Business Machines Corporation	NMI	non-maskable interrupt
IBS	International Business System	NONRED	not redundant, the non-redundant
ICC	International Competency Center	NUMA	Non-Uniform Memory Access
ID	identifier	OEM	other equipment manufacturer
IP	Internet Protocol	OLAP	online analytical processing
IPMI	Intelligent Platform Management Interface	OLTP	online transaction processing
IPX	internetwork packet exchange	OS	operating system
ISO	International Organization for Standards	PAE	Physical Address Extension
ISU	Industry Solutions Utilities	PC	personal computer
IT	information technology	PCI	peripheral component interconnect
ITSO	International Technical Support Organization	PFA	Predictive Failure Analysis
IXA	Integrated xSeries Adapter	PKT	packet
JDE	JD Edwards	POST	power on self test
KB	Kilobyte	PXE	Pre-boot-execution
KDE	K Desktop Environment	RAC	Real Application Clusters
KVM	keyboard video mouse	RAID	redundant array of independent disks
LAN	local area network	RAM	random access memory
LDAP	Lightweight Directory Access Protocol	RBS	redundant bit steering
LED	light emitting diode	RHEL	Red Hat Enterprise Linux
LILO	Linux loader	ROLAP	Relational Online Analytical Processing
MAC	media access control	ROM	read-only memory
MB	megabyte	RPM	Red Hat Package Manager
		RSA	Remote Supervisor Adapter

SA	solution assurance	SSE	Streaming SIMD Extensions
SAC	Special Administration Console	SSL	Secure Sockets Layer
SAN	storage area network	TB	terabyte
SAPR	Solution Assurance Product Review	TCO	Total Cost of Ownership
SAR	Solution Assurance Review	TCP/IP	Transmission Control Protocol/Internet Protocol
SAS	Serial Attached SCSI	UDB	DB2 Universal Database™
SCM	Supply Chain Management	URL	Uniform Resource Locator
SCON	server consolidation	USB	universal serial bus
SCSI	small computer system interface	VMFS	virtual machine file system
SDRAM	static dynamic RAM	VRM	voltage regulator module
SEL	System Event Log	WAN	wide area network
SIMD	Single Instruction Multiple Data	WMI	Windows Management Instrumentation
SL	SlimLine	XML	Extensible Markup Language
SLES	SUSE LINUX Enterprise Server	XON	transmitter on
SLIP	Serial Line Internet Protocol		
SMB	server message block		
SMBIOS	system management BIOS		
SMI	Structure of Management Information		
SMP	symmetric multiprocessing		
SMS	System Managed Space		
SMTP	simple mail transfer protocol		
SNA	systems network architecture		
SNMP	Simple Network Management Protocol		
SOL	Serial over LAN		
SPEC	Standard Performance Evaluation Corporation		
SPORE	ServerProven Opportunity Request for Evaluation		
SQL	structured query language		
SRAT	Static Resource Allocation Table		

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

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

The following Redbooks, Redpapers, and Hints and Tips are relevant. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *IBM Director 5.10*, SG24-6188
- ▶ *IBM @server xSeries 260 Solution Assurance Product Review Guide*, REDP-4007
- ▶ *IBM @server xSeries and BladeCenter Server Management*, SG24-6495
- ▶ *IBM @server xSeries BMC — Firmware and Drivers Quick Reference*, TIPS0532
- ▶ *Implementing VMware ESX Server 2.1 with IBM TotalStorage FASTT*, SG24-6434
- ▶ *Integrating IBM Director with Enterprise Management Solutions*, SG24-5388
- ▶ *Introducing Windows Server x64 on IBM @server xSeries Servers*, REDP-3982
- ▶ *Remote Supervisor Adapter II Family — Firmware and Drivers Quick Reference*, TIPS0534
- ▶ *ServeRAID Adapter Quick Reference*, TIPS0054
- ▶ *Server Consolidation with the IBM @server xSeries 440 and VMware ESX Server*, SG24-6852
- ▶ *Server Consolidation with VMware ESX Server*, REDP-3939
- ▶ *Tuning IBM @server xSeries Servers for Performance*, SG24-5287
- ▶ *VMware ESX Server: Scale Up or Scale Out?*, REDP-3953

Other publications

These publications are also relevant as further information sources. These are available from the URLs listed below or generally from:

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

- ▶ *IBM @server xSeries Solution Assurance Trigger Criteria (Worldwide)*, found at:

<http://w3.ibm.com/support/assure/assur30i.nsf/PubAllNum/SA200?OpenDocument>

- ▶ *xSeries 260, xSeries 366, xSeries 460, and xSeries MXE-460 FAQ — Hints and Tips*, found at:

<http://www.pc.ibm.com/support?page=MIGR-61395>

- ▶ *xSeries Configuration and Options Guide*, found at:

<http://www.pc.ibm.com/support?page=SCOD-3ZVQ5W>

- ▶ *Solving POST 1801 PCI Allocation Errors*, found at:

<http://www.pc.ibm.com/support?page=MIGR-61663>

- ▶ *OSA System Management Bridge User's Guide*, found at:

<http://www.ibm.com/pc/support/site.wss/MIGR-57816.html>

xSeries 460 and MXE-460 publications

- ▶ *IBM @server xSeries 460 and MXE 460 Enablement Kit for export-restricted configurations*, found at:

<http://www.pc.ibm.com/support?page=MIGR-60747>

- ▶ *IBM @server xSeries 460 Type 8872 and xSeries MXE 460 Type 8874 Installation Guide*, found at:

<http://www.pc.ibm.com/support?page=MIGR-59865>

- ▶ *IBM @server xSeries 460 Type 8872 and xSeries MXE 460 Type 8874 Problem Determination and Service Guide*, found at:

<http://www.pc.ibm.com/support?page=MIGR-59864>

- ▶ *IBM @server xSeries 460 Type 8872 and xSeries MXE 460 Type 8874 User's Guide*, found at:

<http://www.pc.ibm.com/support?page=MIGR-59863>

- ▶ *Rack Installation Instructions*, found at:

<http://www.pc.ibm.com/support?page=MIGR-59862>

xSeries 366 publications

- ▶ *IBM @server xSeries 366 Type 8863 Installation Guide*, found at:
<http://www.pc.ibm.com/support?page=MIGR-58972>
- ▶ *IBM @server xSeries 366 Type 8863 Problem Determination and Service Guide*, found at:
<http://www.pc.ibm.com/support?page=MIGR-58887>
- ▶ *IBM @server xSeries 366 Type 8863 User's Guide*, found at:
<http://www.pc.ibm.com/support?page=MIGR-58973>
- ▶ *Rack Installation Instructions*, found at:
<http://www.pc.ibm.com/support?page=MIGR-58974>

xSeries 260 publications

- ▶ *IBM @server xSeries 260 Type 8865 Installation Guide*, found at:
<http://www.pc.ibm.com/support?page=MIGR-61926>
- ▶ *IBM @server xSeries 260 Type 8865 Problem Determination and Service Guide*, found at:
<http://www.pc.ibm.com/support?page=MIGR-61239>
- ▶ *IBM @server xSeries 260 Type 8865 User's Guide*, found at:
<http://www.pc.ibm.com/support?page=MIGR-61110>

Online resources

The following Web sites and URLs are also relevant as further information sources.

IBM Web pages

Here are some useful IBM web links.

Solutions

- ▶ IBM-SAP alliance
<http://www.ibm.com/solutions/sap/us/en>
- ▶ IBM-SAP alliance (IBM internal network users only)
<http://w3.ncs.ibm.com/solutions/SAP>

- ▶ SAP sizing documents
<http://www.ibm.com/support/techdocs/atmastr.nsf/PubAllNum/PRS261>
- ▶ IBM hardware sizing for applications
<http://www.ibm.com/servers/solutions/finder/portal/hw-sizing.html>
- ▶ SAP Benchmark results
<http://www.sap.com/solutions/benchmark>
- ▶ IBM-Siebel alliance
<http://www.ibm.com/solutions/businesssolutions/siebel>
- ▶ IBM-Siebel alliance (IBM internal network users only)
<http://w3.ncs.ibm.com/solutions/Siebel>
- ▶ Siebel sizing document
<http://www.ibm.com/support/techdocs/atmastr.nsf/PubAllNum/PRS593>
- ▶ IBM SAP International Competence Center (IBM internal network users only)
<http://w3.ncs.ibm.com/solution.nsf/SOL/BTIY-5PBJ73>
- ▶ IBM-VMware alliance
<http://www.ibm.com/eserver/xseries/vmware>
- ▶ IBM Datacenter Solutions with High Performance xSeries Servers
<http://www.ibm.com/servers/eserver/xseries/windows/datacenter.html>

Configurations

- ▶ ServerProven
<http://www.ibm.com/servers/eserver/serverproven/compat/us>
- ▶ ServerProven operating system compatibility matrix
<http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/matrix.shtml>
- ▶ DataCenter Supported Configurations
<http://www.pc.ibm.com/support?page=MIGR-4P7RG3>
- ▶ IBM Export Regulations Office (IBM internal network users only)
<http://w3.ibm.com/chq/ero>
- ▶ Solution Assurance (IBM internal network users only)
<http://w3.ibm.com/support/assure>
- ▶ Solution Assurance (PartnerWorld users only)
<http://www.ibm.com/partnerworld/techsupport>

- ▶ Solution Assurance Trigger Criteria (IBM internal network users only)
<http://w3.ibm.com/support/assure/assur30i.nsf/PubAllNum/SA200>

Support and downloads

- ▶ AP Technical Sales Support (IBM internal network users only)
<http://w3.ibm.com/support/ap>
- ▶ Contact PartnerWorld
http://www.ibm.com/partnerworld/pwhome.nsf/weblook/cpw_index.html
- ▶ xSeries driver matrix
<http://www.pc.ibm.com/support?page=MIGR-4JTS2T>
- ▶ UpdateXpress
<http://www.pc.ibm.com/support?page=MIGR-53046>
- ▶ OS installation instructions for xSeries systems
<http://www.pc.ibm.com/support?page=MIGR-4QLNTQ>
- ▶ Installing VMware ESX Server 2.5.1
<http://www.pc.ibm.com/support?page=MIGR-60546>
- ▶ Windows 2000 Server HAL for the x460
<http://www.pc.ibm.com/support?page=MIGR-61296>
- ▶ ServerGuide
<http://www.pc.ibm.com/support?page=MIGR-4ZKPPT>
- ▶ IBM Director
<http://www.pc.ibm.com/support?page=SERV-DIRECT>
- ▶ Dynamic System Analysis
<http://www.ibm.com/pc/support/site.wss/document.do?lndocid=SERV-DSA>
- ▶ SMBridge BMC utility
<http://www.pc.ibm.com/support?page=MIGR-57729>
- ▶ IBM Director downloads
http://www.ibm.com/servers/eserver/xseries/systems_management/xseries_sm/dwnl.html
- ▶ System Management Bridge Baseboard Management Controller CLI and Remote Console Utility - Servers
<http://www.ibm.com/pc/support/site.wss/MIGR-57729.html>

Intel Web pages

- ▶ EM64T technology
<http://www.intel.com/technology/64bitextensions>
- ▶ Hyper-Threading Technology
<http://www.intel.com/technology/hyperthread>
- ▶ Execute Disable bit
<http://www.intel.com/cd/ids/developer/asmo-na/eng/149308.htm>
- ▶ *Intelligent Platform Management Interface Specification V1.5*, found at:
ftp://download.intel.com/design/servers/ipmi/IPMIV1_5rev1_1-012904markup.pdf

Microsoft Web pages

- ▶ Microsoft Support on VMware
<http://support.microsoft.com/kb/897615>
- ▶ Windows Server Catalog
<http://www.microsoft.com/windows/catalog/server>
- ▶ Microsoft Emergency Messaging Service
http://www.microsoft.com/resources/documentation/WindowsServ/2003/standard/proddocs/en-us/EMS_topnode.asp
- ▶ Special Administration Console commands
http://www.microsoft.com/resources/documentation/WindowsServ/2003/standard/proddocs/en-us/EMS_SAC_commands.asp

Others

- ▶ “x86 Servers Brace for a Hurricane” by David Kanter, found at:
<http://www.realworldtech.com/page.cfm?ArticleID=RWT042405213553>
- ▶ “Error Correcting Memory, Part 1”, by David Wang, which describes memory hot-swapping, scrubbing, and bit steering:
<http://www.realworldtech.com/page.cfm?ArticleID=RWT121603153445>
- ▶ VMware Download Center
<http://www.vmware.com/download>
- ▶ Java Runtime Environment Version 5.0 Update 6 Manual Download
<http://www.java.com/en/download/manual.jsp>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Numerics

05K9276, USB diskette drive 141
110V power 127
13M7409, memory card 95
13M7414, 2.3m scalability cable 107
13M7416, 2.9m scalability cable 107
13N0694, 3.16 GHz Cranford processor 90
13N0695, 3.66 GHz Cranford processor 90
13N0713, 3.33 GHz Potomac processor 90
13N0714, 3.00 GHz Potomac processor 90
13N0715, 2.83 GHz Potomac processor 90
13N2227, ServeRAID-8i 123
19K4164, USB-parallel port cable 35
25R8941, 2.67 GHz Paxville processor 90
25R8942, 3.00 GHz Paxville processor 90
30R5145, two 4 GB DIMMs 96
31R1558, export enablement CD-ROM 119
32R0719, Tower To Rack Conversion Kit 14
39Y6580, Dual Core X3 Upgrade Kit 19, 91
40K2450, memory card for 32-core 119
4816, Software Update Subscriptions 150
64-bit computing, benefits 26
64-bit processor 23
73P2865, two 512 MB DIMMs 96
73P2866, two 1 GB DIMMs 96
73P2867, two 2 GB DIMMs 96
8863 13
8865 14
8872 12, 149, 151
8874 12, 149, 151

A

Active Memory 29
Adaptec AIC-9410 31, 122
adjacent sector prefetch 93
AMD64 23
application integration 58
applications 52, 68
 scalability 68

B

Baseboard Management Controller

See BMC

benchmarks 45, 69
benefits 44
BIOS
 BMC configuration 122, 179
 Clustering Technology 93
 ESX Server 158
 High Performance Memory Array 29
 Hyper-Threading 92
 memory configuration 29, 102
 OS USB Selection 121
 Periodic SMI Generation 121
 Processor Adjacent Sector Prefetch 93
 Processor Exectue Disable Bit 94
 Processor Hardware Prefetcher 94
 processor settings 92
 Remote Console Redirection 196
 RSA II 166
 RSA II setup 120
 SMBridge configuration 196
 updating 140
 Windows Server 2003 146
block diagram
 SAS 32
 X3 Architecture 8
blue screen of death 165
BMC 37, 177
 add to IBM Director 187
 add users 188
 alert forwarding 193
 BIOS settings 179
 clock 181
 connectivity 177
 drivers 182
 event log 180
 features 177
 firmware 140, 178
 IBM Director integration 187
 MPCLI 193
 ports 186
 remote control 182
 SEL 180
 SMBridge 194
 TCP/IP ports 186

BMC configuration 122
bootcfg command 208
Broadcom Ethernet 35
Business Intelligence 53
buttons 15

C

cabling 107
cache 19, 28
 XceL4v 27
Calgary 8
Certificate of Authenticity 148
chassis 4
Chassis ID 113
Chipkill memory 30
chipset
 comparison 42
 SAS 33
 X3 Architecture 3
 XceL4v 27
Clustering Technology 93
comparison 4, 45
compatibility mode, EM64T 24
complex 5
composite theoretical performance 118
connectors 15
console redirection 196
consolidation 56
cores 5
CPUs 18
Cranford 4, 18
 memory addressing 27
create a partition 112
CRM 53

D

data integration 58
database applications 53
Datacenter
 See Windows Datacenter
Datacenter High Availability Program Offering 148
Datacenter Scalable Offering 150
DB2 80
DDR2 memory 28, 96
default IP address, RSA II 166
default stripe size 123
diagnostics 140
differences 4

DIMMs 9, 28, 96
diskette drive 141
driver matrix 141
drivers
 BMC 182
 IPMI device drivers 182
 Red Hat 154
 RSA II 171
 RSA II Linux driver 173
 SAS 34
 ServeRAID 34
 SUSE LINUX 154
 Windows 2000 Server 152
 Windows Server 2003 145
DS300 126
Dual Core X3 Upgrade Kit 19
dual-core processors 5, 20
Dynamic System Analysis 220

E

eCommerce 54
EIDE interface 35
EM64T 23
Emergency Messaging Service 201
enterprise management 162
EnterpriseOne 73
ERP 52
event log 180
expanders, SAS 32
Export Regulation Office 119
export regulations 118
express models 13

F

far memory 10
fault tolerance 98
FC2-133 HBA 126
features 3, 44
firmware 106
 BMC 178
 RSA II 170
front view 15
frontside bus 21

G

Gallatin
 memory addressing 27

General Purpose Registers 24
grocery store analogy 10

H

HAL
 Windows 2000 Server 153
 Windows Server 2003 145
hardware prefetch 94
High Performance Memory Array 29
hot-add memory 30, 100
hot-replace
 memory 30
hot-swap memory 99
Hurricane 8–9
Hyper-Threading 21, 92
 optimized 138
 OS support 137

I

IA-32e mode, EM64T 25
IBM Director 162
 add users 188
 adding service processor 187
 alert forwarding 193
 integration with RSA and BMC 187
 MPCLI 193
 padlock icon 188
 Server Configuration Manager 188
 upward integration 162
IBM Export Regulation Office 119
identical processors 18
installation
 Red Hat 154
 SUSE LINUX 154
 VMware ESX Server 158
 Windows 2000 Server 152
 Windows Server 2003 145
Integrated xSeries Adapter 178
interleaving 9, 28
IPMI device drivers 182
Itanium 2
 64-bit implementation 23
 memory addressing 27

J

J.D. Edwards 73
Java runtime 174

K

KVM 174

L

latency 10
light path diagnostics 38, 220
local memory 10
logical processors 21

M

Maintenance Update Subscription 150
matrix 141
memory 28
 64-bit addressing 24
 Active Memory 29
 addressability by processor 27
 BIOS settings 102
 Chipkill memory 30
 disabling during POST 30
 far 10
 hot-add 30, 100
 hot-swap 30, 99
 interleaving 28
 latency 10
 local 10
 memory card 13M7409 95
 memory mirroring 30, 98
 Memory ProteXion 29, 101
 memory scrubbing 29
 multi-node configurations 106
 NUMA 10
 performance 129
 ports 9
 remote 10
 request queue 10
 rules 97
 throughput 9
 use by XceL4v cache 141
merge timeout 114
merging nodes 107
MIB files, RSA II 176
Microsoft SQL Server 79
Microsoft Windows Server support 136
MIOC 8
models 12
 express models 13
 Windows Datacenter 149, 151
 x260 14

- x366 13
- x460 12
- modular expansion enclosure 5
- MPCLI
 - add user 193
 - alert forwarding 193
- MTOPS 118
- multi-node
 - cables 107
 - concept 4
 - configurations 104
- MXE-460 5
 - comparison with x460 66
 - drive bays 33
 - export regulations 118
 - front view 15
 - memory cards 28
 - memory configuration 96
 - models 12
 - OS support 136
 - power supplies 127
 - processor upgrades 90
 - rear view 16
 - removable media 125
 - x460 comparison 66
- mySAP 72

N

- NetWare support 136
- Nocona 23
 - memory addressing 27
- nodes 4, 67
- NUMA 9

O

- operating systems 136
 - Hyper-Threading 137
 - Red Hat 153
 - scalability 137
 - SUSE LINUX 153
 - VMware ESX Server 155
 - Windows 2000 Server 152
 - Windows Datacenter 147
 - Windows Server 2003 142
- Opteron 23
 - memory addressing 27
- Oracle 77
- OS USB Selection 121

- OSA SMBridge 194
- Other OS setting, RSA II 167
- overview 1

P

- PAE switch 147
- parallel ports 35
- partial merge 115
- partial mirroring 98
- partitioning 6, 67, 116
- Paxville 4, 18, 20
 - memory addressing 27
- pay-as-you-grow 65
- PCI bridge 9
- PCI scan order 35
- PCI-X slots 34
- PeopleSoft 77
- performance 129
 - CPU settings 130
 - memory 9
 - memory cards 129
 - memory settings 129
 - number of DIMMs 129
 - PCI-X in multi-node 130
- Periodic SMI Generation 121
- PFA 220
- physical consolidation 58
- PKT files
 - RSA II 170
- ports
 - BMC 186
 - memory 9
 - RSA II 166, 176
 - SAS 32
 - SMBridge 198
- positioning 41, 66
- Potomac 4, 18
 - memory addressing 27
- power control 107
 - RSA II 165
 - SMBridge 205
- power supply considerations 126
- primary nodes 7
- Primary Scalable Node 113
- Processor Adjacent Sector Prefetch 93
- processor cache 19
- Processor Exectue Disable Bit 94
- Processor Hardware Prefetcher 94

- processor tray 19
- processors 18
 - 64-bit 23
 - BIOS options 92
 - Clustering Technology 93
 - dual-core 20
 - dual-core upgrade 91
 - hardware prefetch 94
 - Hyper-Threading 21, 92, 137
 - memory addressability 27
 - multi-node configurations 106
 - Processor Adjacent Sector Prefetch 93
 - rules 91

Q

- queuing time 11

R

- RAID levels 123
- rear view 15
- Red Hat 153
 - drivers 154
 - installation 154
 - RSA II drivers 173
 - SMBridge installation 200
 - SMBridge support 196
 - SOL, enabling 212
 - support 136
 - telnet 202
- Redbooks Web site 229
 - Contact us xi
- redundant bit steering 101
- remote console redirection 196
- remote control
 - BMC 182
 - RSA II 174
- remote memory 10
- Remote Supervisor Adapter II SlimLine 35–36
- request queue 10
- RSA II 164
 - add to IBM Director 187
 - add users 188
 - alert forwarding 193
 - alerts 164
 - BIOS settings 166
 - blue screen of death 165
 - configuration 120
 - connectivity 166

- default IP address 166
- default user 170
- DHCP 167
- driver 171
- event log 164
- features 164
- firmware 140, 170
- health monitoring 164
- IBM Director integration 187
- installation 173
- installing 165
- IP address 166
- Java runtime 174
- KVM 174
- MIB files 176
- MPCLI 193
- multi-node configuration 110
- network settings 166
- OS USB Selection 121
- Other OS setting 167
- partitioning 6
- Periodic SMI Generation 121
- ports 176
- remote control 174
- remote media 165
- SNMP MIB files 176
- static address 167
- TCP ports 176
- video speed 174
- Web interface 168
- Windows drivers 172

- rules
 - hot-add memory 100
 - memory 97
 - memory mirroring 98
 - multi-node configurations 112
 - processors 91
- RXE-100
 - support 13

S

- SAC commands 206
- SAP 71
- SAS 31
- scalability 9, 137
- scalability directory 27
- Scalability Enablement/Memory Adapter 119
- Scalable Partition Number 113

- Scalable Systems Manger 163
- scale-up 42
- scale-up versus scale-out 82
- scan order, PCI 35
- scrubbing, memory 29
- SEL 180
- Serial Attached SCSI 31
- Serial over LAN 194
 - SMBridge 204
- serial port 35
- Server Configuration Manager 188
- server consolidation 54, 56
- ServeRAID-6M 126
- ServeRAID-8i 33, 123
- ServerProven 130, 136
- shopping cart analogy 10
- Siebel 75
- SMBridge 194
 - authentication 203
 - BIOS settings 196
 - bootcfg command 208
 - CLI 195, 217
 - commands 204, 218
 - connecting 203
 - console command 204
 - console redirection 196
 - daemon 200
 - EMS 201, 206
 - EMS, enabling 208
 - event log 206
 - exiting SOL 207
 - F1 key in telnet 202
 - GRUB 214
 - installation 197
 - LILO 212
 - Linux support 196
 - ports 198
 - power control 205
 - PXE boot 196
 - Red Hat, enabling SOL 212
 - remote console redirection 196
 - SAC 206
 - SAC commands 206
 - SAC support 201
 - Serial over LAN 204
 - service 200
 - telnet clients 201
 - telnet server 194, 201
 - timeout values 199

- Windows SOL 206
- SMP 9
- snoop filter 27
- Software Update Subscription
 - Datacenter 150–151
- Solution Assurance 131
- Special Administration Console 201
- Special Mode 93
- SQL Server 79
- SRAT table 11
- SSE 24
- stripe sizes 123
- Supply Chain Management 53
- support
 - VMware, applications running on 157
- SUSE LINUX
 - drivers 154
 - installation 154
 - RSA II drivers 173
 - RSA II information 155
 - SOL, enabling 215
 - support 136
- system boards 3

T

- target applications 52
- TCP/IP ports
 - BMC 186
 - RSA II 176
- Total Cost of Ownership 61
- TotalStorage DS300/400 126
- Tower To Rack Conversion Kit 14
- two-way interleaving 9

U

- Ultra320 SCSI 31, 126
- upward integration 162
- USB diskette drive 141
- USB ports 35

V

- video PCI adapters 35
- VMware ESX Server 155
 - BIOS settings 158
 - installation 159
 - Microsoft support 157
 - pre-install 158

support 136
VRMs 19, 91

W

Web interface 168

Windows

- bootcfg command 208
- EMS 201, 208
- Hyper-Threading, affect of 23
- RSA II drivers 172
- SAC 201, 206
- SMBridge
 - installation 198
 - support for 196
- SOL support 206

Windows 2000 Server

- drivers 152
- HAL 153
- installation 152

Windows Datacenter 147

- Datacenter High Availability Program Offering 148
- Datacenter Scalable Offering 150
- x460 models 149, 151

Windows Server 2003

- BIOS settings 146
- boot from internal drives 146
- drivers 145
- installation 145
- PAE switch 147
- support 142

X

x255 comparison 50

x260

- drive bays 33
- front view 17
- memory cards 28
- memory configuration 96
- OS support 136
- power supplies 127
- processor upgrades 90
- processors 18
- rack models 14
- rear view 18
- removable media 125
- target applications 54
- tower models 14

tower-to-rack conversion 14

x255 comparions 50

X3 Architecture 1, 42

x365 comparison 46

x366

- drive bays 33
- dual-core upgrade 19
- front view 16
- memory cards 28
- memory configuration 96
- models 13
- OS support 136
- power supplies 127
- processor upgrades 90
- processors 18
- rear view 17
- removable media 125
- target applications 55
- x365 comparison 46

x445 comparison 48

x460 66

- drive bays 33
- dual-core upgrade 19
- export regulations 118
- front view 15
- memory cards 28
- memory configuration 96
- models 12, 149, 151
- OS support 136
- power supplies 127
- processor upgrades 90
- processors 18
- rear view 16
- removable media 125
- Scalability Enablement/Memory Adapter 119
- target applications 55
- Windows Datacenter models 149, 151
- x445 comparison 48

XA-64e chipset 8

XceL4v 27, 43

- memory use 141

Xeon 7020/7040 20

Xeon MP 18

XpandOnDemand 43, 64, 104



Planning and Installing the IBM @server X3 Architecture Servers

(0.2" spine)
0.17" <-> 0.473"
90 <-> 249 pages



Redbooks

Planning and Installing the IBM *e*server X3 Architecture Servers

**Covers the xSeries
460, xSeries 366, and
xSeries 260**

**Learn the technical
details of these
high-performance
servers**

**How to configure,
install, and manage
the servers**

The IBM *e*server X3 Architecture servers are the new third-generation Enterprise X-Architecture servers from IBM. The xSeries 366 was announced in February 2005, the xSeries 460 was announced in May, and the xSeries 260 was announced in August. These X3 Architecture servers are ideal for random commercial workloads with high processor and memory bandwidth requirements.

Although aimed at different client sets, they share common components and options. The x460 is targeted at clients who need to implement a large single image “scale-up” configuration of up to 32 processors and 512 GB of RAM. The x366 is targeted at the high-performance environment where rack space is a premium. The x260 is for clients who need high performance but also require large amounts of internal disk storage.

This IBM Redbook provides a detailed technical description of the three servers and explains how to plan, install, configure, and manage these high-performance servers running 32-bit and 64-bit versions of Windows Server 2003, Red Hat Enterprise Linux, SUSE Linux Enterprise Server, and VMware ESX Server.

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**