IBM

IBM BladeCenter Products and Technology

- Describes the BladeCenter chassis and blade server technology
- Provides details of available I/O modules and expansion options
- Explains networking and storage configurations

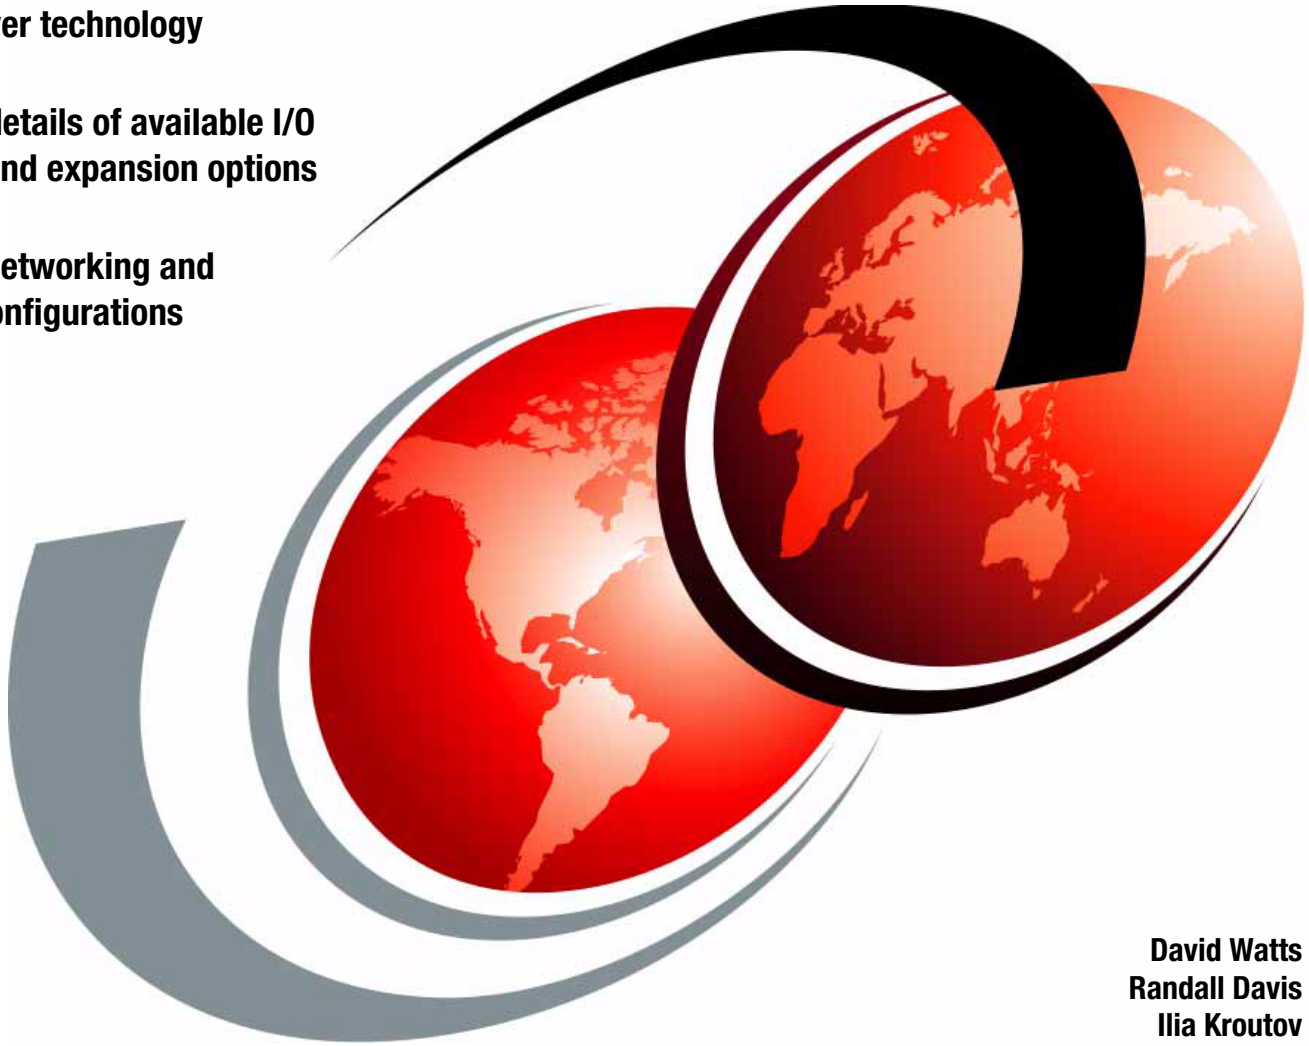

David Watts
Randall Davis
Ilia Kroutov

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM BladeCenter Products and Technology

March 2009

SG24-7523-03

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

Fourth Edition (March 2009)

This edition applies to the following products:

- IBM BladeCenter E, 8677
- IBM BladeCenter H, 8852
- IBM BladeCenter T, 8720 and 8730
- IBM BladeCenter HT, 8740 and 8750
- IBM BladeCenter S, 8886
- IBM BladeCenter HC10, 7996
- IBM BladeCenter HS12, 8014, 8028
- IBM BladeCenter HS20, 8843
- IBM BladeCenter HS21, 8853
- IBM BladeCenter HS21 XM, 7995
- IBM BladeCenter LS21, 7971
- IBM BladeCenter LS22, 7901
- IBM BladeCenter LS41, 7972
- IBM BladeCenter LS42, 7902
- IBM BladeCenter JS12, 7998-60X
- IBM BladeCenter JS21, 8844
- IBM BladeCenter JS22, 7998
- IBM BladeCenter QS21, 0792
- IBM BladeCenter QS22, 0793
- IBM BladeCenter PN41, 3020

Contents

Notices

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate 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.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

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

1350™
AIX 5L™
AIX®
AS/400®
BladeCenter®
Calibrated Vectored Cooling™
DPI®
DS4000®
DS6000™
DS8000®
EnergyScale™
Enterprise Workload Manager™
eServer™
FlashCopy®
i5/OS®
IBM PowerXCell™
IBM®
iSeries®
Micro-Partitioning™
Netfinity®
Power Systems™
POWER6™
PowerExecutive™
PowerPC®
PowerVM™
POWER®
pSeries®
Redbooks®
Redbooks (logo) ®
RETAIN®
ServerProven®
Solid®
System i®
System p®
System Storage™
System x®
Tivoli®
TotalStorage®
WebSphere®
xSeries®

The following terms are trademarks of other companies:

Advanced Micro Devices, AMD, AMD Opteron, ATI, Catalyst, ES1000, HyperTransport, Radeon, the AMD Arrow logo, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

InfiniBand, and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.

Novell, SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

Enterprise Fabric Suite, QLogic, QuickTools, SANbox, SANsurfer, and the QLogic logo are registered trademarks of QLogic Corporation. SANblade is a registered trademark in the United States.

Cell Broadband Engine and Cell/B.E. are trademarks of Sony Computer Entertainment, Inc., in the United States, other countries, or both and is used under license therefrom.

PC-over-IP, and the Teradici logo are trademarks or registered trademarks of Teradici Corporation in the United States, other countries, or both.

VMware, the VMware "boxes" logo and design are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

Brocade, Fabric OS, SilkWorm, and the Brocade logo are trademarks or registered trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries.

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

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

Celeron, Intel Celeron, Intel Core, Intel Xeon, Intel, 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, or service names may be trademarks or service marks of others.

Preface

IBM® BladeCenter® remains an innovative solution to running business solutions. IBM BladeCenter builds on the IBM commitment to integrating server, storage, and networking functionality with technology exchange and heterogeneous management. IBM BladeCenter offers the ease, density, availability, affordability, and scalability that are central to the blade technology promise.

Blade servers have captured industry focus because of their modular design, which can reduce cost with a more efficient use of valuable floor space, and its simplified management, which can help to speed up such tasks as deploying, reprovisioning, updating, and troubleshooting hundreds of blade servers. In addition, blade servers provide improved performance by doubling current rack density. By integrating resources and sharing key components, not only will costs be reduced but also availability will be increased.

This IBM Redbooks® publication describes IBM BladeCenter and discusses the technology and features of the different chassis, blade server models, and connectivity options. We go into details about every major component and provide guidance as to networking and storage connectivity.

This fourth edition includes new products that were announced in the second half of 2008.

The team that wrote this book

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

David Watts is a Consulting IT Specialist at the IBM ITSO Center in Raleigh. He manages residencies and produces IBM Redbooks publications on hardware and software topics related to IBM System x® and BladeCenter servers and associated client platforms. He has authored over 80 books, papers, and technotes. He holds a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM both in the US and Australia since 1989. David is an IBM Certified IT Specialist.

Randall Davis is a Senior IT Specialist working in the System x pre-sales team for IBM Australia as a Field Technical Sales Support (FTSS) specialist. He regularly performs System x, BladeCenter and Storage demonstrations for customers at the IBM Demonstration Centre in Melbourne, Australia. He also helps instruct Business Partners and customers on how to configure and install the BladeCenter. His areas of expertise are the IBM BladeCenter, System x servers, VMware®, and Linux®. Randall started at IBM as a System 36 and AS/400® Engineer in 1989

Ilia Kroutov is an Advisory IT Specialist with IBM Russia and has been with IBM since 1998. He is currently the Brand Manager for IBM System x and BladeCenter offerings for Russia and CIS countries. Prior to this, he was a Field Technical Sales Support (FTSS) specialist for System x products with a major focus on IBM BladeCenter, and an instructor at IBM Learning Services Russia, conducting training classes related to Netfinity® and System x hardware, Microsoft® server operating systems, and Cisco networking equipment topics. He has IBM Certified Advanced Technical Expert - IBM BladeCenter certification and certifications from Cisco Systems (CCNP, CCSP, CCDA). He graduated from the Moscow Engineering and

Physics Institute, and holds a Master's degree in Computer Engineering. Ilia is an Accredited IT Specialist.

The team (l–r): Randall, Ilia, David

The authors of the third edition of this book were:

- David Watts
- Randall Davis
- Ilia Kroutov
- Kevin Galloway

The authors of the second edition of this book were:

- David Watts
- Randall Davis
- Ilia Kroutov

The authors of the first edition of this book were:

- David Watts
- Randall Davis
- Ilia Kroutov

Thanks to the authors of the IBM Redpaper of the same name that was the predecessor to this book:

- David Watts
- Randall Davis
- Ilia Kroutov
- Mollie Tucker

Thanks to the authors of the first edition of the Solution Assurance Product Review Guide upon which this publication is based:

- Rufus Credle
- Yoko Fujiwara
- Marc Muehlhoff

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

ITSO:

- Tamikia Barrow
- Rufus Credle
- Linda Robinson

- Margaret Ticknor

BladeCenter and System x Marketing:

- Mark Chapman
- John Fitzgerald
- Kyle Hampton
- Hal Humrickhouse
- Harsh Kachhy
- Botond Kiss
- Richard Mancini
- Shekhar Mishra
- Ishan Sehgal
- David Tareen
- Scott Tease

BladeCenter Development:

- Alfredo Aldereguia
- Ganesh Balakrishnan
- Wayne Banks
- Brian Baker
- Paul Bashor
- Matt Belanger
- Dave Bentlage
- Jon Bitner
- Patrick Caporale
- Charles Clifton
- Denny Cole
- Matthew Eckl
- Lesley Estroff
- Nathan Flowers
- Dottie Gardner
- Parker Grannis
- Sam Graver
- Tim Green
- Darren Gross
- Jeff Hamilton
- Bill Holland
- Dick Hunter
- Andy Huryn
- Robert Jakes
- Brooks Johnston
- Bret Lehman
- Seth Lewis
- Ryan McArthur
- Mark McColgan
- Robyn McGlotten
- Caroline Metry
- Rob Ord
- Douglas Pase
- Dina Requena
- Ann Richter
- Jonri Rothwell
- Craig Rose
- Victor Stankevich
- Tung Tran
- Jimmy Thomas
- Steve Vanderlinden
- Doug Vassello
- Mark Welch
- Jack Wong
- Jeff Werth
- Wilson Velez

Technical support:

- Sharon Alonzo
- Khalid Ansari
- Bill Champion
- JJ Mangum
- John Sixsmith

IBM employees from around the world:

- Andreas Grottel, System Storage™ Product Manager, Germany
- Hal Jennings, IBM System p® Executive Briefing Center, Austin TX
- Fabiano Matassa, IT Specialist, UK
- Olaf Menke, IBM Pre-Sales Support, Germany
- Massimo Re Ferre', IBM Systems Group Technical Sales, Italy

Other people offering their generous support:

- Barbara Craig, QLogic® Corporation
- Matt Slavin, Blade Network Technologies
- Doug Toth, ITX Solutions

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will 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 books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks in one of the following ways:

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

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Summary of changes

This section describes the technical changes made in this edition of the book and in the IBM Redpaper that this book is built on. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for SG24-7523-03
for IBM BladeCenter Products and Technology
as created or updated on March 27, 2009.

March 2009, Fourth Edition

This edition has the following new and changed information.

New products and information

- BladeCenter LS22, page 44 and page 313
- BladeCenter LS42, page 45 and page 319
- BladeCenter PN41, page 53 and page 353
- New models of the BladeCenter HS12, page 264
- New models of the BladeCenter HS21, page 289
- New models of the BladeCenter HS21 XM, page 276
- New models of the BladeCenter QS22, page 350
- QLogic 20-Port 8 Gb SAN Switch Module, page 158
- QLogic 8 Gb Intelligent Pass-thru Module, page 161
- SAS RAID Controller Module, page 166
- Near Line SAS drives, page 211
- 2/4 Port Ethernet Expansion Card (CFFh), page 242
- Broadcom 2-port 10 Gb Ethernet Expansion Card (CFFh), page 244
- Broadcom 4-port 10 Gb Ethernet Expansion Card (CFFh), page 245
- QLogic Ethernet and 8Gb Fibre Channel Expansion Card (CFFh), page 247
- ServeRAID MR10ie (CIOv) Controller, page 255
- Support for the ServeRAID MR10ie in the HS12, page “Storage controllers” on page 269
- 300 GB SATA drive in HS12, page 270
- 300 GB SATA drive in HS21 (page 293) and HS21 XM (page 285)
- 16 GB (2x 8GB DIMMs) memory option for HS21 (page 292) and HS21 XM (page 279)
- Low power memory options for HS21 (page 289) and HS21 XM (page 276)
- EXP3000 attachment using ServeRAID MR10ie, page 420
- BladeCenter S storage configurations, covering the integrated Disk Storage Modules, EXP3000 enclosures, SAS I/O modules, and SAS expansion cards including the ServeRAID MR10ie, page 423

- Operating system support matrix, page 6

Updates to existing products and information

- Added feature codes to relevant options throughout the book. Clarified the feature codes of the I/O modules, page 3
- Updated I/O topology diagrams for each chassis, page 58
- BladeCenter HT InterSwitch Links, page 67
- Clarified the use of 4 Gbps and 8 Gbps Fibre Channel modules in BladeCenter T (page 67) and BladeCenter HT (page 69)
- Updated information about the MSIM for BladeCenter HT, page 178
- QS21 and QS22 both support the 4X InfiniBand® DDR CFFh card, part number 43W4423
- Clarified the local storage options and boot options for the QS22, page 352
- The QS22 supports the PEU3e expansion blade, page 353
- Updated operating system support of the BladeCenter Boot Disk System, page 410
- Added and corrected information about the DS3200, page 413
- Updated information about SAS tape storage, page 416

August 2008, Third Edition

This edition has the following new and changed information.

New products and information

- BladeCenter H - new 4YU model, page 19
- BladeCenter S - new 1NG model, page 33
- BladeCenter HS12, page 35 and 261
- BladeCenter JS12, page 40 and 329
- BladeCenter QS22, page 52 and 348
- Cisco Catalyst™ Switch Module 3110G, page 89
- Cisco Catalyst Switch Module 3110X, page 97
- Cisco Catalyst Switch Module 3012, page 106
- Nortel 1/10 Gb Uplink Ethernet Switch Module, page 130
- Intelligent Copper Pass-thru Module, page 171
- Multi-switch Interconnect Module for BladeCenter HT, page 178
- SAS Connectivity Card, page 254
- BladeCenter HC10 - new models, page 256
- BladeCenter HS21 XM - new models, page 276
- BladeCenter HS21 - new models, page 289
- PCI Express I/O Expansion Unit, page 368
- Cisco Virtual Blade Switch, page 384
- IBM System Storage DS3200, page 413
- SAS tape storage, page 432

Updates to existing products and information

- Added feature codes for options, where possible
- Removed details about HS20, LS20, HS40 models (except long-life models)
- BladeCenter Open Fabric Manager (OFM) with Nortel/BNT Extensions, page 115
- Connectivity options for the SAS Connectivity Module, page 164
- The Cisco 4X InfiniBand DDR Expansion Card has been withdrawn, page 251
- Updated information about the SAS Expansion Card, page 252
- Models of the HC10 with 60 GB drives have been withdrawn, page 256
- Clarified the use of the MPE in an LS41 when no processors are installed in the MPE, page 309
- Updated BladeCenter Open Fabric Manager information, page 387
- Power management information, page 218

February 2008, Second Edition

New products and information

- BladeCenter S chassis, page 30
- BladeCenter JS22 server, page 41 and 340
- BladeCenter QS21 server, page 51 and 345
- I/O paths for BladeCenter S, page 72
- SAS Connectivity Module, page 164
- 4X InfiniBand Pass-thru Module, page 182
- Fan module for BladeCenter S, page 195
- Power modules for BladeCenter S, page 204
- BladeCenter S Serial Pass-thru Module, page 207
- BladeCenter S Office Enablement Kit, page 209
- BladeCenter S storage modules and supported drives, page 211
- BladeCenter S physical site plans, page 213
- Emulex 4 GB Fibre Channel Expansion Card, page 250
- InfiniBand DDR Host Channel Adapters, page 251
- SAS Expansion Card, page 252
- New models and new processor options of the HS21 XM, page 264
- New 31.4 GB 2.5" Solid® State Drive for HS21 XM, HS21, LS21 and LS41, page 272
- New 8 GB Modular Flash Drive for HS21 XM, page 284
- New models and new processor options of the HS21, page 289
- New models and new processor options of the LS21, page 301
- New memory option for the LS21 and LS41, page 303 and page 310
- New models and new processor options of the LS41, page 307
- BladeCenter Open Fabric Manager, page 386
- IBM BladeCenter Boot Disk System, page 410

- IBM BladeCenter S integrated storage, page 423

Updates to existing products and information

- Added BladeCenter S, HC10, JS22, QS21 to the support matrixes, page 2
- JS21 and JS22 are now supported in the BladeCenter S chassis, page 2
- The 4X InfiniBand Switch Module is not supported in BladeCenter HT, page 3 and 138
- Added BladeCenter HC10 to the Product Overview, page 34
- Updated the I/O module support matrixes in each chassis, page 58
- Added BladeCenter S to management modules, page 184
- Added HC10, QS21 and JS22 to the Serial Port Breakout Cable, page 205
- Rearranged the section on Expansion Cards, page 234
- Added HC10, QS21 and JS22 to the expansion card matrix, page 235
- Added support of the 15.8 GB 2.5” solid state drive to the HS21, LS21 and LS41, page 294
- Updated the support matrix for withdrawn expansion blades, page 361
- Added more information about external storage options, page 396
- Added information about boot from iSCSI, iSCSI BladeBoot, page 406
- Added information about boot from SAN, page 409

August 2007, First Edition

New products and information

- BladeCenter HT chassis, page 25
- Nortel 10 Gigabit Ethernet High Speed Switch Module, page 134
- QLogic InfiniBand Ethernet Bridge Module, page 141
- QLogic InfiniBand Fibre Channel Bridge Module, page 143
- Cisco Systems 4 Gb 20-port and 10-port Fibre Channel Switch Modules, page 149
- QLogic 20-port and 10-port 4 Gb SAN Switch Modules, page 153
- QLogic 4 Gb Intelligent Pass-thru Module, page 156
- Multi Switch Interconnect Module, page 174
- Fan modules for BladeCenter HT, page 194
- Power modules for BladeCenter HT, page 202
- New section on the expansion card form factors, including the new CFFv and CFFh cards, page 234
- BladeCenter HS21 XM, page 261
- Solid State Drive options for the HS21 XM, page 271
- New models of various blade servers
- BladeCenter HC10, page 256
- New information about the use of InfiniBand and Ethernet using bridge modules, page 392
- New information about the use of InfiniBand and Fibre Channel SANs using a bridge module, page 432
- BladeCenter HT physical specifications and power considerations, page 213

Updates to existing products

- Clarifications and updates about the features of the Cisco IGESM and Cisco Fiber IGESM, pages 79 and 84.
- The Nortel Networks Layer 2-7 GbE Switch supports 1024 VLANs not 128, page 123.
- The Cisco 4X InfiniBand Switch Module internal ports run at 4X speeds with is 10 Gbps, not 2.5 Gbps.
- When used in a BladeCenter T chassis, the internal ports all 4 Gb Fibre Channel switch modules operate at 2 Gbps and the external ports operate at 4 Gbps. This applies to the Brocade (page 146), Cisco (page 149), QLogic (page 151), and McDATA switch modules.
- The Brocade 10-port 4 Gb Fibre Channel Switch Module now supports Dynamic Ports on Demand with the application of v5.0.5a of Brocade Fabric OS, page 146.
- The QLogic 10-port 4 Gb FC Switch module now supports the Dynamic Ports function with the use of firmware release 5.5.2.05, page 151.
- The McDATA 10-port 4 Gb FC Switch module now supports the Dynamic Ports function with the use of firmware release 5.5.2.05.
- Added text to further explain the use of the Copper Pass-thru Module, page 170.
- Added a table showing what servers support the BladeCenter H Serial Breakout Cable, page 205.
- The JS21 servers support the Cisco InfiniBand 4X Expansion Card, page 339.
- Updated the information about the cKVM card, page 356.
- Updated the information about expansion blades, page 358.
- Corrected the RAID levels that the ServeRAID 8k supports, page 365.
- The Cisco InfiniBand 1X Expansion Card is a standard form factor (StFF) card, not a small form factor (SFF) card.
- Updated Table 4-1 on page 372 to include selection criteria for new high speed switch modules.
- Updated the discussion on expansion blade units (MIO, SIO, and PEU2), page 358.
- Updated networking information related to 10 Gb Ethernet, page 371.
- Updated 5.3, “FC switch selection and fabric interoperability rules” on page 397, to include a discussion on the types of fabric and NPIV technology.
- Protected Mode feature of the Cisco IGESM, page 383.
- The JS21 now offers local keyboard and video console support with the June 26 GA3 firmware update, page 228.

1

Product overview

Blade servers are thin servers that insert into a single rack-mounted chassis which supplies shared power, cooling, and networking infrastructure. Each server is an independent server with its own processors, memory, storage, network controllers, operating system, and applications. Blade servers came to market around 2000, initially to meet clients' needs for greater ease of administration and increased server density in the data center environment.

When IBM released the IBM BladeCenter in November 2002, it quickly changed the industry with its modular design. The IBM BladeCenter provides complete redundancy in a chassis, and enables network and storage integration.

This chapter includes the following topics:

1.1 Support matrixes

This section provides compatibility matrixes:

1.1.1 Chassis-server compatibility

Table 1-1 lists the blade servers supported in each BladeCenter chassis.

Table 1-1 The blade servers supported in each BladeCenter chassis

Blade	Machine type	Blade width	BC S 8886	BC E 8677	BC T 8720	BC T 8730	BC H 8852	BC HT 8740	BC HT 8750
HS12	8014	1 slot	Yes	Yes[a]	Yes	Yes	Yes	Yes	Yes
HS12	8028	1 slot	Yes	Yes[a]	Yes	Yes	Yes	Yes	Yes
HS20	8678	1 slot	No	Yes	No	No	No	No	No
HS20	8832	1 slot	Yes	Yes	No	Yes	Yes	Yes	Yes
HS20	8843	1 slot	Yes	Yes	Yes[b]	Yes[b]	Yes	Yes	Yes
HS20	7981	1 slot	No	Yes	Yes	Yes	Yes	Yes	Yes
HS21	8853	1 slot	Yes	Yes[c]	Yes[d]	Yes[d]	Yes	Yes[e]	Yes[e]
HS21 XM	7995	1 slot	Yes	Yes[c]	Yes[d]	Yes[d]	Yes	Yes[c]	Yes[c]
HS40	8839	2 slots	No	Yes	No	Yes	No	No	No
LS20	8850	1 slot	Yes	Yes	Yes	Yes	Yes	Yes	Yes
LS21	7971	1 slot	Yes	Yes[f]	Yes[g]	Yes[g]	Yes	Yes	Yes
LS22	7901	1 slot	Yes	Yes[a,h]	No	No	Yes	No	No
LS41	7972	1 or 2 slots[i]	Yes	Yes[f]	Yes[g]	Yes[g]	Yes	Yes	Yes
LS42	7902	1 or 2 slots[i]	Yes	Yes[a,h]	No	No	Yes	No	No
JS12	7998-60X	1 slot	Yes	Yes[a]	Yes	Yes	Yes	Yes	Yes
JS20	8842	1 slot	No	Yes	Yes	Yes	Yes	Yes	Yes
JS21	8844	1 slot	Yes	Yes	Yes	Yes	Yes	Yes	Yes
JS22	7998 (excl 60X)	1 slot	Yes	No	No	No	Yes	Yes[j]	Yes[k]
QS21	0792	1 slot	Yes	No	No	No	Yes	Yes	Yes
QS22	0793	1 slot	Yes	No	No	No	Yes	Yes	Yes
PN41	3020	1 slot	No	No	No	No	Yes	Yes	Yes
HC10	7996	1 slot	Yes	Yes[a]	Yes	Yes	Yes	Yes	Yes

a. Supported with Advanced Management Module only.
b. Due to limited power availability some higher power models require installation at lower densities.
c. Models with quad-core processors are supported up to 80 W.

d. Up to eight blades with up to 65 W dual-core processors are supported. Up to seven blades with 80 W dual-core or quad-core processors are supported. Dual-core processors are supported up to 65W and quad-core processors are supported up to 50W.
e. Up to ten HS21 with 120 W quad-core processors are supported.
f. Versions of the LS21/41 that use the 95W processors (LS21 - 3AY, 6AY and LS41 - 3AY, 6AY, 3BY, 6BY) are designed for use in the BladeCenter H chassis. Due to increased power and thermal needs, these 95W versions will run at 'n-2' processor speed in the BladeCenter E. In most cases this represents a speed reduction of 400 MHz. The versions of the LS21/41 that use the High Efficiency (HE) 68W processors (LS21 - 31Y, 51Y and LS41 - 31Y, 51Y, 32Y, 52Y) will run at full speed in all chassis.
g. LS21/41 are supported up to 68W.
h. Only LS22/42 models with HE processors (55W ACP/79W TDP) are supported in BladeCenter E chassis.
i. The LS41/LS42 base unit requires one slot. With the Multi Processor Expansion (MPE) unit attached, the LS41/LS42 requires two slots.
j. Up to ten JS22s are supported in BC HT chassis.
k. Up to ten JS22s are supported in BC HT chassis.

Note: The support matrix for the expansion blades is shown in 3.18, “Expansion blades” on page 358.

1.1.2 Chassis-I/O module compatibility

Table 1-2 lists the I/O modules supported in each BladeCenter chassis.

Table 1-2 The I/O modules supported in each BladeCenter chassis

I/O module	Part Number	Feature code (BC-EHT/BC-S/ Power)[a]	BC S	BC E	BC T	BC H	BC HT	MSIM	MSIM-HT	Page
Ethernet switch modules										
Cisco Fiber Intelligent Gigabit ESM	32R1888	1498/none/none	N	Y	Y	Y	Y	Y	N	79
Cisco Intelligent Gigabit ESM	32R1892	1497/none/none	N	Y	Y	Y	Y	Y	Y	84
Cisco Catalyst Switch Module 3110G[b]	41Y8523	2989/none/3173	N	Y	Y	Y	Y	Y	N	89
Cisco Catalyst Switch Module 3110X[b]	41Y8522	2988/none/3171	N	Y	Y	Y	Y	Y	N	97
Cisco Catalyst Switch Module 3012[b]	43W4395	5450/3174/3174	Y	Y	Y	Y	Y	Y	N	106
IBM Server Connectivity Module[b]	39Y9324	1484/3220/3220	Y	Y	Y	Y	Y	Y	N	113
Nortel L2/3 Copper GbE Switch Module	32R1860	1495/3212/3212	Y	Y	Y	Y	Y	Y	Y	115
Nortel L2/3 Fiber GbE Switch Module	32R1861	1496/3213/3213	Y	Y	Y	Y	Y	Y	N	120
Nortel Networks Layer 2-7 Gb ESM	32R1859	1494/1494/3211	Y	Y	Y	Y	Y	N	N	123
Nortel Layer 2/3 10 Gigabit Uplink ESM[b]	32R1783	1493/3210/3210	Y	Y	Y	Y	Y	Y	N	126
Nortel 1/10 Gb Uplink ESM[b]	44W4404	1590/6980/1590	Y	Y	Y	Y	Y	Y	Y	130
Nortel 10 Gb High Speed Switch Module	39Y9267	2952/none/none	N	N	N	Y	Y	N	N	134
InfiniBand Modules										
Cisco Topspin InfiniBand Switch Module	32R1900	1533/none/none	N	Y	Y	Y	Y	N	N	None[c]
Cisco 4X InfiniBand Switch Module	32R1756	1574/none/3217	N	N	N	Y	Y	N	N	138
QLogic InfiniBand Ethernet Bridge Module	39Y9207	2941/none/none	N	N	N	Y	Y	N	N	141

I/O module	Part Number	Feature code (BC-EHT/BC-S/ Power)[a]	BC S	BC E	BC T	BC H	BC HT	MSIM	MSIM-HT	Page
QLogic InfiniBand FC Bridge Module	39Y9211	2942/none/none	N	N	N	Y	Y	N	N	143
Fibre Channel Switch Modules										
Brocade 4 Gb 20-port SAN Switch Module	32R1812	1569/none/3206	N	Y	Y	Y	Y	Y	N	146
Brocade 4 Gb 10-port SAN Switch Module	32R1813	1571/3207/3207	Y	Y	Y	Y	Y	Y	N	146
Cisco 4 Gb 20 port FC Switch Module[b]	39Y9280	2983/none/3242	N	Y	Y	Y	Y	Y	N	149
Cisco 4 Gb 10 port FC Switch Module[b]	39Y9284	2984/3241/3241	Y	Y	Y	Y	Y	Y	N	149
QLogic 4 Gb 20-port FC Switch Module	26R0881	1560/none/none	N	Y	Y	Y	Y	Y	N	151
QLogic 4 Gb 10-port FC Switch Module	32R1904	1575/none/none	N	Y	Y	Y	Y	Y	N	None[c]
QLogic 20-port 4 Gb SAN Switch Module[b]	43W6725	2987/3244/3244	N	Y	Y	Y	Y	Y	Y	153
QLogic 10-port 4 Gb SAN Switch Module[b]	43W6724	2986/3243/3243	Y	Y	Y	Y	Y	Y	Y	153
QLogic 4 Gb Intelligent Pass-thru Module[b]	43W6723	2985/3245/3245	Y	Y	Y	Y	Y	Y	N	156
McDATA 4 Gb 20-port FC Switch Module	32R1833	1562/none/none	N	Y	Y	Y	Y	Y	N	None[c]
McDATA 4 Gb 10-port FC Switch Module	32R1905	1576/none/none	N	Y	Y	Y	Y	Y	N	None[c]
QLogic 20-Port 8 Gb SAN Switch Module[b]	44X1905	5478/5866/3284	Y	Y	Y	Y	Y	Y	N	158
QLogic 8 Gb Intelligent Pass-thru Module[b]	44X1907	5482/5868/5449	Y	Y	Y	Y	Y	Y	N	161
SAS modules										
SAS Connectivity Module	39Y9195	2980/3267/3267	Y	Y	Y	Y	Y	N	N	164
SAS RAID Controller Module	43W3584	none/3734/none	Y	N	N	N	N	N	N	166
Pass-thru and interconnect modules										
IBM BladeCenter Optical Pass-thru Module	39Y9316	1556/3218/3218	Y	Y	Y	Y	Y	Y	N	173
IBM BladeCenter Copper Pass-thru Module	39Y9320	2900/3219/3219	Y	Y	Y	Y	Y	Y	N	170
Intelligent Copper Pass-thru Module[b]	44W4483	5452/5865/5452	Y	Y	Y	Y	Y	Y	N	171
QLogic 4 Gb Intelligent Pass-thru Module[b]	43W6723	2985/3245/3245	Y	Y	Y	Y	Y	Y	N	156
QLogic 8 Gb Intelligent Pass-thru Module[b]	44X1907	5482/5868/5449	Y	Y	Y	Y	Y	Y	N	161
IBM Multi-Switch Interconnect Module	39Y9314	1465/none/3239	N	N	N	Y	N	N	N	174
Multi-Switch Interconnect Module for BC HT	44R5913	5491/none/none	N	N	N	N	Y	N	N	178
4X InfiniBand Pass-thru Module	43W4419	2990/none/3216	N	N	N	Y	N	N	N	182

a. Feature codes are listed in the form of three groups separated by a slash mark (/). The first feature code is for BladeCenter E-, T-, H-, and HT-based configurations available through System x. The second feature code is for BladeCenter S-based configurations available through the System x sales channel. The third feature code is for BladeCenter S- and BladeCenter H-based configurations available through the Power Systems™ sales channel when applicable.

b. These switches require an Advanced Management Module to be installed in the BladeCenter E chassis.

c. None in the page column means that the module has been withdrawn from marketing and is not covered in this guide.

1.1.3 Server-expansion card compatibility

Table 1-3 lists the available expansion cards that can be installed in blade servers and in which specific servers they are a supported option. For details of an expansion card, see the page number referenced in the last column of the table.

Note: Table 1-3 lists which blade servers support a specific expansion card. It does not, however, take into account which chassis the blade is installed in. For details about chassis support, see the page reference.

For the latest support information, see one of the following resources:

- ServerProven®:

  http://www.ibm.com/servers/eserver/serverproven/compat/us/eserver.html
- Configuration and Option Guide

  http://www.ibm.com/support/docview.wss?rs=1201&uid=psg1SCOD-3ZVQ5W
- *BladeCenter Interoperability Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5073016

Table 1-3 The expansion cards supported in each blade server

Expansion cards (Y = supported, N = not supported)	Part	Feature code	HS12 — 8014, 8028	HS20 — 8843	HS21 — 8853	HS21 XM - 7995	LS21 — 7971	LS22 — 7901	LS41 — 7972	LS42 — 7902	JS12 — 7998-60X	JS20 — 8842	JS21 — 8844	JS22 — 7998	QS21 — 0792	QS22 — 0793	Page
Ethernet expansion cards																	
Gigabit Ethernet (StFF)	73P9030	1547	N	Y	Y	Y	Y	N	N	N	N	Y	Y	N	N	N	None[a]
Gigabit SFF Ethernet (SFF)	39R8624	1548	Y	Y	Y	Y	Y	Y	Y	Y	N	Y	Y	N	N	N	243
Ethernet Expansion Card (CFFv)	39Y9310	2969	Y	N	Y	Y	Y	Y	Y	Y	Y	N	Y	Y	N	N	242
QLogic Ethernet/4 Gb FC (CFFh)	39Y9306	2968	Y	N	Y	Y	Y	Y	Y	Y	Y	N	Y	Y	N	N	246
NetXen 10 Gb Ethernet (CFFh)	39Y9271	2967	Y	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N	243
2/4 Port Ethernet Exp Card (CFFh)	44W4479	5476	Y	N	Y	Y	Y	Y	Y	N	N	N	N	Y	N	N	242
Broadcom 10 Gb 4-pt Eth. (CFFh)	44W4465	5479	Y	N	Y	Y	Y	Y	N	N	N	N	N	N	N	N	245
Broadcom 10 Gb 2-pt Eth. (CFFh)	44W4465	5489	Y	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N	244
QLogic Eth. and 8 Gb FC (CFFh)	44X1940	5485	Y	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N	247
Fibre Channel expansion cards																	
BladeCenter FC (StFF)	13N2203	1519	N	Y	Y	N	Y	N	Y	N	N	Y	N	N	N	N	None[a]
BladeCenter 2 Gb FC (SFF)	26K4841	1549	N	Y	Y	Y	Y	N	Y	N	N	Y	Y	N	N	N	248
Emulex 4 Gb SFF FC	39Y9186	2925	N	Y	Y	Y	Y	N	Y	N	N	Y	Y	N	N	N	249
Emulex 4 Gb CFFv FC	43W6859	2994	Y	N	Y	Y	Y	Y	Y	Y	Y	N	N	Y	N	N	250
QLogic 4 Gb SFF FC	26R0890	1577	N	Y	Y	Y	Y	N	Y	N	N	Y	Y	N	N	N	249

Expansion cards (Y = supported, N = not supported)	Part	Feature code	HS12 — 8014, 8028	HS20 — 8843	HS21 — 8853	HS21 XM - 7995	LS21 — 7971	LS22 — 7901	LS41 — 7972	LS42 — 7902	JS12 — 7998-60X	JS20 — 8842	JS21 — 8844	JS22 — 7998	QS21 — 0792	QS22 — 0793	Page
QLogic 4 Gb StFF FC	26R0884	1597	N	Y	Y	Y	Y	N	Y	N	N	Y	Y	N	N	N	249
QLogic Ethernet/4 Gb FC (CFFh)	39Y9306	2968	Y	N	Y	Y	Y	Y	Y	Y	Y	N	Y	Y	N	N	246
QLogic 4 Gb FC (CFFv)	41Y8527	2970	Y	N	Y	Y	Y	Y	Y	Y	Y	N	Y	Y	N	N	248
QLogic Eth. and 8 Gb FC (CFFh)	44X1940	5485	Y	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N	247
InfiniBand expansion cards																	
Cisco InfiniBand 1X (StFF)	32R1896	1461	N	Y	Y	Y	Y	N	Y	N	N	Y	Y	N	N	N	None[a]
Cisco InfiniBand 4X (HSFF)	32R1760	1466	N	N	Y	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	None[a]
4X InfiniBand DDR (CFFh)	43W4423	2991	N	N	Y	Y	Y	Y	Y	Y	Y	N	N	Y	Y	Y	251
Cisco 4X InfiniBand DDR (CFFh)	43W4421	2992	N	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N	None[a]
Voltaire 4X InfiniBand DDR (CFFh)	43W4420	2993	N	N	Y	Y	Y	Y	Y	Y	Y	N	N	Y	N	N	251
Other expansion cards																	
Myrinet Cluster (StFF)	73P6000	1546	N	Y	Y	Y	Y	Y	Y	Y	N	Y	Y	N	N	N	None[a]
QLogic iSCSI (StFF)	32R1923	1458	N	Y	Y	Y	Y	N	Y	N	Y	Y	Y	Y	N	N	252
IBM Concurrent KVM (Custom)	26K5939	1568	N	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N	356
SAS Expansion Card (CFFv)	39Y9190	2979	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	252
SAS Connectivity Card (CFFv)	43W3974	1591	Y	N	N	Y	N	N	N	N	N	N	N	N	N	N	254
ServeRAID MR10ie[b] (CIOv)	46C7167	5752	Y	N	N	N	N	N	N	N	N	N	N	N	N	N	255

a. This expansion card has been withdrawn, so it is not covered in further detail in this guide.
b. This expansion card is only supported in the HS12 type 8028.

1.1.4 Supported operating systems

This section contains compatibility information for IBM blades and operating systems. Operating systems are grouped by vendor and then by operating system version. All operating systems that support the current product line of blade servers are listed in the following tables. The information represents a snapshot of IBM ServerProven at the time of writing. Visit the IBM Server Proven Web site for the latest information.

http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/ematrix.shtml

The information in Table 1-4 on page 7 through Table 1-10 on page 11 is for general purposes. Refer to the IBM ServerProven site for any additional information and limitations.

The operating systems covered here are:

- “Microsoft”
- “Novell NetWare” on page 8
- “Novell SUSE Linux” on page 9

Microsoft

Table 1-4 summarizes support for operating systems from Microsoft.

Table 1-4 Supported Microsoft operating systems

Operating system	HC10 7996	HS12 8014	HS12 8028	HS20 8843	HS21 XM 7995	HS21 8853	LS21 7971	LS22 7901	LS41 7972	LS42 7902	JS12 7998-60X	JS21 8844	JS22 7998	QS21 0792	QS22 0793	PN41 3020
Microsoft Windows® Server 2008 x64																
Web Edition	N	Y	Y	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
Standard Edition	N	Y	Y	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
Enterprise Edition	N	Y	Y	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
DataCenter Edition	N	N	N	N	Y	Y	N	Y	N	Y	N	N	N	N	N	N
Microsoft Windows Server® 2008 x86																
Web Edition	N	Y	Y	N	Y	Y	Y	N	Y	N	N	N	N	N	N	N
Standard Edition	N	Y	Y	N	Y	Y	Y	N	Y	N	N	N	N	N	N	N
Enterprise Edition	N	Y	Y	N	Y	Y	Y	N	Y	N	N	N	N	N	N	N
DataCenter Edition	N	N	N	N	Y	Y	N	N	N	N	N	N	N	N	N	N
Microsoft Windows Essentials Business Server 2008																
Standard Edition	N	N	N	N	N	N	N	Y	N	Y	N	N	N	N	N	N
Premium Edition	N	N	N	N	N	N	N	Y	N	Y	N	N	N	N	N	N
Microsoft Windows Compute Cluster Server 2003																
Compute Cluster Edition	N	Y	Y	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
Microsoft Windows Server 2003/2003 R2 x64																
Standard Edition	N	Y	Y	Y	Y	Y	Y	N	Y	N	N	N	N	N	N	N
Enterprise Edition	N	Y	Y	Y	Y	Y	Y	N	Y	N	N	N	N	N	N	N
Enterprise Edition with MSCS	N	N	N	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
Microsoft Windows Server 2003/2003 R2																
Web Edition	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
Standard Edition	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
Enterprise Edition	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
Enterprise Edition with MSCS	N	N	N	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N

Operating system	HC10 7996	HS12 8014	HS12 8028	HS20 8843	HS21 XM 7995	HS21 8853	LS21 7971	LS22 7901	LS41 7972	LS42 7902	JS12 7998-60X	JS21 8844	JS22 7998	QS21 0792	QS22 0793	PN41 3020
Microsoft Windows Storage Server 2003/2003 R2																
Standard Edition	N	N	N	Y	N	N	N	N	N	N	N	N	N	N	N	N
Windows Small Business Server 2003/2003 R2																
Standard Edition	N	Y	Y	N	Y	Y	Y	Y	N	Y	N	N	N	N	N	N
Premium Edition	N	Y	Y	N	Y	Y	Y	Y	N	Y	N	N	N	N	N	N
Microsoft Windows 2000																
Server	N	N	N	Y	N	Y	Y	N	Y	N	N	N	N	N	N	N
Advanced Server	N	N	N	Y	N	Y	Y	N	Y	N	N	N	N	N	N	N
Microsoft Windows XP																
XP Professional (x86 and x64)	Y	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Microsoft Windows Vista®																
Vista Business (x86 and x64)	Y	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Vista Enterprise (x86 and x64)	Y	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Vista Ultimate (x86 and x64)	Y	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N

Novell NetWare

Table 1-5 summarizes support for NetWare operating systems from Novell®.

Table 1-5 Supported Novell operating systems

Operating system	HC10 7996	HS12 8014	HS12 8028	HS20 8843	HS21 XM 7995	HS21 8853	LS21 7971	LS22 7901	LS41 7972	LS42 7902	JS12 7998-60X	JS21 8844	JS22 7998	QS21 0792	QS22 0793	PN41 3020
NetWare 6.5, Novell Cluster Services (NCS) v1.8	N	N	N	Y	N	N	N	N	N	N	N	N	N	N	N	N
NetWare 6.5	N	N	N	Y	Y	Y	Y	N	Y	N	N	N	N	N	N	N

Novell SUSE Linux

Table 1-6 summarizes support for SUSE® Linux operating systems from Novell.

Table 1-6 Supported SUSE Linux operating systems

Operating system	HC10 7996	HS12 8014	HS12 8028	HS20 8843	HS21 XM 7995	HS21 8853	LS21 7971	LS22 7901	LS41 7972	LS42 7902	JS12 7998-60X	JS21 8844	JS22 7998-61X	QS21 0792	QS22 0793	PN41 3020
SUSE Linux Enterprise Server 10																
Real Time for AMD64/EM64T	N	N	N	N	Y	N	Y	Y	N	N	N	N	N	N	N	N
with Xen for AMD64/EM64T	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
with Xen for x86	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
for AMD64/EM64T	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
for x86	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
for IBM POWER®	N	N	N	N	N	N	N	N	N	N	Y	N	Y	N	N	N
SUSE Linux Enterprise Server 9																
for AMD64/EM64T	N	Y	Y	Y	Y	Y	Y	N	Y	N	N	N	N	N	N	N
for x86	N	Y	Y	Y	Y	Y	Y	N	Y	N	N	N	N	N	N	N
for IBM POWER	N	N	N	N	N	N	N	N	N	N	N	Y	N	N	N	N
SUSE Linux Enterprise Server 8																
for x86	N	N	N	Y	N	N	N	N	N	N	N	N	N	N	N	N

Red Hat Linux

Table 1-7 summarizes support for Red Hat Linux operating systems from Red Hat.

Table 1-7 Supported Red Hat Linux operating systems

Operating system	HC10 7996	HS12 8014	HS12 8028	HS20 8843	HS21 XM 7995	HS21 8853	LS21 7971	LS22 7901	LS41 7972	LS42 7902	JS12 7998-60X	JS21 8844	JS22 7998	QS21 0792	QS22 0793	PN41 3020
Red Hat Enterprise Linux 5																
Server x64 Edition	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
Server with Xen x64 Edition	N	Y	Y	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
Server Edition	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
Server with Xen Edition	N	Y	Y	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
for System i® and System p	N	N	N	N	N	N	N	N	N	N	Y	N	Y	Y	Y	N
Red Hat Enterprise Linux 4																
AS for AMD64/EM64T	N	Y	Y	Y	Y	Y	Y	N	Y	N	N	N	N	N	N	N

Operating system	HC10 7996	HS12 8014	HS12 8028	HS20 8843	HS21 XM 7995	HS21 8853	LS21 7971	LS22 7901	LS41 7972	LS42 7902	JS12 7998-60X	JS21 8844	JS22 7998	QS21 0792	QS22 0793	PN41 3020
AS for iSeries® and pSeries®	N	N	N	N	N	N	N	N	N	N	Y	Y	Y	N	N	N
AS for x86	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
ES for AMD64/EM64T	N	Y	Y	Y	Y	Y	Y	N	Y	N	N	N	N	N	N	N
ES for x86	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
WS/HPC for AMD64/EM64T	N	Y	Y	Y	Y	Y	Y	N	Y	N	N	N	N	N	N	N
WS/HPC for x86	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
Red Hat Enterprise Linux 3																
AS for AMD64/EM64T	N	N	N	Y	N	Y	Y	N	Y	N	N	N	N	N	N	N
AS for x86	N	N	N	Y	N	Y	Y	N	Y	N	N	N	N	N	N	N
ES for AMD64/EM64T	N	N	N	Y	N	Y	Y	N	Y	N	N	N	N	N	N	N
ES for x86	N	N	N	Y	N	Y	Y	N	Y	N	N	N	N	N	N	N
WS for AMD64/EM64T	N	N	N	Y	N	Y	Y	N	Y	N	N	N	N	N	N	N
WS for x86	N	N	N	Y	N	Y	Y	N	Y	N	N	N	N	N	N	N
Red Hat Enterprise Linux 2.1																
AS for x86	N	N	N	Y	N	N	N	N	N	N	N	N	N	N	N	N

Sun Solaris

Table 1-8 summarizes support for Sun™ Solaris™ 10 operating systems from Sun.

Table 1-8 Supported Sun Solaris operating systems

Operating system	HC10 7996	HS12 8014	HS12 8028	HS20 8843	HS21 XM 7995	HS21 8853	LS21 7971	LS22 7901	LS41 7972	LS42 7902	JS12 7998-60X	JS21 8844	JS22 7998	QS21 0792	QS22 0793	PN41 3020
Solaris 10	N	N	N	Y	Y	Y	Y	N	Y	Y	N	N	N	N	N	N

VMware

Table 1-9 summarizes support for VMware operating systems from VMware.

Table 1-9 Supported VMware operating systems

Operating system	HC10 7996	HS12 8014	HS12 8028	HS20 8843	HS21 XM 7995	HS21 8853	LS21 7971	LS22 7901	LS41 7972	LS22 7902	JS12 7998-60X	JS21 8844	JS22 7998	QS21 0792	QS22 0793	PN41 3020
VMware ESX 3.5	N	N	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	N
VMware ESX Server 3.0	N	N	N	Y	Y	Y	Y	N	Y	N	N	N	N	N	N	N
VMware ESXi 3.5	N	N	Y	N	Y	Y	Y	N	Y	N	N	N	N	N	N	N
VMware ESX Server 2.5	N	N	N	Y	N	Y	Y	N	Y	N	N	N	N	N	N	N

IBM

Table 1-10 summarizes support for the AIX® and i5/OS® operating systems from IBM.

Table 1-10 Supported IBM AIX and i5/OS operating systems

Operating system	HC10 7996	HS12 8014	HS12 8028	HS20 8843	HS21 XM 7995	HS21 8853	LS21 7971	LS22 7901	LS41 7972	LS22 7902	JS12 7998-60X	JS21 8844	JS22 7998	QS21 0792	QS22 0793	PN41 3020
AIX Version 6.1	N	N	N	N	N	N	N	N	N	N	Y	N	Y	N	N	N
AIX 5L™ Version 5.3	N	N	N	N	N	N	N	N	N	N	Y	Y	Y	N	N	N
AIX 5L Version 5.2	N	N	N	N	N	N	N	N	N	N	N	Y	N	N	N	N
IBM i Operating System 6.1	N	N	N	N	N	N	N	N	N	N	Y	N	Y	N	N	N
IBM Virtual I/O Server 1.5	N	N	N	N	N	N	N	N	N	N	Y	N	Y	N	N	N

1.2 BladeCenter chassis

There are five chassis in the BladeCenter family:

- IBM BladeCenter E provides the greatest density and common fabric support and is the lowest entry cost option. See 1.2.1, “BladeCenter E” on page 12.
- IBM BladeCenter H delivers high performance, extreme reliability, and ultimate flexibility for the most demanding IT environments. See 1.2.2, “BladeCenter H” on page 15.
- IBM BladeCenter T models are designed specifically for telecommunications network infrastructures and other rugged environments. See 1.2.3, “BladeCenter T” on page 19.
- IBM BladeCenter HT models are designed for high-performance flexible telecommunications environments by supporting high-speed internet working technologies such as 10G Ethernet, and provide a robust platform for NGNs. See 1.2.4, “BladeCenter HT” on page 25.

- IBM BladeCenter S combines the power of blade servers with integrated storage, all in an easy-to-use package designed specifically for the office and distributed enterprise environment. See 1.2.5, “BladeCenter S” on page 30.

All five chassis share a common set of blades and standard switch modules, and additionally, BladeCenter H and HT offer high-speed I/O bays for high-speed switches, such as Nortel 10 Gb Ethernet High Speed Switch Module.

1.2.1 BladeCenter E

IBM designed the IBM BladeCenter E (machine type 8677) to be a highly modular chassis to accommodate a range of diverse business requirements. BladeCenter supports not only blade servers, but also a wide range of networking modules, including Gigabit Ethernet, Fibre Channel, and InfiniBand for high-speed connectivity to the client’s existing network environment. BladeCenter E also supports a redundant pair of Management Modules for comprehensive systems management.

Providing a wide selection of integrated switching options, BladeCenter systems help you lower the Total Cost of Ownership (TCO) by eliminating the need to purchase additional keyboards, videos, and mice (KVM), Ethernet and Fibre Channel switches, or the cumbersome and expensive cabling required by the switches. BladeCenter is a leader in the industry in providing flexibility and a variety of integration choices with components that fit into your infrastructure and deliver a comprehensive blade solution.

BladeCenter E’s superior density and feature set are made possible by the BladeCenter E innovative chassis architecture. Because BladeCenter E uses super energy-efficient components and shared infrastructure architecture, clients can realize lower power consumption when compared to their most likely alternative, that is, non-blade server designs. BladeCenter E’s lower power consumption and Calibrated Vectored Cooling™ allow more servers to fit in a tight power or cooling environment.

Figure 1-1 displays the front view of an IBM BladeCenter E.

Figure 1-1 BladeCenter E front view

Figure 1-2 on page 13 displays the front view showing key chassis features.

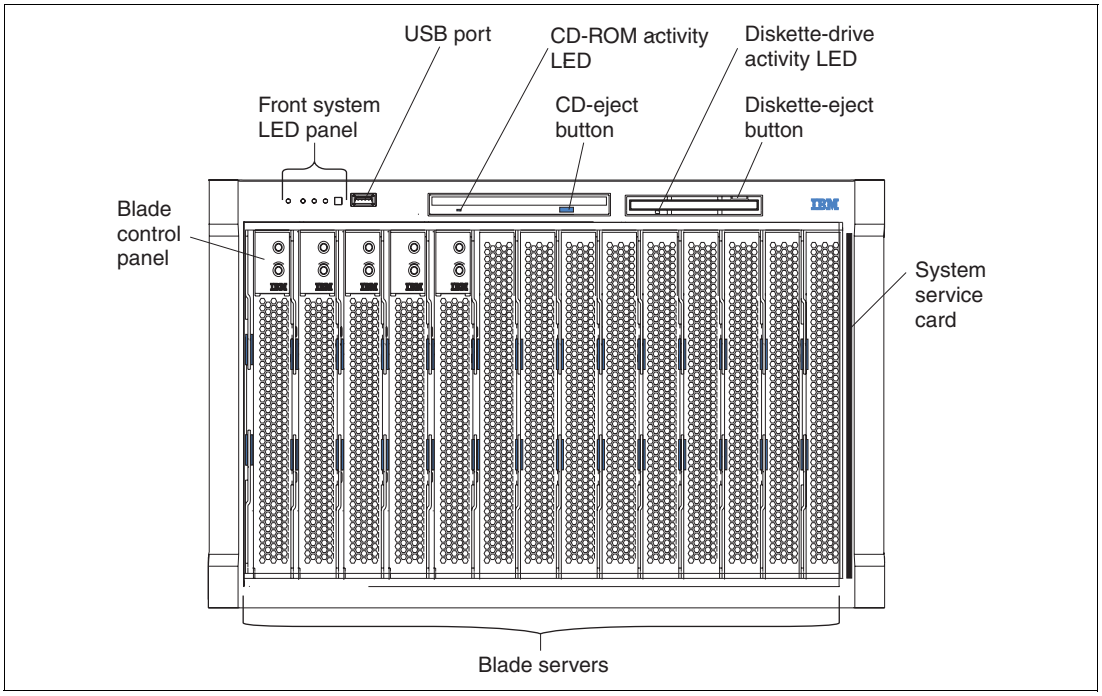

Figure 1-2 BladeCenter E front view showing the key features of the chassis

The following features are key features on the front of IBM BladeCenter E chassis:

- A media tray at the front with a diskette drive, a DVD drive (slim IDE), one USB V1.1 port, and a system-status panel
- 14 hot swap blade bays supporting different blade server models

Note: For a list of blade servers supported in the IBM BladeCenter E, refer to Table 1-1 on page 2.

Figure 1-3 displays the rear view of an IBM BladeCenter E.

Figure 1-3 BladeCenter E rear view

Figure 1-4 on page 14 displays the rear view showing key chassis features.

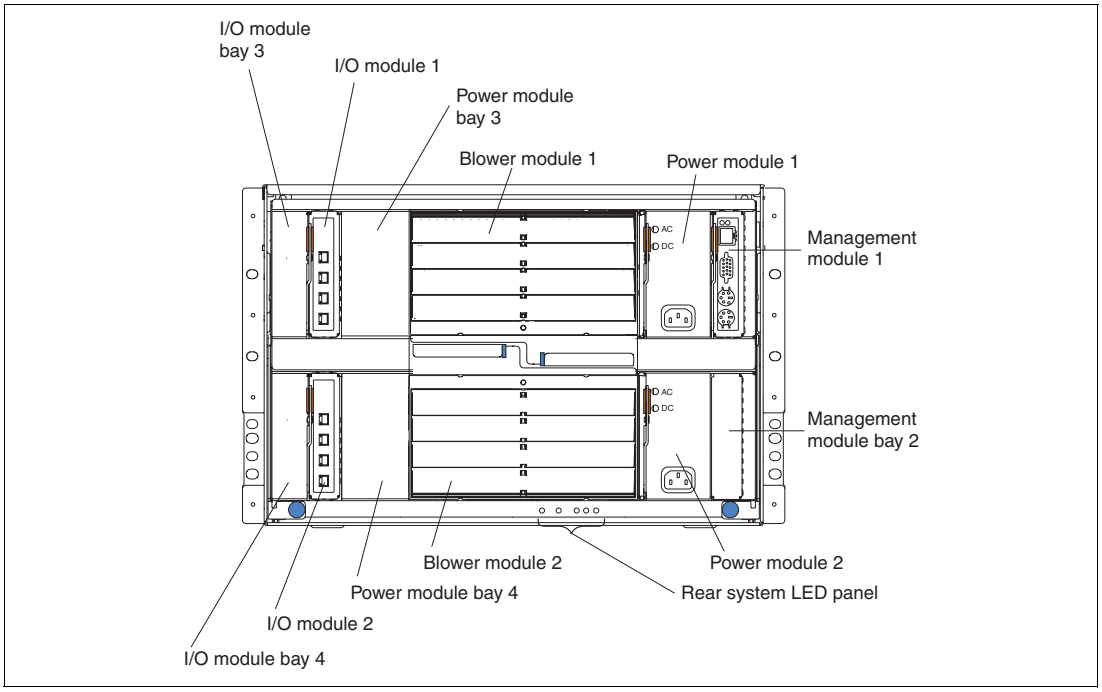

Figure 1-4 BladeCenter E rear view showing the key features of the chassis

The following features are key features on the rear of IBM BladeCenter E:

- Four hot swap I/O module bays
- Two hot swap Management Module bays - with one Management Module as standard
- Four hot swap power module bays - with two power modules as standard
- Two hot swap blowers

The BladeCenter E chassis allows for either 14 single-slot blade servers or seven double-slot blade servers. However, you can mix different blade server models in one chassis to meet your requirements, subject to power and cooling requirements, and some other configuration rules that we discuss in later chapters of this document.

The current model BladeCenter E chassis (8677-3Rx) ships standard with one Advanced Management Module. This module provides the ability to manage the chassis, as well as providing the local KVM function.

An optional redundant Advanced Management Module provides BladeCenter E with higher levels of resiliency. While in the chassis, the second module is in passive or standby mode. If the active or primary module fails, the second module is automatically enabled with all of the configuration settings of the primary module. This function provides clients with easy remote management and connectivity to the BladeCenter E chassis for their critical applications.

Note: The BladeCenter E chassis started shipping standard with the Advanced Management Module in 1Q06. Prior to 1Q06, it had shipped with the original Management Module.

Further details regarding the features and functions of the management modules can be found in 2.7, “Management modules” on page 184.

The BladeCenter E does not ship standard with any I/O modules. You need to choose these I/O modules depending on your connectivity needs. An Ethernet Switch Module (ESM) will be

required in I/O module bays 1 and 2, to enable the use of both Ethernet ports on a Blade Server. The I/O modules required in I/O module bays 3 and 4 will depend on the I/O Expansion Card installed in the Blade Servers.

We discuss guidelines about how to install I/O module options in detail in Chapter 2, “BladeCenter chassis and infrastructure configuration” on page 57.

The BladeCenter E chassis ships standard with:

- One Advanced Management Module
- Two hot swap power supply modules
- Two hot swap blower modules
- one USB v1.1 port
- One DVD-ROM drive
- One 1.44 MB diskette drive

Table 1-11 lists the major features of the IBM BladeCenter E.

Table 1-11 BladeCenter E features at a glance

Feature	Specification
Machine type	8677-3RU, 3RX
Rack form factor (H x D)	7U x 28 inches (711 mm)
DVD/CD drives standard	1x DVD-ROM (in Media Tray)
Diskette drives standard	1x 1.44 MB diskette drive (in Media Tray)
Number of blade server slots	14 (30 mm blade servers)
Number of switch module slots	4 hot-swap
Switch modules standard	None
Power supply size standard	2000 Watts AC
Number of power supplies (standard/maximum)	2/4
Number of blowers (standard/maximum)	2/2
Dimensions	Height: 12.0” (305 mm) Width: 17.5” (4429 mm) Depth: 28.0” (711 mm)

1.2.2 BladeCenter H

IBM BladeCenter H delivers high performance, extreme reliability, and ultimate flexibility to even the most demanding IT environments. In 9 U of rack space, the BladeCenter H chassis can contain up to 14 blade servers, 10 switch modules, and four power supplies to provide the necessary I/O network switching, power, cooling, and control panel information to support the individual servers.

The chassis supports up to four traditional fabrics using networking switches, storage switches, or pass through devices. The chassis also supports up to four high-speed fabrics for support of protocols like 4X InfiniBand or 10 Gigabit Ethernet. The built-in media tray includes light path diagnostics, two front USB inputs, and a optical drive.

Figure 1-5 displays the front view of an IBM BladeCenter H.

Figure 1-5 BladeCenter H front view

Figure 1-6 displays the front view showing key chassis features.

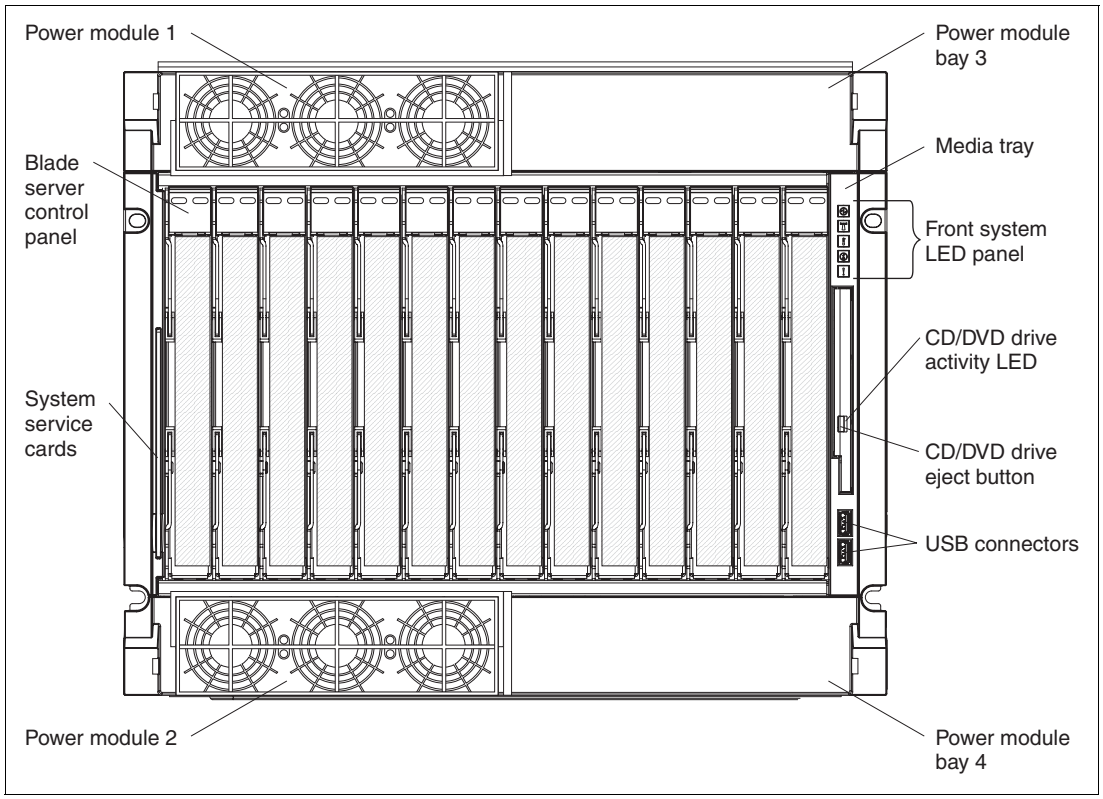

Figure 1-6 BladeCenter H front view showing the key features of the chassis

The following features are the key features on the front of the BladeCenter H:

- A media tray at the front right, with an optical drive, two USB v2.0 ports, and a system status LED panel.
- One pair of 2,900-watt power modules. An additional power module option (containing two 2,900 W power modules) is available.
- Two hot swap fan modules (two extra hot swap fan modules are included with the additional power module option).
- 14 hot swap blade server bays supporting different blade server types

> **Note:** For a list of blade servers supported in the IBM BladeCenter H, see Table 1-1 on page 2.

Figure 1-7 displays the rear view of an IBM BladeCenter H.

Figure 1-7 BladeCenter H rear view

Figure 1-7 displays the rear view showing key chassis features.

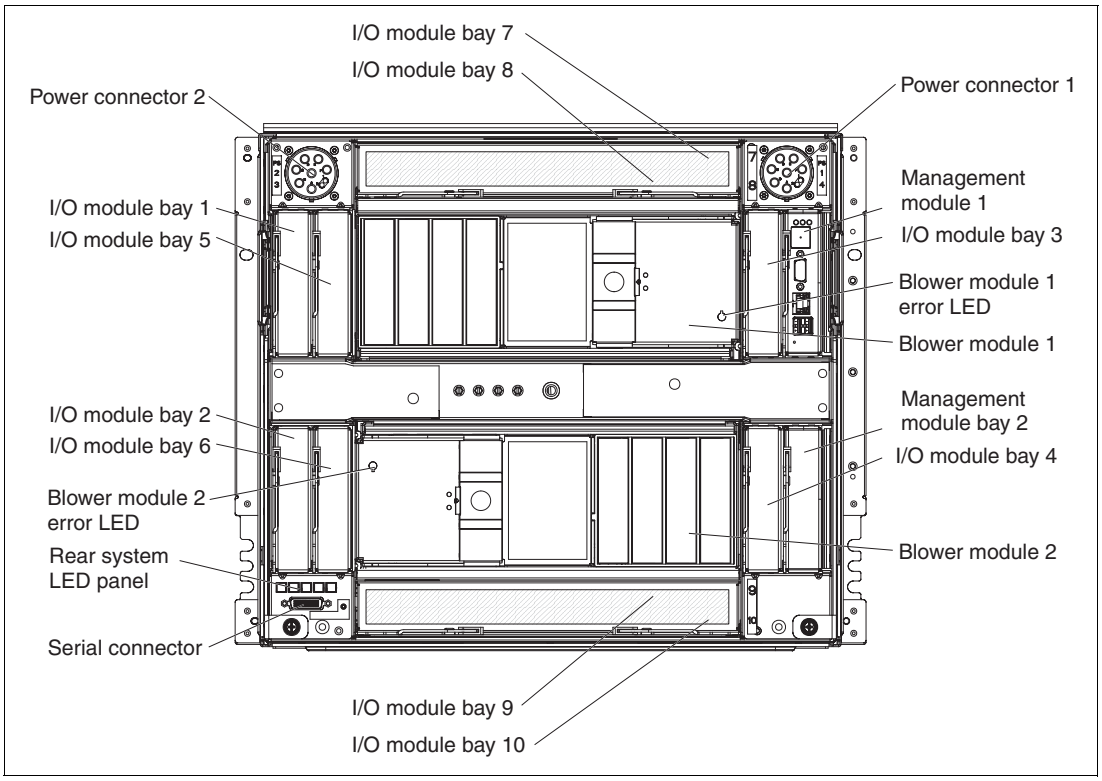

Figure 1-8 BladeCenter H rear view showing the key features of the chassis

The following features are the key features on the rear of the BladeCenter H:

- Two hot-swap blower modules as standard
- Two hot swap management-module bays - with one management module as standard
- Four traditional fabric switch modules
- Four high speed fabric switch modules
- Serial port breakout connector to give direct serial connection to installed blades (for those blades with the functionality)

The BladeCenter H chassis allows for either 14 single-slot blade servers or seven double-slot blade servers. However, you can mix different blade server models in one chassis to meet your requirements.

The BladeCenter H chassis ships standard with one Advanced Management Module. This module provides the ability to manage the chassis, as well as providing the local KVM function.

The optional redundant Advanced Management Module provides the IBM BladeCenter H with higher levels of resiliency. While in the chassis, the second module is in passive or standby mode. If the active or primary module fails, the second module is automatically enabled with all of the configuration settings of the primary module. This function provides clients with easy remote management and connectivity to the BladeCenter H chassis for their critical applications.

The BladeCenter H does not ship standard with any I/O modules. You choose these I/O modules based on your connectivity needs. An Ethernet Switch Module (ESM) is required in I/O module bays 1 and 2, to enable the use of both Ethernet ports on a Blade Server. The I/O modules required in I/O module bays 3 and 4 will depend on the I/O Expansion Card installed

in the blade servers. The I/O modules required in the high speed I/O module bays 7, 8, 9, and 10 will depend on the I/O Host Channel Adapter cards installed in the Blade Servers.

Guidelines about how to install I/O module options are discussed in detail in Chapter 2, “BladeCenter chassis and infrastructure configuration” on page 57.

The BladeCenter H chassis requires 200 to 240 V AC power.

Note: The BladeCenter H chassis comes standard with no power cord. Power cords need to be ordered separately.

The BladeCenter H chassis ships standard with:

- One advanced management module
- Two blower modules
- Two power supply modules (one pair of 2,900-watt power modules).
- Two hot swap power supply fan modules.
- Two USB v2.0 ports
- One optical drive (DVD-ROM or multi-burner depending on the mode)

The chassis does not have a diskette drive. An optional USB-attached 1.44 MB diskette drive is available.

Table 1-12 lists the major features of the IBM BladeCenter H.

Table 1-12 BladeCenter H features at a glance

Feature	Specifications	
Machine type	8852-4XU	8852-4YU
Rack form factor (H x D)	9U x 28 inches (711 mm)	
DVD/CD drives standard (in Media Tray)	1x DVD-ROM	1x UltraSlim Enhanced Multi-Burner Drive
Diskette drives standard	None	
Number of Blade Server slots	14 (30 mm blade servers)	
Number of switch module slots	10 hot-swap (4 x high speed, 4 x standard, 2 x bridge)	
Switch modules standard	None (in standard chassis offerings)	
Power supply size standard	2900W AC	
Number of power supplies (standard/maximum)	2/4	
Number of blowers (standard/maximum)	2/2	
Dimensions	Height: 15.75” (400 mm) Width: 17.4” (442 mm) Depth: 28.0” (711 mm)	

1.2.3 BladeCenter T

The IBM BladeCenter T is a compact eight-server blade chassis designed for high-density server installations, typically for telecommunications use. This 8U high chassis with DC or AC power supplies provides a cost-effective, high-performance, high-availability solution for

telecommunication network and other “rugged” non-telco environments. The IBM BladeCenter T chassis is positioned for expansion, capacity, redundancy, and carrier-grade NEBS level/ETSI compliance in DC models.

BladeCenter T provides a solid foundation for next-generation networks enabling service providers to become on demand providers. Coupled with technological expertise within the enterprise data center, IBM is leveraging the industry know-how of key Business Partners to jointly deliver added value within service provider networks.

Figure 1-9 displays the front view of the BladeCenter T.

Figure 1-9 BladeCenter T front view

Figure 1-10 displays the front view showing the key chassis features.

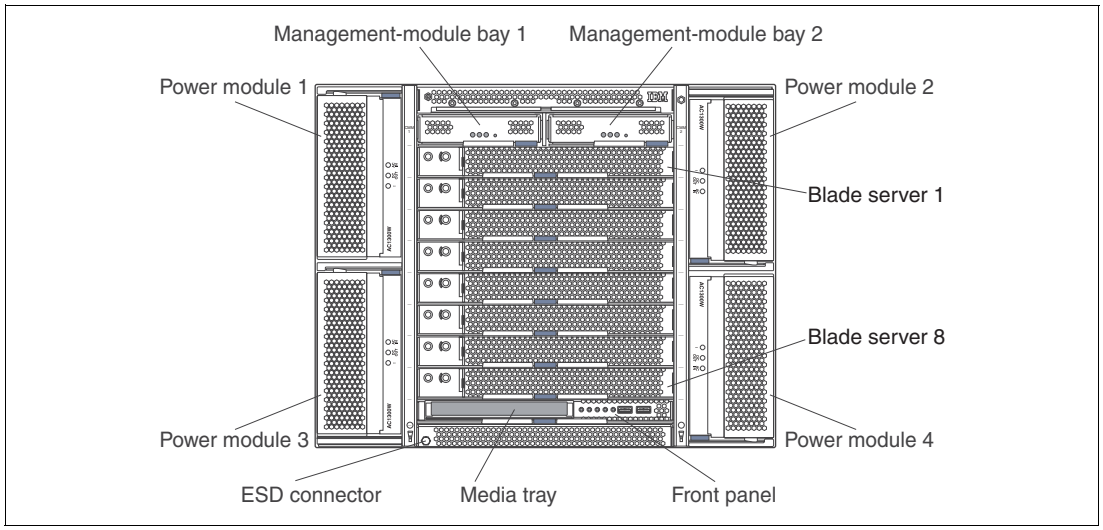

Figure 1-10 The IBM BladeCenter T chassis

The following features are the key features on the front of the BladeCenter T:

- A media tray at the front with a DVD/CD-RW drive (slim IDE), two USB v1.1 ports, and a system-status panel
- Two hot swap management-module bays
- Four hot swap power-module bays
- Front bezel with changeable filter
- Eight hot swap blade bays supporting different blade server types

Note: For a list of blade servers supported in the IBM BladeCenter T 8720 and 8730. see Table 1-1 on page 2.

Figure 1-9 on page 20 displays the rear view of the BladeCenter T.

Figure 1-11 BladeCenter T rear view (machine type 8730 with AC power inputs)

Figure 1-12 displays the rear view showing the key chassis features.

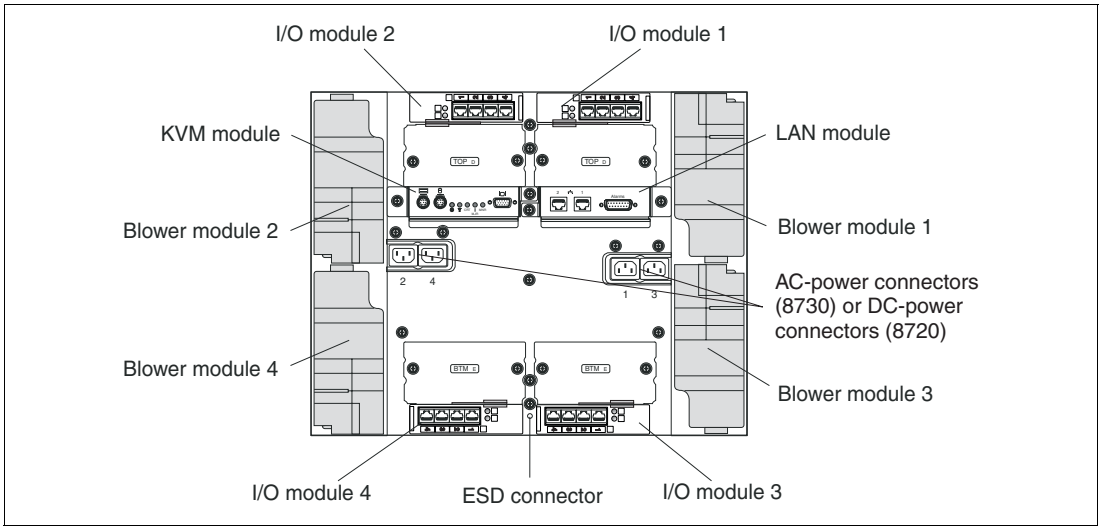

Figure 1-12 BladeCenter T rear view (type 8730 AC)

The following features are the key features on the rear of the BladeCenter T:

- Four hot swap I/O module bays
- One hot swap KVM (keyboard, video, and mouse) module
- One hot swap LAN module
- Four variable speed hot swap blowers
- AC power (machine type 8730) or DC power (machine type 8720) connectors

BladeCenter T delivers rich telecom features and functionality, including integrated servers, storage and networking, fault-tolerant features, optional hot-swappable redundant DC or AC power supplies and cooling, and built-in system management resources in a 20-inch (50.8 cm) deep chassis. The result is a Network Equipment Building Systems (NEBS-3) and ETSI-compliant server platform optimized for next-generation networks.

The BladeCenter T applications suited for these servers include:

- Network management and security
  - Network management engine
  - Internet cache engine
  - RSA encryption
  - Gateways
  - Intrusion detection
- Network infrastructure
  - Softswitch
  - Unified messaging
  - Gateway/Gatekeeper/SS7 solutions
  - VOIP services and processing
  - Gateways
  - Voice portals
  - IP translation database

The BladeCenter T chassis allows for either eight single-slot blade servers or four double-slot blade servers. However, you can mix different blade server models in one chassis to meet your requirements. Not all available Blade Servers are supported in the BladeCenter T chassis; refer to Table 1-1 on page 2 for details.

The BladeCenter T chassis ships standard with one advanced Management Module. This module provides the ability to manage the chassis, as well as providing the local KVM function.

The optional redundant advanced Management Module provides the IBM BladeCenter T with higher levels of resiliency. While in the chassis, the second module is in passive or standby mode. If the active or primary module fails, the second module is automatically enabled with all of the configuration settings of the primary module. This function provides clients with easy remote management and connectivity to the BladeCenter T chassis for their critical applications.

The two Ethernet connections from the two Management Modules are routed through to the LAN module at the rear of the chassis; see Figure 1-11 on page 21. There are no connections to the front of the chassis.

Note: The BladeCenter T (8720 and 8730) chassis started shipping standard with the Advanced Management Module in 1Q06. Prior to 1Q06, it shipped with the original management module.

The BladeCenter T does not ship standard with any I/O modules. You need to choose these I/O modules depending on your connectivity needs. An Ethernet Switch Module (ESM) is required in I/O module bays 1 and 2, to enable the use of both Ethernet ports on a Blade Server. The I/O modules required in I/O module bays 3 and 4 will depend on the I/O Expansion Card installed in the Blade Servers.

Guidelines about how to install I/O module options are discussed in detail in Chapter 2, “BladeCenter chassis and infrastructure configuration” on page 57.

The KVM module is a hot swap module at the rear of the chassis providing two PS/2 connectors for keyboard and mouse, a systems-status panel, and a HD-15 video connector.

Tip: In contrast to the standard BladeCenter chassis, the KVM module and the Management Module in the BladeCenter T chassis are separate components. The front of the Management Modules only feature the LEDs for displaying status. All Ethernet and KVM connections are fed through to the rear to the LAN and KVM modules.

The hot-swap LAN module at the rear of the chassis provide two Ethernet connections for the Management Modules. The hot-swap LAN module also provides a DSUB 15P telco alarm connector (male) for critical, major, and minor telco alarms as shown in Figure 2-43 on page 189. Each of the alarms has a relay that enables multiple systems alarm indicators to be daisy-chained together.

The BladeCenter T chassis is currently available with two kinds of power supplies:

- 8720-1RX, 8720-2RX: This type comes standard with two 1300 W DC (-48 to -60 V) power supplies.
- 8730-1RX, 8730-2RX: This type comes standard with two 1300 W AC (200-240 V) power supplies.

Both chassis types come standard with:

- One management (maximum of two supported)
- One KVM module
- One LAN module
- Four blowers
- Two USB 1.1 ports
- One DVD optical drive

The chassis does not have a diskette drive. An external FDD is supported attached to one of the two front-mounted USB ports. The BladeCenter T does not ship standard with any rails. Rails need to be ordered separately.

Note: The internal cabling, as shipped, is different inside the two chassis. Therefore, it is impossible to convert an existing chassis to the other flavor of power supply. An AC chassis cannot be changed into a DC chassis. A DC chassis cannot be changed in an AC chassis.

BladeCenter T is compliant with the NEBS 3 and ETSI core network specifications.

NEBS and ETSI compliance

NEBS stands for Network Equipment Building System and is a set of specifications originating from Bell Labs in the 1970s. It allows you to operate devices and equipment from different manufacturers in what is called a central office. A central office is a hosting site similar to a small data center room. However, as these central offices are distributed throughout the country the environmental conditions are much more extreme than in a

controlled data center environment. NEBS sets the criteria to which the manufacturers have to adhere.

There are three levels of NEBS specification.

he NEBS-Level 3 is the level with the strongest requirements, encompassing:

- Spatial requirements (equipment size and construction, central office physical layout, and so on)
- Temperature, humidity, and altitude criteria
- Fire resistance
- Equipment handling criteria
- Earthquake, office vibration, and transportation vibration
- Airborne contaminants
- Acoustical noise
- Electrostatic Discharge (ESD)
- Electromagnetic interference (emissions and immunity)
- Lightning and AC power fault
- Steady state power induction
- Electrical safety
- Bonding and grounding

The specification is now maintained by Telcordia at:

http://www.telcordia.com

The European Telecommunication Standards Institute (ETSI) maintains a set of specifications which is very similar to the NEBS specification. You can find more information at:

http://www.etsi.org

Note that the NEBS-3 and ETSI specifications are by far surpassing the operational and environmental criteria of ordinary data center requirements.

Tip: Although the BladeCenter T is NEBS-3 compliant, it does not automatically mean that every possible configuration is compatible. Each component and option have to be tested separately. However, some options, for example, hard disks or HS20 type 8843 blade servers with 3.6 GHz CPUs, do not meet the NEBS certification requirements. For this reason, a NEBS compatibility list is maintained at:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-54720

Table 1-13 BladeCenter T features at a glance

Feature	Specification
Machine type	8720-2RX (DC) 8730-2RX (AC)
Rack form factor (H x D)	8U x 20 inches (508 mm)
DVD/CD drives standard	1x DVD-ROM (in Media Tray)
Diskette drive standard	None
Number of Blade Server slots	8 (30 mm blade servers)
Number of switch module slots	4
Switch modules standard	None (in standard chassis offerings)

Feature	Specification
Power supply size standard	1300W DC (8720-2RX), -38V to -75V (-48V nominal); 1300W AC (8730-2RX)
Number of power supplies (standard/maximum)	2/4
Number of blowers (standard/maximum)	4/4
Dimensions	Height: 13.75” (349 mm) Width: 17.4” (442 mm) Depth: 20.0” (508 mm)

1.2.4 BladeCenter HT

The IBM BladeCenter HT is a 12-server blade chassis designed for high-density server installations, typically for telecommunications use. It offers high performance with the support of 10 G Ethernet installations. This 12U high chassis with DC or AC power supplies provides a cost-effective, high-performance, high-availability solution for telecommunication network and other *rugged* non-telco environments. IBM BladeCenter HT chassis is positioned for expansion, capacity, redundancy, and carrier-grade NEBS level 3/ETSI compliance in DC models.

BladeCenter HT provides a solid foundation for next-generation networks (NGN) enabling service providers to become on demand providers. Coupled with technological expertise within the enterprise data center, IBM leverages the industry know-how of key Business Partners to jointly deliver added value within service provider networks.

Figure 1-13 shows the front view of the BladeCenter HT.

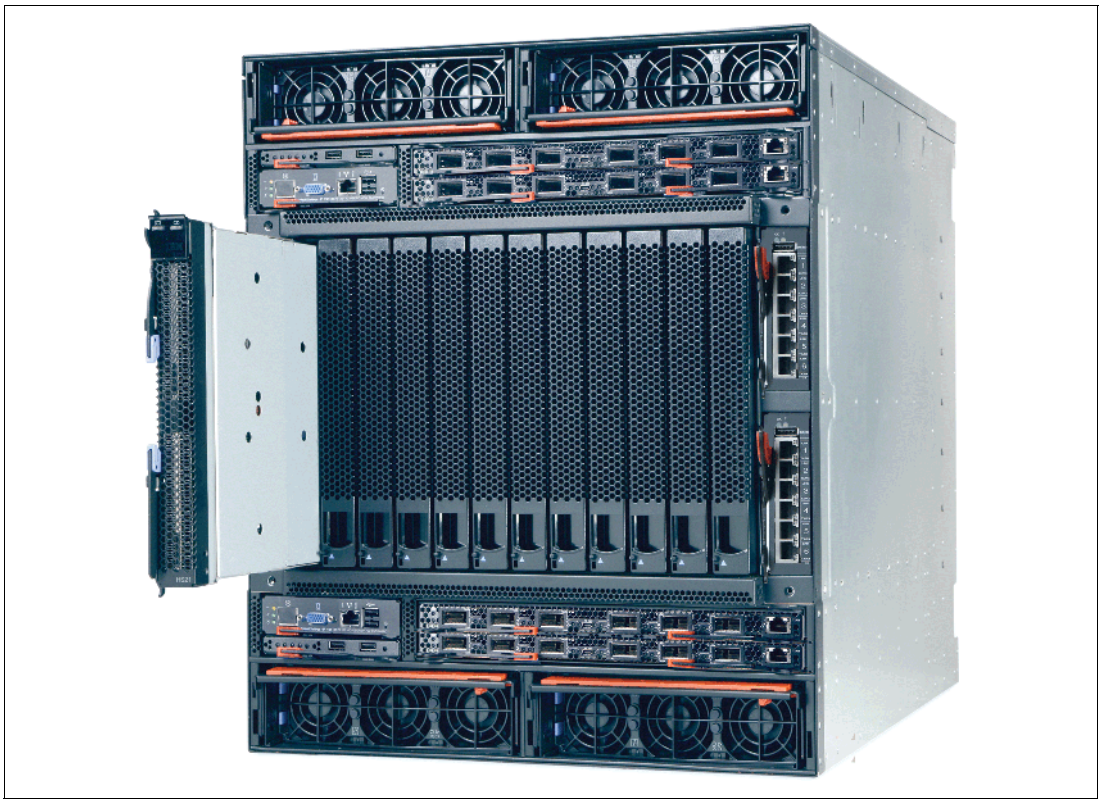

Figure 1-13 BladeCenter HT front view

The BladeCenter HT brings significant new capabilities to the broad IBM ecosystem of hundreds of NGN applications already being deployed on BladeCenter. A key example is the introduction of the Nortel 10 Gb Ethernet Switch Module for BladeCenter, which delivers 10 Gb to each blade server deployed in the BladeCenter H or BladeCenter HT chassis, and six 10 Gb Ethernet uplinks. This capability helps greatly reduce the cost of implementing IPTV and other high bandwidth NGN applications.

Figure 1-14 on page 27 displays the rear view showing the key chassis features.

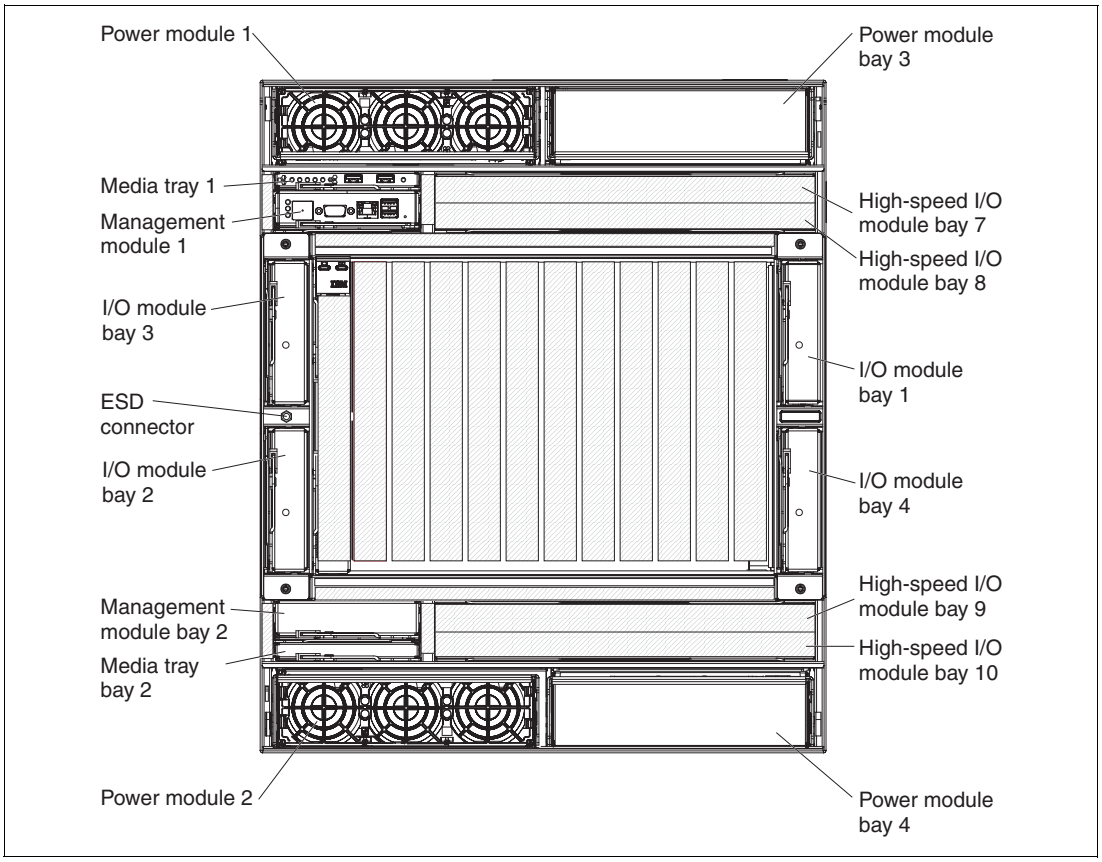

Figure 1-14 The IBM BladeCenter HT chassis

The key features of the BladeCenter HT include:

- Support for up to 12 blade servers, compatible with the other chassis in the BladeCenter family
- Four low-speed and four high-speed I/O module bays, compatible with the other chassis in the BladeCenter family
- A media tray at the front with light path diagnostics, two USB 2.0 ports, and optional compact flash memory module support
- Two hot swap management-module bays (one management module standard)
- Four hot swap power-module bays (two power modules standard)
- New serial port for direct serial connection to installed blades
- Compliance with the NEBS 3 and ETSI core network specifications

Figure 1-15 shows the rear view of the BladeCenter HT.

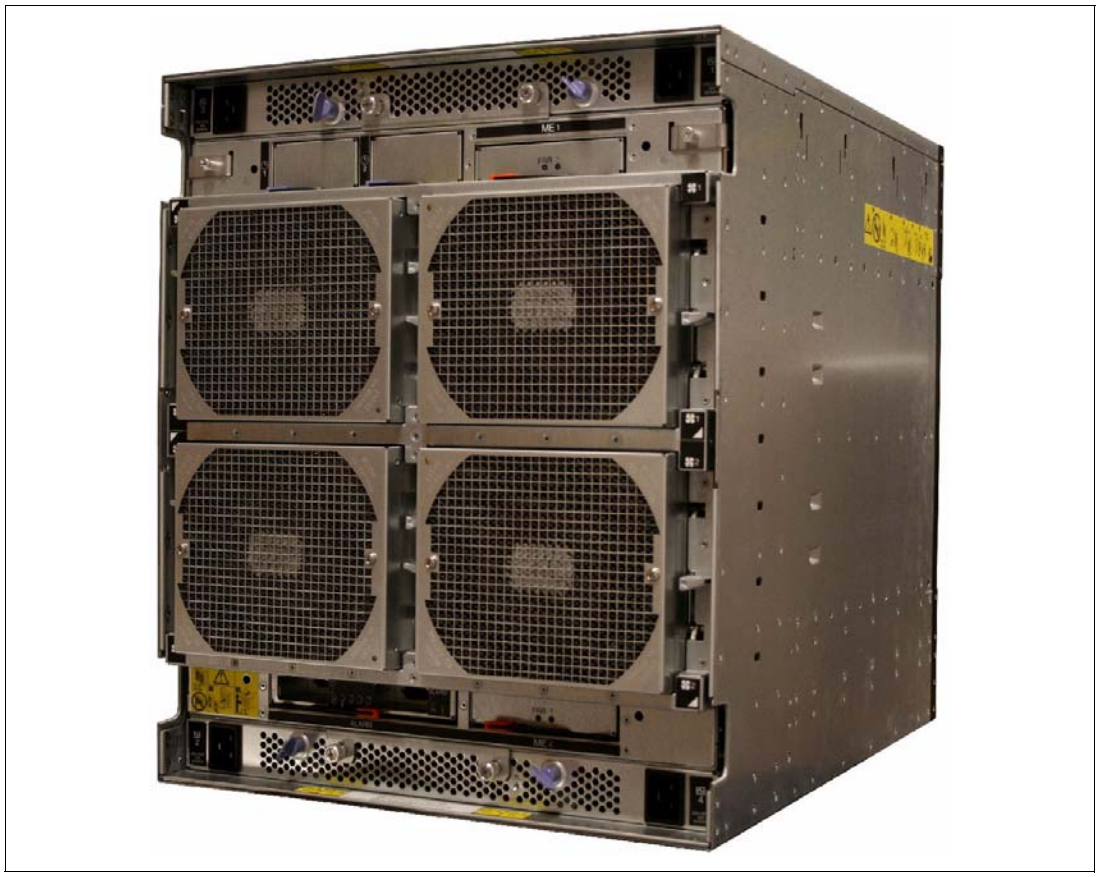

Figure 1-15 BladeCenter HT rear view

Figure 1-16 on page 29 displays the rear view showing key chassis features.

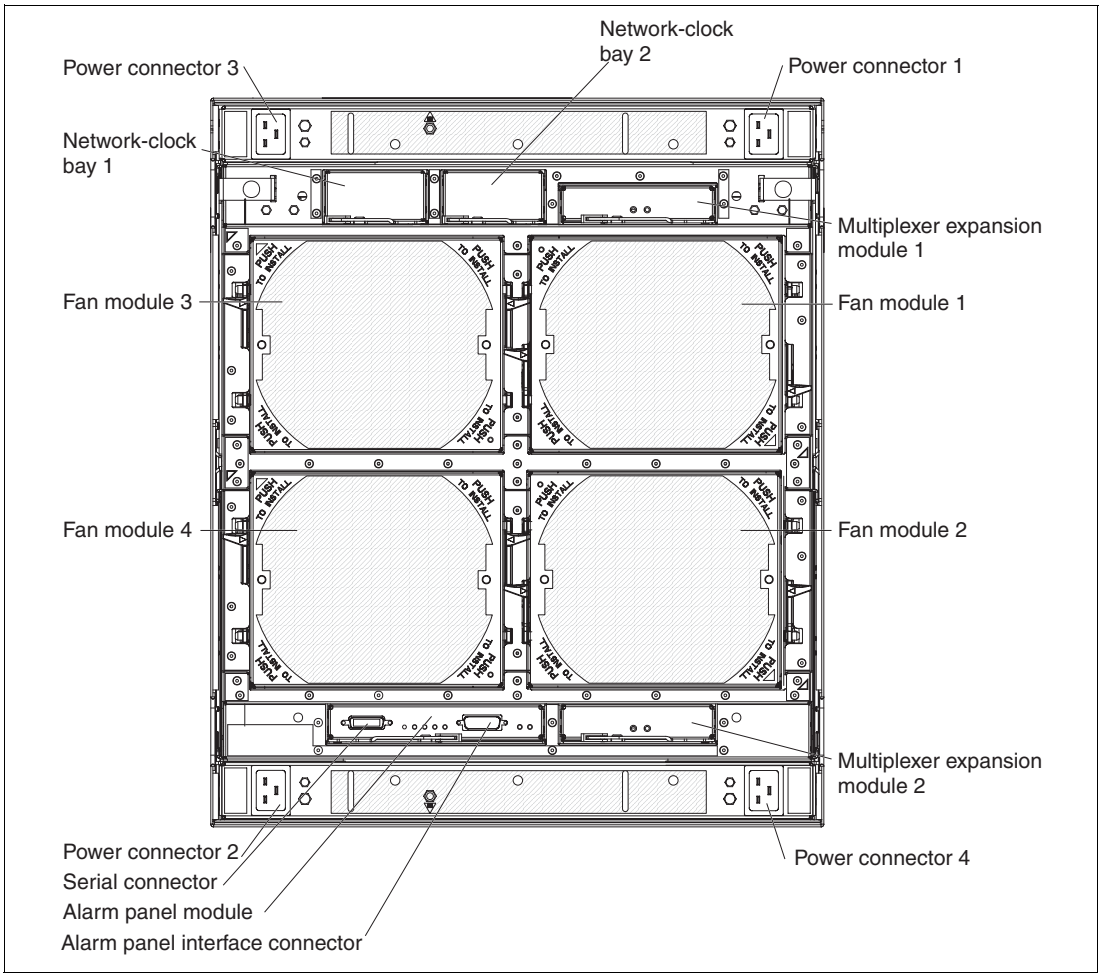

Figure 1-16 Diagram of BladeCenter HT (AC model) rear view

BladeCenter HT delivers rich telecom features and functionality, including integrated servers, storage and networking, fault-tolerant features, optional hot swappable redundant DC or AC power supplies and cooling, and built-in system management resources. The result is a Network Equipment Building Systems (NEBS-3) and ETSI-compliant server platform optimized for next-generation networks.

The BladeCenter HT applications suited for these servers include:

- Network management and security
  - Network management engine
  - Internet cache engine
  - RSA encryption
  - Gateways
  - Intrusion detection
- Network infrastructure
  - Softswitch
  - Unified messaging
  - Gateway/Gatekeeper/SS7 solutions
  - VOIP services and processing
  - Voice portals
  - IP translation database

The BladeCenter HT chassis allows for either 12 single-slot blade servers or six double-slot blade servers. However, you can mix different blade server models in one chassis.

The BladeCenter HT chassis ships standard with one Advanced Management Module. This module provides the ability to manage the chassis as well as providing the local KVM function. The optional redundant Advanced Management Module provides the IBM BladeCenter HT with higher levels of resiliency. This module provides clients with easy remote management and connectivity to the chassis.

The BladeCenter HT does not ship standard with any I/O modules. You need to choose these I/O modules depending on your connectivity needs. An Ethernet Switch Module (ESM) is required in I/O module bays 1 and 2 to enable the use of both Ethernet ports on a blade server. The I/O modules required in I/O module bays 3 and 4 depends on the I/O Expansion Card installed in the blade servers.

High-speed switch modules can be installed into I/O bays 7 through 10, and are used together with high-speed expansion cards installed into blade servers. I/O bays 1 trough 4 can also be used for bridge modules.

Table 1-14 BladeCenter HT features at a glance

Feature	Specification
Machine type	8740-1RY (DC) 8750-1RY (AC)
Rack form factor (H x D)	12U x 27.8 inches (706 mm)
DVD/CD drives standard	None
Diskette drive standard	None
Number of Blade Server bays	12 (30 mm blade servers)
Number of standard switch module bays	4
Number of high-speed switch module bays	4
Switch modules shipped standard	None (in standard chassis offerings)
Number of power supplies (standard/maximum)	2/4
Number of fan modules (standard/maximum)	4/4
Dimensions	Height: 21.0” (528 mm) Width: 17.4” (442 mm) Depth: 27.8” (706 mm)

1.2.5 BladeCenter S

The BladeCenter S chassis is a robust and flexible physical platform. Its modular tool-free design allows easy access and maintenance. All external components (except running blade servers) are hot-swappable and clearly marked with release levers/handles.

The key features on the front of the BladeCenter S are indicated in Figure 1-17.

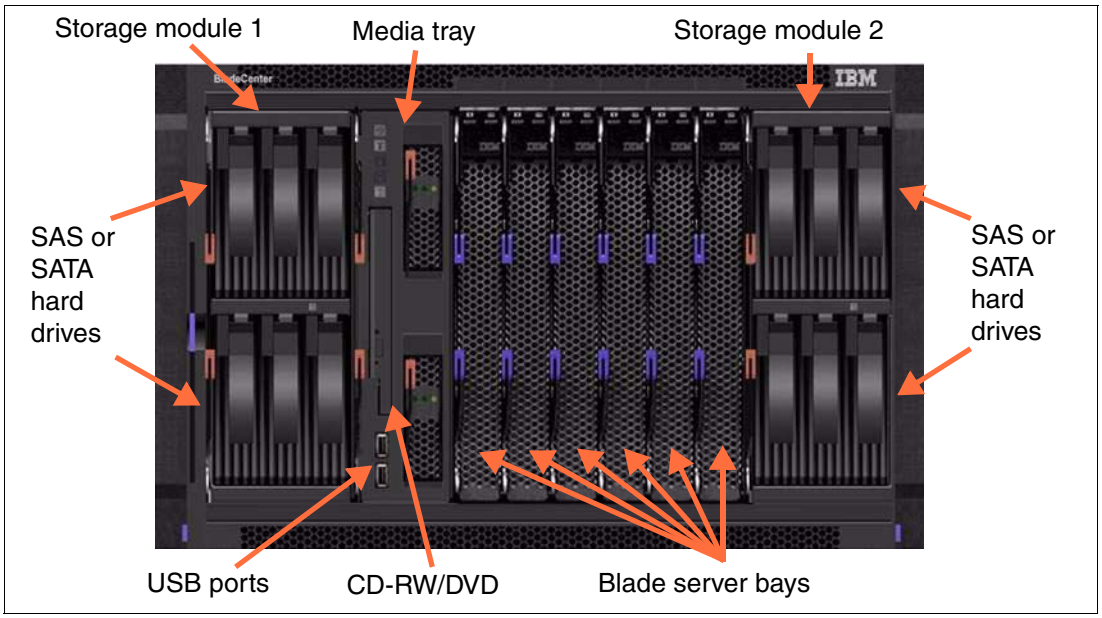

Figure 1-17 The front of the BladeCenter S chassis

The following features are the key features on the front of the BladeCenter S:

- A media tray at the front, with a DVD drive, two USB 2.0 ports, system status LED panel, and two bay locations that are reserved for future use.
- Six hot swap blade server bays supporting different blade server types.
- Two bays for disk storage modules; each storage module can house up to six 3.5-inch disk drives of internal storage. No storage modules are standard with the BladeCenter S chassis.

Note: For a list of blade servers supported in the IBM BladeCenter S, see Table 1-1 on page 2.

Figure 1-18 shows the key features on the back of the chassis.

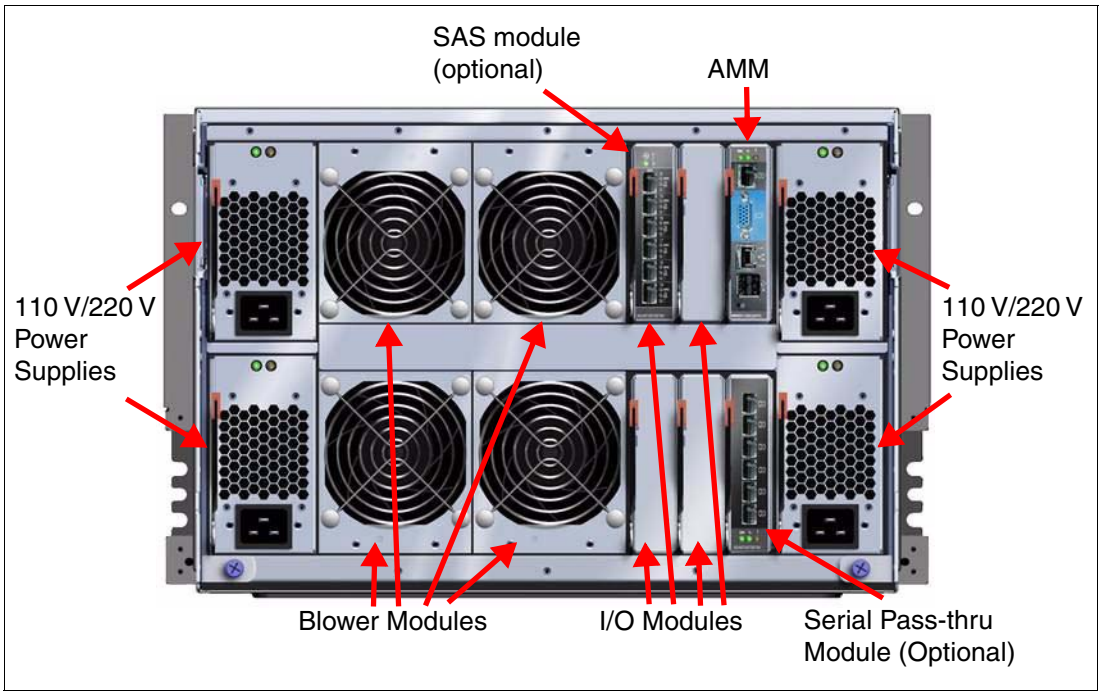

Figure 1-18 The rear of the BladeCenter S chassis

The following features are the key features on the rear of the BladeCenter S:

- Four hot-swap blower modules as standard
- One hot swap management-module bay with one management-module as standard
- Four I/O bays for standard switch modules
- One pair of 950/1450-watt power modules. An additional power module option (containing two 950/1450 W power modules) is available. Models with either C14 and C20 input sockets are available as listed in Table 1-15 on page 33.
- One bay for Serial pass-thru module to give direct serial connection to installed blades (for those blades with the functionality)

The BladeCenter S chassis allows for either six single-slot blade servers or three double-slot blade servers. However, you can mix different blade server models in one chassis to meet your requirements.

The BladeCenter S chassis ships standard with an Advanced Management Module. This module provides the ability to manage the chassis and well as providing the local KVM function. Only one management module is supported with BladeCenter S chassis.

The BladeCenter S does not ship standard with any I/O modules. You choose these I/O modules based on your connectivity needs. An Ethernet Switch Module (ESM) will be required in I/O module bay 1, to enable the use of both Ethernet ports on a Blade Server. The I/O modules required in I/O module bays 3 and 4 depend on the I/O Expansion Card that is installed in the blade servers. Bay 2 is reserved for future use. The chassis does not ship with any storage modules.

We discuss guidelines about how to install I/O module options in detail in Chapter 2, “BladeCenter chassis and infrastructure configuration” on page 57.

The BladeCenter S chassis uses either 100 to 127 V or 200 to 240 V AC power and can be attached to standard office power outlets.

The BladeCenter S chassis ships standard with:

- One advanced management module
- Four blower modules
- Two power supply modules (one pair of 950/1450-watt power modules).
- Two 2.5 m rack jumper power cords (IEC 320 C19 - C20)
- Four country-specific line cords
- One CD-RW/DVD-ROM drive

The chassis does not have a diskette drive. An optional USB-attached 1.44 MB diskette drive is available.

Table 1-15 BladeCenter S features at a glance

Feature	**Specifications**	
Machine type	8886-1MY	8886-1NG
Rack form factor (H x D)	7U x 28.9 inches (733.4 mm)	
DVD/CD drives standard	1x DVD-ROM (in Media Tray)	
Diskette drive standard	None	
Number of Blade Server slots	6 (30 mm blade servers)	
Number of switch module slots	4 hot-swap (3 x standard, 1 x standard is reserved for future use)	
Switch modules (std/max)	0 / 4	
Storage modules (std/max)	0 / 2	
Power supply size standard	950/1450 W AC	
Number of power supplies (standard/maximum)	2/4	
Input connectors	IEC 320 C20	IEC 320 C14
Number of blowers (standard/maximum)	4/4	
Dimensions	Width: 440.0 mm (17.5 in) Depth: 733.4 mm (28.9 in) Height: 306.3 mm (12 in)	
Availability	Worldwide	Denmark, Switzerland and Sweden

1.3 Blade servers

IBM BladeCenter servers supports a wide selection of processor technologies and operating systems to allow clients to run all of their diverse work loads inside a single architecture. The slim, hot-swappable blade servers fit in a single chassis like books in a bookshelf, and each is an independent server, with its own processors, memory, storage, network controllers, operating system and applications. The blade server simply slides into a bay in the chassis and plugs into a midplane or backplane, sharing power, fans, diskette drives, switches, and ports with other blade servers.

The benefits of the blade approach will be obvious to anyone tasked with running down hundreds of cables strung through racks just to add and remove servers. With switches and power units shared, precious space is freed up, and blade servers enable higher density with far greater ease.

Table 1-1 on page 2 shows which blade servers are supported in each of the IBM BladeCenter chassis.

Blade servers covered in this section are:

1.3.1 BladeCenter HC10

BladeCenter HC10 blade workstations, coupled with the BladeCenter chassis, deliver advanced application serving with performance, density, and scalability ideal for enterprise environments.

Figure 1-19 on page 35 shows the features and functions of HC10.

Figure 1-19 IBM BladeCenter HC10

Features and functions of HC10 include:

- Single-wide, high-performance workstation blade
- Powerful dual-core processors Intel® Core 2 Duo
- L2 cache per core (up to 4 MB per processor)
- 1066 MHz front-side bus (FSB)
- Two Broadcom 5708S Gigabit Ethernet connections
  - One for the on-board NIC
  - One for the compressed video connection
- Up to 8 GB of non-buffered, non-ECC, DDR2 high-speed memory (four DIMM slots x 2 GB/DIMM)
- One SATA HDD
- Integrated systems-management processor
- Advanced high-availability and systems-management features

The list of supported operating systems include Windows Vista and Windows XP Professional. See the following ServerProven page for further supported NOS information:

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

For more details about the server, see 3.2, “HC10 workstation blade” on page 256.

1.3.2 BladeCenter HS12

The BladeCenter HS12 blade servers support the economics of application server deployment with power, scalability, control, and serviceability. They are well-suited for Web caching, terminal serving, firewalls, dynamic Web serving, and virtualization.

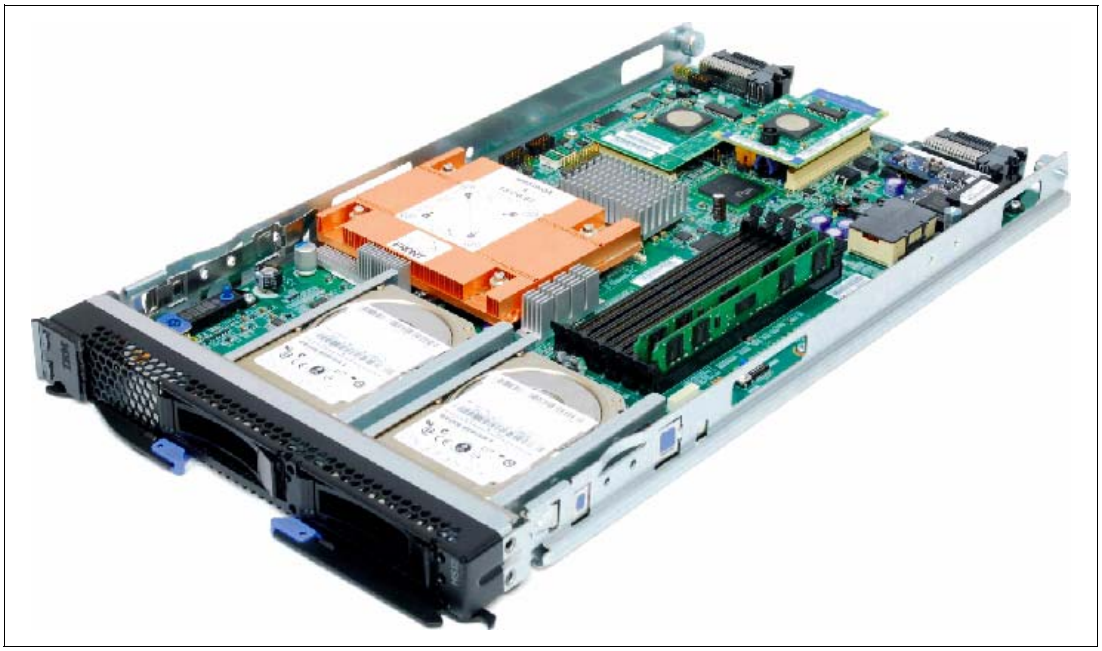

Figure 1-20 IBM BladeCenter HS12

The HS12 includes these features to support its power, scalability, control, and serviceability:

- 1.86 GHz processor with 512 MB or 2 MB L2
- 2.13 GHz dual-core processor with 2 MB L2
- 2.50 GHz quad-core processor with 6 MB L2 or 2.66 and 2.83 GHz quad-core with 12 MB L2
- Dual Gigabit Ethernet connections
- Support for up to non-hot-swap SATA/solid state drives or up to two hot-swap SAS HDDs, and RAID 0 and 1 support with RAID card
- Support for Ethernet or Fibre Channel expansion cards

The supported operating systems for the HS12 blade servers are Microsoft Windows 2003 32-bit and 64-bit, Microsoft Windows Server 2008 Enterprise, Standard, and Web edition, RHEL 5 32-bit and 64-bit edition, and SUSE Linux Enterprise Server 10. See the following ServerProven page for further supported NOS information:

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

For more details about the server, see 3.3, “HS12” on page 261.

1.3.3 BladeCenter HS21 XM

The BladeCenter HS21 XM blade servers are designed for front and mid-tier applications requiring high performance, enterprise class availability, and extreme flexibility and power efficiency.

Figure 1-21 on page 37 shows the HS21 XM blade server.

Figure 1-21 The HS21 XM

The HS21 XM has the ability to grow with your application requirements, thanks to:

- Up to two multi-core Xeon processors (four or eight cores in all).
- Low-voltage processors available that use as little as 20W per core.
- Up to 64 GB of system memory in 8 DIMM slots provides flexibility in memory configurations
- One internal 2.5-inch SAS HDD drive that supports either a single SAS drive or a single or dual solid state drive.
- An optional Modular Flash Drive can be used in place of, or in addition to, the internal SAS HDD, as a boot device (Linux only).
- Two Gigabit Ethernet ports standard; plus more, using either a 2-port Gigabit Ethernet Expansion Card or an I/O Expansion Unit.
- Supports optional expansion blade: IBM BladeCenter Storage and I/O (SIO) Expansion Blade with support for up to three additional hot-swap SAS SFF HDDs with RAID-0 and RAID-1 (through ServeRAID 8k-l adapter as standard), as well as two additional I/O expansion slots. An optional ServeRAID 8k adapter is available that will add RAID-5 capability and a battery backed cache. See 3.18.2, “Storage and I/O Expansion Unit” on page 363, for information about the SIO.
- Advanced high-availability and systems-management features
- Concurrent keyboard, video, mouse (cKVM) support with the addition of the optional IBM BladeCenter Concurrent KVM feature card. See 3.17, “Concurrent KVM Feature Card (cKVM)” on page 356.
- Three-year, onsite limited warranty.

The supported operating systems are Microsoft Windows Server 2003 32-bit and 64-bit, RHEL 4 32-bit and 64-bit, SLES 9/10 32-bit and 64-bit, Novell NetWare 6.5, Sun Solaris 10, VMware ESX Server 3.0. See the following ServerProven page for further supported NOS information:

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

For more details about the server, see 3.4, “HS21 XM type 7995” on page 274.

1.3.4 BladeCenter HS21

The BladeCenter HS21 offerings are positioned as high-density servers. They represent a new approach to the deployment of application servers where 2-way, SMP-capable Xeon processing, high-availability design, systems management, and easy setup features are combined in an extremely dense package.

The HS21 blades require less space and power resources, when compared to an equivalent standard 1U server, because of their high-density design, reduced power requirements and single environment systems management.

The supported operating systems include those from Microsoft, Red Hat, SUSE Linux, VMware and Novell. See the following ServerProven page for further supported operating system information:

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

HS21 type 8853

BladeCenter HS21 type 8853 blade servers deliver advanced application serving with performance, density, and scalability.

Figure 1-22 shows the HS21 type 8853 blade server.

Figure 1-22 The HS21 type 8853

Features, functions, support, and services at your command include:

- Blade servers supported in all BladeCenter chassis.
- Intel Xeon® 5100 and 5300 series dual and quad core processors.
- 4 MB L2 cache (shared between both cores).
- 1066 MHz or 1333 MHz front-side bus speed.
- Dual Gigabit Ethernet connections.
- Four DIMM sockets standard, supporting up to 32 GB of high-speed memory with 8 GB DIMMs.

  When coupled with the Memory and I/O (MIO) Expansion Blade, an additional four DIMM sockets are made available for a total of eight sockets (64 GB using 8 GB DIMMs). See 3.18.1, “BladeCenter Memory and I/O Expansion Unit” on page 361, for information about the MIO.
- Up to two non-hot swap SAS SFF HDDs with RAID-0 or RAID-1 supported via the onboard LSI Logic 53C1046E controller.
- Supports optional expansion blades:
  - IBM BladeCenter Storage and I/O (SIO) Expansion Blade with support for up to three additional hot-swap SAS SFF HDDs with RAID-0 and RAID-1 (via ServeRAID 8k-l adapter as standard) as well as two additional I/O expansion slots. An optional ServeRAID 8k adapter is available that will add RAID-5 capability and a battery backed cache. See 3.18.2, “Storage and I/O Expansion Unit” on page 363, for information about the SIO.
  - IBM BladeCenter Memory and I/O Expansion Blade allows expansion up to 32 GB of high-speed memory with ECC (eight DIMM slots x 4 GB DIMMs), one additional I/O Expansion slot, and additional two Dual Gigabit Ethernet connections (total of four). See 3.18.1, “BladeCenter Memory and I/O Expansion Unit” on page 361.
- Advanced high-availability and systems-management features.

- Concurrent keyboard, video, mouse (cKVM) support with addition of the optional IBM BladeCenter Concurrent KVM feature card. See 3.17, “Concurrent KVM Feature Card (cKVM)” on page 356.
- Three-year, onsite limited warranty.

The supported operating systems are Microsoft Windows Server 2003 32-bit and 64-bit, RHEL 3/4 32-bit and 64-bit, SLES 9 32-bit and 64-bit, Novell NetWare 6.5, Sun Solaris 10, VMware ESX Server 3.0. See the following ServerProven page for further supported NOS information:

http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/ematrix.shtml

For more details about the server, see 3.5, “HS21 type 8853” on page 288.

1.3.5 BladeCenter HS20

The IBM BladeCenter HS20 blades are built for performance, scalability and availability, also allowing complete flexibility in network and storage attachment. Two midplane connections provide durable and reliable connections to all chassis resources. The HS20 contains advanced technologies for increasing uptime such like Predictive Failure Analysis (PFA) on all major components and an integrated systems management processor.

The HS20 supports many operating systems including those from Microsoft, Red Hat, SUSE Linux, VMware and Novell. See the following ServerProven page for detailed support information:

http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/ematrix.shtml

The HS20 type 8843 type has Intel Xeon processors up to 3.8 GHz with 64-bit extensions, a front-side bus speed of 800 MHz, and onboard SCSI hard disks.

You can use the HS20 8843 in all chassis as shown in Table 1-1 on page 2. However, only some models meet the strict criteria for NEBS-3/ETSI certification.

Tip: The HS20 models 8843-JTY and 8843-MTY blade servers for the BladeCenter T and HT chassis have an extended production life cycle of three years from the date of General Availability (GA).

The blade server HS20 type 8843 is an advancement in technology from previous HS20 type blade servers. It offers true 64-bit technology on Intel processors, better performance, enhanced I/O expansion capabilities, and more flexibility.

The key features of the blade server HS20 type 8843 are the following:

- Up to two Intel Xeon processors up to 3.8 GHz with EM64T and 800 MHz front-side bus with 2 MB L2 cache with MicroBurst architecture and Hyper-Threading technology
- Four DIMM sockets, supporting high-speed ECC DDR2 DIMMs, allowing up to 16 GB of memory
- Dual Gigabit Ethernet (Broadcom 5704S) connections with failover support
- Support for an optional small-form factor I/O expansion card which allow additional connectivity options
- Integrated systems management processor

- Integrated SCSI controller (LSI 53C1020) with RAID-1 and RAID-1E (mirroring) capability and connectors for up to two onboard 2.5-inch small form factor (SmFF) SCSI hard disk drives
- Optional Ultra320 SCSI hard disk drive storage and high-availability RAID-1 support (through PCI Expansion Unit II)
- Operating system support includes Windows 2000 and Windows 2003, Novell NetWare, Red Hat Enterprise Linux, SUSE Linux Enterprise Server, VMware ESX Server
- Three-yea, onsite limited warranty

Figure 1-23 shows the HS20 type 8843 blade server.

Figure 1-23 HS20 type 8843

There are some models of HS20 type 8843 server blades that are available with low profile handles to accommodate the front bezel air filter of the BladeCenter T and HT. However, this does not automatically mean that these blade servers are NEBS-3/ETSI compliant. See further details in “NEBS and ETSI compliance” on page 23 to get compatibility information.

For more details about the server, see 3.6, “HS20 type 8843” on page 296.

1.3.6 BladeCenter JS12

The BladeCenter JS12 Express provides excellent processing power, along with maintainability and scalability. Ideal for small database and application servers, the JS12 is made to be an easy-to-use integrated platform.

The JS12 is shown in Figure 1-24.

Figure 1-24 BladeCenter JS12

The JS12 is a single-wide server that offers these features:

- Two 3.8 GHz POWER6™ processor cores
- Dual-core SMP scalability
- Built-in PowerVM™ Standard Edition
- Support for System Cluster 1350™
- High-performance DDR2 memory running at up to 667 MHz
- 4 MB 4-way set-associative L2 cache per processor core
- Up to two 73 or 146 GB 2.5-inch SAS hard drives

The BladeCenter JS12 runs on the following environments: AIX V5.3 or later, AIX V6.1or later, iV6R1 or later, Novell SUSE Linux Enterprise Server 10 for POWER or later, RHEL 6 for POWER V4.6 or later, and RHEL 6 for POWER v5.1 or later. See the following ServerProven page for more information about compatible operation systems:

http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/ematrix.shtml

For more details about the server, see 3.11, “JS12 Type 7998-60x” on page 329.

1.3.7 BladeCenter JS22

BladeCenter JS22 is a highly desirable blade server for organizations of any size, due to its exceptional price/performance in a virtualized environment, breadth of available applications, reliability availability and serviceability (RAS) characteristics, and energy efficiency. The JS22 is powered by IBM POWER6 processor cores, and is suitable for clients who want a durable server solution for consolidating multiple applications, enabling high-performance computing, and implementing commercial applications in an AIX or Linux environment.

Figure 1-25 on page 42 displays details of the BladeCenter JS22.

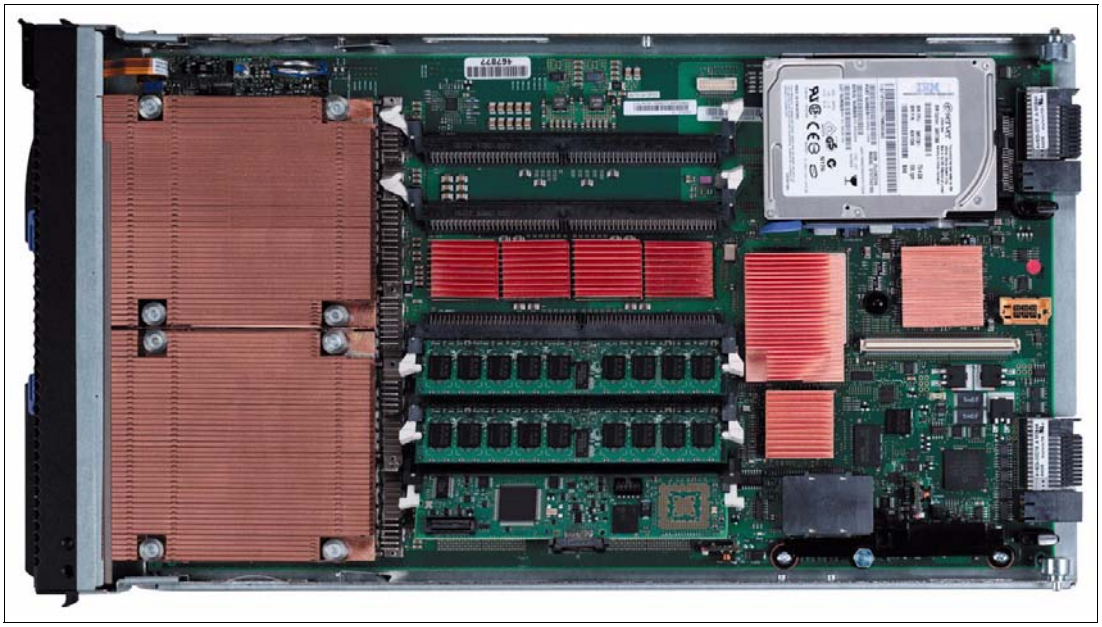

Figure 1-25 The internals of the BladeCenter JS22

The BladeCenter JS22 is a single-width blade server that offers these features:

- Two dual-core processors, using POWER6 technology
- Support for Energy Scale thermal management for power management/oversubscription (throttling) and environmental sensing
- 4.0 GHz POWER6 SCM processors, directly mounted to the planar board
- L1 cache per core: 64 KB I-cache and 64 KB D-cache
- L2 cache: 4 MB per core
- Four DIMM sockets supporting up to 32 GB of RAM (using four 8 GB DIMMs)
- Memory is 667 MHz DDR2 (1, 2, and 4 GB DIMMs) or 533 MHz DDR2 (8 GB DIMMs); supports ECC, Chipkill and bit steering protection features
- One internal connector for a SAS hard disk drive
- Support for PCI-X (StFF, SFF, CFFv) and PCI Express (HSFF, CFFh) expansion cards
- Advanced POWER Virtualization hardware feature supports Integrated Virtualization Manager (IVM) and Virtual I/O Server
- Integrated BMC service processor plus pluggable management card (next to DIMM sockets in Figure 1-25)
- ATI™ RN 50 ES1000™ video controller.

The list of supported OSes include AIX 6.1, SLES 10 for IBM POWER, and RHEL 5 for System i and System p.

See the following ServerProven page for further supported NOS information:

http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/ematrix.shtml

For more information about JS22, refer to 3.13, “JS22 Type 7998-61x” on page 340.

1.3.8 BladeCenter JS21

Built for speed and reliability, the JS21 blade server supports a choice of operating systems for running HPC Linux clusters as well as AltiVec optimized applications on AIX or Linux. BladeCenter packaging makes POWER technology more affordable.

Figure 1-26 shows the IBM BladeCenter JS21.

Figure 1-26 IBM BladeCenter JS21

With virtualization support standard on the JS21, server and workload consolidation of multiple independent applications on a single blade can be provided by PowerVM for AIX 5L V5.3 and Linux environments by ordering Virtual I/O Server (VIOS) V1.2.1.

The BladeCenter JS21 is a single-width blade server that offers these features:

- Up to two 2.5 or 2.7 GHz PowerPC® 970MP microprocessors.
- Up to 16 GB of GB ECC Chipkill DDR2 memory (4 DIMM sockets) for memory-intensive applications.
- Integrated dual Gigabit Ethernet Broadcom 5780.
- Integrated systems management processor
- Integrated SAS controller with up RAID 0 or RAID 1 support
- Support for up to two 2.5 inch SAS hard disk drives
- Support for Broadcom 5780: PCI-Express and dual gigabit Ethernet; Fibre Channel, Myrinet, InfiniBand, and iSCSI TOE connections; AltiVec optimized application development; PowerVM (enabled through Virtual I/O Server)
- Operating system support includes SUSE Linux Enterprise Server 9 SP3, or later, Red Hat Enterprise Linux AS 4 U3, or later, for POWER, and AIX 5L V5.2 and V5.3.

See the following ServerProven page for further supported NOS information:

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

For more details about the server, see 3.12, “JS21 type 8844” on page 335.

1.3.9 BladeCenter LS22

The LS22 single slot blade server is capable of supporting up to two quad core AMD™ Opteron™ 2300 series processors and up to 32 GB DDR2 memory (in 8 DIMM slots with 4 GB DIMM modules) with Chipkill protection, for high performance and reliability. The latest AMD Opteron processors are designed with 2 MB of L2 cache and another 6 MB of L3 cache, 64-bit extensions, and HyperTransport™ technology, to help provide the computing power you require to match your business needs and growth.

The LS22 supports up to two non-hot swap 2.5-inch SAS or SATA HDDs or up to two non-hot swap solid state drives. A dual port Gigabit Ethernet controller is standard, providing high-speed data transfers and offering TCP Offload Engine (TOE) support, load-balancing, and failover capabilities.

With optional expansion cards, the LS22 can connect to additional Ethernet, Myrinet, Fibre Channel, InfiniBand and other high-speed communication switches housed in the chassis. Optional expansion cards can add additional fabrics to the LS22 as needed. The IBM BladeCenter LS22 is shown in Figure 1-27.

The BladeCenter LS22 is a single-width blade server that offers these features:

- Up to two high-performance, AMD Quad-Core Opteron processors
- A system board containing eight DIMM connectors, supporting up to 32 GB of system memory with 4 GB DIMMs
- An SAS controller, supporting up to two internal SAS, SATA, or solid state disk drives
- Two Gigabit Ethernet ports (Broadcom 5709S) standard with TOE support
- An internal USB connector
- Support for concurrent KVM (cKVM) and concurrent USB/DVD (cMedia) via Advanced Management Module and optional daughter card
- Support for the PCI Express I/O expansion (PEU3e) unit and PCI expansion (PEU2) unit II
- A three-year, onsite limited warranty

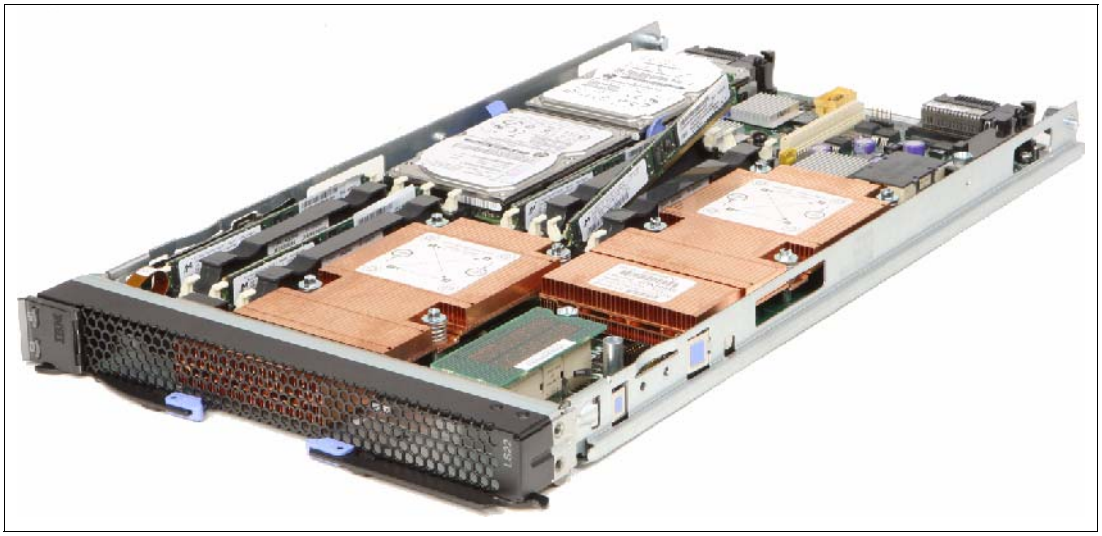

Figure 1-27 IBM BladeCenter LS22

The LS22 supports the Windows and Linux operating systems. See the following ServerProven page for further supported NOS information:

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

Table 1-16 lists the LS22 features and specifications.

Table 1-16 Features of the LS22

Feature	Specification
Processor	AMD Opteron 2300 series quad core
Number of processors (std/max)	1/2
Cache	512 KB L2 per processor core, 2 MB or 6 MB shared L3 per processor
Memory	8 VLP DIMM slots/DDR2 667Mhz or 800Mhz/32 GB maximum
Internal hard disk drives (standard/maximum)	0/2
Maximum internal storage	On board: two 2.5” Non-Hot-Swap SAS, SATA, or solid state disk drives
Network	Two ports, Integrated Dual Gigabit Ethernet (Broadcom 5709S) TOE
I/O upgrade	One PCI-X expansion connector and one PCI-Express expansion connector

For more details about the server see 3.9, “LS22 type 7901” on page 313.

1.3.10 BladeCenter LS42

The BladeCenter LS42 blade servers are scalable AMD Opteron-based servers. The LS42 is a two-socket offering that can be expanded to become a four-socket offering. It is optimized for scalable enterprise workloads and scientific computing.

When fully configured, the LS42 can support up to 64 GB of DDR2 ECC memory in 16 DIMM slots and up to four quad core AMD Opteron 8300 series processors.

The LS42 supports up to two SAS or SATA disk drives or up to two solid state drives. The onboard SAS controller provides mirroring capability for these internal drives.

A dual port Gigabit Ethernet controller is standard in both the base unit and the multi-processor expansion unit, for a total of four built-in Gigabit Ethernet ports offering TCP/IP Offload Engine (TOE) support, load-balancing, and failover capabilities. Optional I/O Expansion Cards add additional fabrics to the LS42 as needed. Each LS42 blade server can connect to additional Ethernet, SAS, Fibre Channel, iSCSI, InfiniBand and other high-speed communication switches housed in the chassis. The IBM BladeCenter LS42 is shown in Figure 1-28 on page 46.

The LS42 type 7902 is available in two form factors:

- A single slot blade server with two CPU sockets (LS42 base unit)
- A double slot blade server with four CPU sockets (LS42 base unit plus a Multi-Processor Expansion unit)

The single slot LS42 blade server models are fully functional two CPU socket blade servers. The single slot LS42 blade server models can also be upgraded to a double slot blade server, with four CPU sockets, by installing a Multi Processor Expansion (MPE) unit. The MPE includes two processors. All that is required for a fully functional four CPU socket blade server

is to install memory in the MPE. The single slot LS42 blade server offers investment protection by offering “pay as you grow” flexibility.

The double slot LS42 blade server models come standard with two processors installed on the base unit sockets. The included Multi Processor Expansion unit has two empty CPU sockets. Two additional processors and additional memory need to be ordered separately to make the LS42 a fully functional four CPU socket blade server.

The LS42 offers these features:

- Up to four high-performance, AMD quad-core Opteron processors.
- Up to two internal SAS or SATA hard disk drives.
- 8 or 16 DIMM sockets with up to 64 GB maximum memory.
- Two Gigabit Ethernet ports (Broadcom 5709S) standard with TOE support. Additional two Gigabit Ethernet (Broadcom 5709S) ports in a four-socket LS42.
- An internal USB connector.
- Expansion capabilities for Fibre Channel, Ethernet, Myrinet and InfiniBand.
- Support for concurrent KVM (cKVM) and concurrent USB and DVD (cMedia) via Advanced Management Module and optional daughter card.
- Support for the PCI Express I/O Expansion Unit (PEU3e) and PCI Expansion Unit II (PEU2).
- Three-year, onsite limited warranty.

Figure 1-28 shows the IBM BladeCenter LS42 and MPE unit.

Figure 1-28 IBM BladeCenter LS42 and MPE unit

Table 1-17 lists the LS42 features.

Table 1-17 Features of the LS42

	LS42 single slot (30mm)	LS42 double slot (60mm)
LS42 Model	1Sx, 3Sx	3Qx, CQx
Form Factor	30 mm blade (expandable to 60 mm using MPE unit)	60 mm blade
Processors (Std/Max)	2/2[a]	2/4
Upgradeable with MPE	Yes	No
Memory (Std/Max)	4 GB/32 GB[b]	4 GB/64 GB
DIMM sockets	8[c]	16
Gigabit Ethernet Controllers Standard	1 x Broadcom BCM5709S[d]	2 x Broadcom BCM5709S
2.5-inch drive bays	2	2
Adapter slots total	1 PCI-X slot[e] or 1 PCI Express slot	2 PCI-X slots or 1 PCI Express slot

a. With the addition of the MPE, the maximum number of processors installable is four.
b. With the addition of the MPE, the maximum memory installable is 64 GB.
c. With the addition of the MPE, the total DIMM sockets is 16.
d. With the addition of the MPE, the total number of Ethernet controllers is two.
e. With the addition of the MPE, the total number of PCI-X adapter slots is two.

The LS42 supports the Windows and Linux operating systems. See the following ServerProven page for further supported NOS information:

http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/ematrix.shtml

For more details about the server, see 3.10, “LS42 type 7902” on page 319.

1.3.11 BladeCenter LS21

The BladeCenter LS21 servers are high-throughput, two-socket, SMP-capable, AMD Opteron-based blade servers. This two-socket AMD Opteron blade server is well suited to HPC and other applications requiring high performance and high-speed fabrics.

The processors used in these blades are standard- and low-power, full-performance Opteron processors. The standard AMD Opteron processors draw a maximum of 95 W. Specially manufactured low-power processors operate at 68 W or less without any performance trade-offs. This savings in power at the processor level combined with the smarter power solution that BladeCenter delivers make these blades very attractive to clients who have limited power and cooling resources. The IBM BladeCenter LS21 is shown in Figure 1-29 on page 48.

The BladeCenter LS21 is a single-width blade server that offers these features:

- Up to two high-performance, AMD Dual-Core Opteron processors
- System board containing eight DIMM connectors, supporting 512 MB, 1 GB, 2 GB, or 4 GB DIMMs. Up to 32 GB of system memory is supported with 4 GB DIMMs.
- SAS controller, supporting one internal SAS drive (36 or 73 GB) and up to three additional SAS drives with optional SIO blade.

- Two TCP/IP Offload Engine-enabled (TOE) Gigabit Ethernet controllers (Broadcom 5706S) standard, with load balancing and failover features
- Support for concurrent KVM (cKVM) and concurrent USB/DVD (cMedia) via Advanced Management Module and optional daughter card
- Support for Storage and I/O Expansion (SIO) Unit
- Three-year, onsite limited warranty

Figure 1-29 IBM BladeCenter LS21

The LS21 supports the Windows and Linux operating systems. See the following ServerProven page for further supported NOS information:

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

Table 1-18 Features of the LS21

Feature	Specification
Processor	AMD Opteron Rev F Model 2212HE, 2216HE, 2218, 2218HE, and 2220
Number of processors (std/max)	1/2
Cache	1 MB L2 per processor core
Memory	8 VLP DIMM slots/DDR2 667/32 GB maximum
Internal hard disk drives (standard/maximum)	0/1
Maximum internal storage	On board: one 2.5” Non Hot Swap SAS HDD Optional: SIO blade offers support for 3 additional 2.5” hot-swap SAS drives
Network	Two ports, Integrated Dual Gigabit Ethernet (Broadcom 5706S) TOE
I/O upgrade	1 PCI-X expansion connector and 1 PCI-Express expansion connector

For more details about the server, see 3.7, “LS21 type 7971” on page 300.

1.3.12 BladeCenter LS41

The BladeCenter LS41 blade servers are scalable AMD Opteron-based servers. The LS41 is a 2-way-capable offering that can be expanded to become a four-way-capable offering. It is optimized for scalable enterprise workloads and scientific computing.

The processors used in these blades are standard and low power, full performance Opteron processors. The standard AMD Opteron processors draw a maximum of 95 W. Specially manufactured low-power processors operate at 68 W or less without any performance trade-offs. This savings in power at the processor level combined with the smarter power solution that BladeCenter delivers make these blades very attractive to clients who have limited power and cooling resources. The IBM BladeCenter LS41 is shown in Figure 1-30 on page 50.

The LS41 blade server can be purchased in one of two ways:

- As a two-socket, single-slot (30 mm) blade server that can be upgraded to a four-socket dual-slot (60 mm) server at a later date, with the addition of a multi processor expansion (MPE) unit. The LS41 main unit and the MPE are joined by simply clipping the two units together and tightening a thumb screw.

  The MPE unit does not include any processors or memory as standard. These need to be added separately. The processors will need to match the existing processors on the LS41 main unit. The memory should be split evenly between the two units.
- As a four-socket, dual-slot blade server (LS41 base unit with the multi processor expansion unit attached). These double-slot LS41 models come standard with two processors.

The BladeCenter LS41 offers these features:

- Up to four high-performance, AMD Dual-Core Opteron processors
- Up to two internal SAS hard disk drives
- 8 or 16 DIMM sockets with up to 64 GB maximum memory.
- Two Gigabit Ethernet ports (Broadcom 5706S) standard with TOE support. Additional two Gigabit Ethernet (Broadcom 5708S) ports in a four-socket LS41
- Expansion capabilities for Fibre Channel, Ethernet, Myrinet and InfiniBand
- Support for concurrent KVM (cKVM) and concurrent USB and DVD (cMedia) via Advanced Management Module and optional daughter card
- Support for Storage and I/O Expansion (SIO) Unit
- Three-year, onsite limited warranty

Figure 1-30 IBM BladeCenter LS41

Table 1-19 lists the LS41 features.

Table 1-19 Features of the LS41

Feature	LS41 single slot (30 mm)	LS41 double slot (60 mm)
Model number[a]	7972 31Y, 51Y, 61Y, 6AY, 7AY	7972 32Y, 52Y, 62Y, 6BY, 7BY
Processor	AMD Opteron Rev F Model 8212HE, 8216, 8216HE, 8218, 8218HE and 8220	
Multi Processor Expansion unit	Optional	Standard
Number of processors (standard/maximum)	1/2 (upgrade to 4 by adding the MPE unit)	2/4
Cache	1 MB L2 per processor core	
Memory	8 VLP DIMM slots/DDR2 667/32 Gb max	16 VLP DIMM slots/DDR2 667/64 GB max
Internal hard disk drives (std/max)	0/1	0/2
Maximum internal storage	one 2.5” Non Hot Swap SAS HDD	two 2.5” Non Hot Swap SAS HDD
	Optional: SIO blade offers support for 3 additional 2.5” hot-swap SAS drives	
Network	Two ports, Integrated Dual Gigabit Ethernet (Broadcom 5706S) TOE	Four ports, Integrated Dual Gigabit Ethernet (Broadcom 5706S and Broadcom 5708S) TOE
I/O upgrade	1 PCI-X expansion connector and 1 PCI-Express expansion connector	2 PCI-X expansion connector and 1 PCI-Express expansion connector

a. Model numbers provided are worldwide generally available variant (GAV) model numbers which are not orderable as listed, and need to be modified by country. Substitute the last character with a country-specific designator.

The LS41 supports the Windows and Linux operating systems. See the following ServerProven page for further supported NOS information:

http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/ematrix.shtml

For more details about the server, see 3.8, “LS41 type 7972” on page 306.

1.3.13 BladeCenter QS21

The IBM BladeCenter QS21 is based on the innovative Cell Broadband Engine™ (Cell/B.E.™) processor, providing a new level of parallelism and performance to targeted workloads. Figure 1-31 shows the IBM BladeCenter QS21.

The Cell/B.E. processor chip uses a multicore architecture optimized for high-performance computing and media-rich applications. The new QS21 Cell/B.E. technology-based system is designed for businesses that need the dense computing power and unique capabilities of the Cell/B.E. processor to tackle tasks involving graphic-intensive, numeric applications.

The QS21 is particularly well suited for image processing, video processing, and signal processing applications in Aerospace and Defense, Medical Imaging, EDA, Digital Video Surveillance (DVS), Seismic, Financial markets, Digital Media industries.

The system will rely on the Cell/B.E. processor to help accelerate key algorithms like 3D rendering, compression, and encryption, to help enable companies to create and run highly visual, immersive, real-time applications.

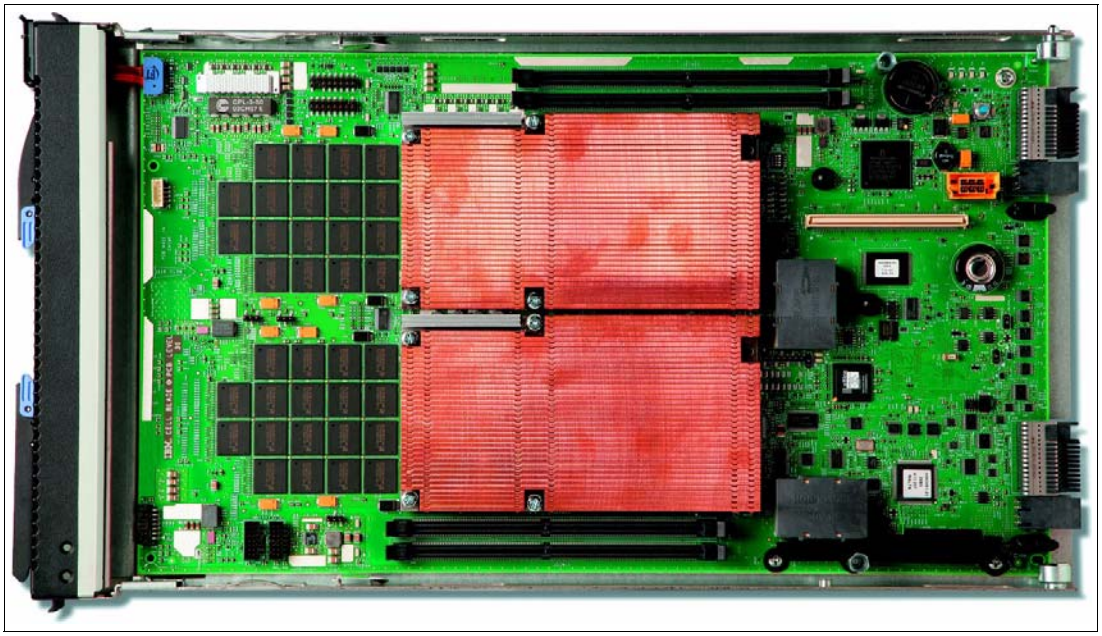

Figure 1-31 IBM BladeCenter QS21

Following is a list of the BladeCenter QS21 features:

- Two 3.2 GHz Cell/B.E. processors
- 2 GB XDR memory
- Integrated dual GB Ethernet
- IBM Enhanced I/O Bridge Chip
- Optional two or four 512 MB VLP DIMMs as I/O buffer
- Optional InfiniBand adapter
- Serial over LAN
- Single-wide blade

- Supported in the BladeCenter H and HT chassis
- Support by IBM SDK for Multicore Acceleration v3.0

The supported operating system for the QS21 is Red Hat Enterprise Linux, RHEL 5u1.

See the following ServerProven page for further supported NOS information:

http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/ematrix.shtml

For more information about QS21, refer to 3.14, “QS21 type 0792” on page 345.

1.3.14 BladeCenter QS22

The BladeCenter QS22 is based on the second-generation processor of the Cell Broadband Engine Architecture (Cell/B.E), and offers flexibility along with strong parallelization. It performs double precision floating point operations at five times the speed of previous generations QS20 and QS21.

Because of its parallel nature and quick computing speed, the BladeCenter QS22 is ideal for High Performance Computing (HPC) solutions. Its PowerXCell 8i multi-core processor architecture helps it accelerate 3D rending, compression, and encryption algorithms.

The QS22 is shown in Figure 1-32.

Figure 1-32 IBM BladeCenter QS22

Features of the QS22 include:

- Two 3.2 GHz IBM PowerXCell™ 8i processors
- Up to 32 GB of PC2-6400 800 MHz DDR2 memory
- 460 single-precision GFlops per blade (peak)

- 217 double-precision GFlops per blade (peak)
- Integrated dual 1Gb Ethernet
- IBM Enhance I/O Bridge Chip
- Serial Over LAN

The QS22 is a single-wide blade server. Operating system support at the time of writing is RHEL 5.2 or later.

See the following ServerProven page for further supported NOS information:

http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/ematrix.shtml

For more details about the server, see 3.15, “QS22 type 0793” on page 348.

1.3.15 BladeCenter PN41

The IBM BladeCenter PN41 Deep Packet Inspection (DPI®) Blade is a content-processing blade. DPI looks beyond addresses and ports into packets to help deliver visibility, service control, and security of network traffic. Additionally, PN41 offers a development environment using Eclipse IDE that enables rapid application development.

BladeCenter PN41 is a highly programmable high-performance blade that helps to improve detection and management of network traffic; enable advanced visibility, service management, and security of network traffic; and provide improved protection of mission-critical network infrastructure. These are key concerns in today's telecommunications and governmental environments.

This blade combines the powerful capabilities of BladeCenter with the performance of CloudShield, which is an industry-leading network service control and security provider. This combination creates an optimal platform for building and operating next-generation network service management solutions.

BladeCenter PN41 is shown in Figure 1-33 on page 54.

Figure 1-33 IBM BladeCenter PN41

The key features of the PN41 are:

- Four internal 1 Gbps Ethernet ports for management
- Four internal 10 Gbps Ethernet port for real-time data flow analysis
- One external SFP-based 1 Gbps Ethernet port for capturing data
- One external XFP-based 10 Gbps Ethernet port for real-time data flow analysis
- Eclipse IDE for rapid application development

With the IBM BladeCenter PN41, organizations can gain more insight into their network traffic to save costs. Specifically, network operators can verify that their own traffic is fast and efficient while concurrently identifying and shaping peer-to-peer traffic that may be draining resources. The IBM BladeCenter PN41 helps manage this process by delivering advanced traffic visibility. Using this visibility, organizations can recognize where low quality of service or limited bandwidth potentially threatens customer satisfaction, and then actively address this situation.

CloudShield sells a key application, known as Subscriber Services Manager, that addresses the problem of increased bandwidth used by peer-to-peer applications. These applications are rarely revenue generators. In fact, they often steal valuable bandwidth and performance from latency-sensitive applications like VoIP, VoD, and online gaming applications. The Subscriber Services Manager application detects and manages traffic from today's P2P protocols and applications to help carriers control bandwidth-hungry applications and deliver high quality of service to their customers.

Data security is an issue that every information-based organization must deal with. Internal and external attacks, which are executed by everything from hackers to automated bots, can compromise network infrastructure and render a business virtually helpless. Thus, enterprises have little choice but to secure their networks. The IBM BladeCenter PN41 delivers outstanding security to help protect a network infrastructure.

For more information about PN41, refer to 3.16, “PN41 type 3020” on page 353.

2

BladeCenter chassis and infrastructure configuration

This chapter contains information about different infrastructure options available for IBM BladeCenter such as management, power, cooling, network, and storage connectivity modules. It includes the following topics:

Note about feature code: All I/O modules can have up to three different feature codes, depending on the chassis the module is installed in, and the sales channel the chassis is ordered through:

- BladeCenter E, T, H, and HT configurations ordered through the System x sales channel
- BladeCenter S configurations ordered through the System x sales channel
- BladeCenter S and H ordered through the Power Systems sales channel

In this chapter, when we list a feature code, it is the one for BladeCenter E, T, H, and HT configurations ordered through the System x sales channel. The only exception to this is the SAS RAID Controller Module, which is only supported in BladeCenter S. Table 1-2 on page 3 lists these feature codes.

2.1 I/O architecture

Each BladeCenter chassis can support different types and quantities of I/O modules, and each blade server in the chassis can be connected to the I/O modules by a number of ways. This section gives you an overview of internal I/O architecture and supported combinations of I/O modules and expansion cards for each chassis.

- 2.1.1, “BladeCenter E I/O topology” on page 59
- 2.1.2, “BladeCenter H I/O topology” on page 62
- 2.1.3, “BladeCenter T I/O topology” on page 66
- 2.1.4, “BladeCenter HT I/O topology” on page 67
- 2.1.5, “BladeCenter S I/O topology” on page 72

I/O modules can be grouped into several categories (or types):

- Ethernet switch modules (which include both standard and high-speed Ethernet switch modules) are covered in 2.2, “Ethernet switch modules” on page 76.
- InfiniBand switch modules (which include high-speed InfiniBand switch modules as well as InfiniBand bridge modules) are covered in 2.3, “InfiniBand switch modules” on page 137.
- Fibre Channel switch modules (which include standard FC switch modules) are covered in 2.4, “Fibre Channel switch modules” on page 144.
- SAS I/O modules are covered in 2.5, “SAS I/O modules” on page 164.
- Pass-thru and interconnect modules (which include copper and optical pass-thru modules, 4X InfiniBand Pass-thru Module, and multi-switch interconnect modules) are covered in 2.6, “Pass-thru and interconnect modules” on page 169.

We use the following abbreviations when referring to the specific I/O module or category of I/O modules in the compatibility tables later in this section:

- ESM - Standard Ethernet switch modules
  - Cisco Systems Fiber Intelligent Gigabit ESM (32R1888, feature code 1497)
  - Cisco Systems Intelligent Gigabit ESM (32R1892, feature code 1498)
  - Cisco Catalyst Switch Module 3110G (41Y8523, feature code 2989)
  - Cisco Catalyst Switch Module 3110X (41Y8522, feature code 2988)
  - Cisco Catalyst Switch Module 3012 (43W4395, feature code 5450)
  - IBM Server Connectivity Module (39Y9324, feature code 1484)
  - Nortel Networks L2/3 Copper Gigabit ESM (32R1860, feature code 1495)
  - Nortel Networks L2/3 Fiber Gigabit ESM (32R1861, feature code 1496)
  - Nortel Networks Layer 2-7 Gigabit ESM (32R1859, feature code 1494)
  - Nortel 10 Gb Uplink Ethernet Switch Module (32R1783, feature code 1493)
  - Nortel 1/10Gb Uplink Ethernet Switch Module (44W4404, feature code 1590)
- HSESM - High-speed Ethernet switch modules
  - Nortel 10 Gb Ethernet Switch Module (39Y9267, feature code 2952)
- HSIBSM - High-speed InfiniBand switch module
  - Cisco Systems 4X InfiniBand Switch Module (32R1756, feature code 1574)
- IBBM - InfiniBand bridge modules
  - QLogic InfiniBand Ethernet Bridge Module (39Y9207, feature code 2941)
  - QLogic InfiniBand Fibre Channel Bridge Module (39Y9211, feature code 2942)

- FCSM - Standard Fibre Channel switch modules
  - Brocade 20-port SAN Switch Module (32R1812, feature code 1569)
  - Brocade 10-port SAN Switch Module (32R1813, feature code 1571)
  - Cisco 20-port 4 Gb Fibre Channel Switch Module (39Y9280, feature code 2983)
  - Cisco 10-port 4 Gb Fibre Channel Switch Module (39Y9284, feature code 2984)
  - QLogic 20-port 4 Gb Fibre Channel Switch Module (26R0881, feature code 1560)
  - QLogic 20-port 4 Gb SAN Switch Module (43W6725, feature code 2987)
  - QLogic 10-port 4 Gb SAN Switch Module (43W6724, feature code 2986)
  - QLogic 4 Gb Intelligent Pass-thru Module (43W6723, feature code 2985)
  - QLogic 20-Port 8 Gb SAN Switch Module (44X1905, feature code 5478)
  - QLogic 8 Gb Intelligent Pass-thru Module (44X1907, feature code 5482)
- SASM - standard SAS I/O modules
  - SAS Connectivity Module (39Y9195, feature code 2980)
  - SAS RAID Controller Module (43W3584, feature code 3734)
- CPM - Copper Pass-thru Module (39Y9320, feature code 2900)
- ICPM - Intelligent Copper Pass-thru Module (44W4483, feature code 5452)
- OPM - Optical Pass-thru Module (39Y9316, feature code 1556)
- HSIBPM - 4X InfiniBand Pass-thru Module (43W4419, feature code 2990)
- MSIM - Multi-switch Interconnect Module (39Y9314, feature code 1465)
- MSIM-HT - Multi-switch Interconnect Module for BladeCenter HT (44R5913, feature code 5491)

2.1.1 BladeCenter E I/O topology

The IBM BladeCenter E chassis provides connection paths between server blade bays and I/O or switch bays (see Figure 1-1 on page 12 and Figure 1-3 on page 13) through a hard-wired dual redundant midplane.

The I/O topology is shown in Figure 2-1. Each blade bay has four dedicated I/O connections (two per midplane) linked to four I/O bays (one blade bay connection to one I/O bay). Thus, each I/O bay has a total of 14 I/O connections (to 14 blade bays).

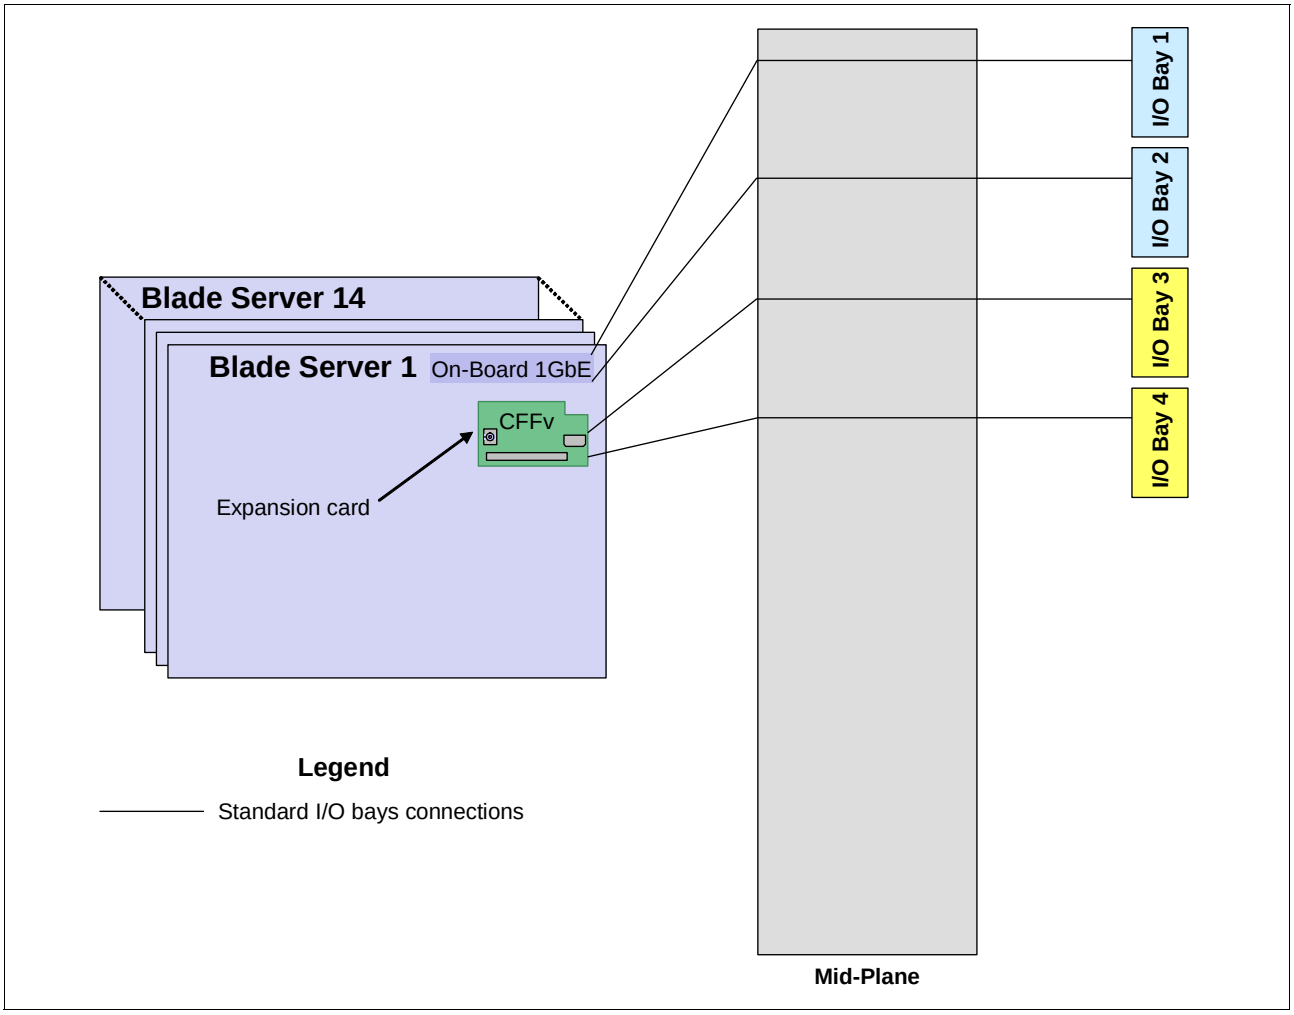

Figure 2-1 IBM BladeCenter E I/O topology

Each blade server has at least two Ethernet ports (Ethernet 1 and Ethernet 2). They are routed to I/O bay 1 and I/O bay 2 respectively, which means that only Ethernet-capable I/O modules can be used in I/O bays 1 and 2. I/O expansion cards that you optionally install in blade servers have two ports (except for the Myrinet Cluster Expansion Card which has only one port) which are usually routed to I/O bays 3 and 4. Thus, depending on the expansion card installed, compatible I/O modules must be chosen for those I/O bays. Unless an expansion card is installed in one or more processor blades, there is no need for I/O modules 3 and 4.

The primary purpose of having redundant midplane and redundant blade connections (that is, dual-port Gigabit Ethernet or dual-port FC) is to avoid a single point of failure for infrastructure by providing redundancy of key components such as links and network and storage switches. Redundant connections can also be used for load balancing purposes, thus increasing system performance.

The I/O modules must be compatible with the I/O interfaces present in the blade servers. For example, when a Fibre Channel expansion card is installed in a blade server, I/O modules 3 and 4 must also be Fibre Channel-based (that is, a FC switch module or Optical Pass-thru

module, OPM), and vice versa, if you install FC switches in bays 3 and 4, then any expansion cards installed in all *other* blade servers in the same chassis must be Fibre Channel.

There is one exception to this rule: if you are using an Optical Passthru Module (OPM), you can have both Ethernet and FC expansion cards for different blades in the same chassis because OPM is compatible with Ethernet and FC expansion cards.

Note: It is possible to use I/O modules of the same type (for example, Ethernet) but from different vendors in the adjacent pair (pair of bays 1 and 2 or pair of bays 3 and 4). However, we recommend that you use I/O modules from the same vendor in the adjacent pair to avoid external interoperability issues (in case of Fibre Channel) or management complexity, or both.

See Table 2-1 for list of supported combinations of I/O modules and I/O expansion cards for IBM BladeCenter E chassis.

Table 2-1 Supported combinations of I/O modules and I/O expansion cards - IBM BladeCenter E

Part Number	Feature code	Expansion card	ESM, CPM, ICPM		FCSM	SASM[a]	OPM	
		I/O module bay number ♦	1, 2	3, 4	3, 4	3, 4	1, 2	3, 4
None	None	Integrated Gigabit Ethernet	Yes	No	No	No	Yes	No
39R8624	1548	SFF Gigabit Ethernet Expansion Card	Yes[b]	Yes	No	No	Yes[b]	Yes
39Y9310	2969	Ethernet Expansion Card (CFFv)	Yes[b]	Yes	No	No	Yes[b]	Yes
26K4841	None	IBM SFF FC Expansion Card	No	No	Yes	No	No	Yes
26R0890	1577	QLogic 4 Gb SFF FC Expansion Card	No	No	Yes	No	No	Yes
26R0884	None	QLogic 4 Gb StFF FC Expansion Card	No	No	Yes	No	No	Yes
41Y8527	2970	QLogic 4 Gb FC Expansion Card (CFFv)	No	No	Yes	No	No	Yes
39Y9186	2925	Emulex 4 Gb SFF FC Expansion Card	No	No	Yes	No	No	Yes
43W6859	2994	Emulex 4 Gb FC Expansion Card (CFFv)	No	No	Yes	No	No	Yes
32R1923	1458	QLogic iSCSI Expansion Card	No	Yes	No	No	No	Yes
39Y9190	2979	SAS Expansion Card (CFFv)	No	No	No	Yes	No	No
43W3974	1591	SAS Connectivity Card (CFFv)	No	No	No	Yes	No	No
46C7167	5490	ServeRAID-MR10ie (CIOv)	No	No	No	Yes[c]	No	No

a. Only the SAS Connectivity Module (39Y9195) is meant under the SASM abbreviation in this table because SAS RAID Controller Module (43W3584) is not supported in the BladeCenter E chassis.
b. Supported only if the expansion card is installed in slot 1 of either the BladeCenter SCSI Storage Expansion Unit II (39R8625) or BladeCenter Storage and I/O Expansion Blade (39R7563, feature code 4265).
c. This card only supports SAS Connectivity Module (39Y9195, feature code 2980). Currently, only HS12 (8028) supports this card, and HS12 requires the SAS Connectivity Card (43W3974, feature code 1591) to be installed into it to support external attachments.

Look at the top data row, Integrated Gigabit Ethernet. Table 2-1 shows that onboard Ethernet ports on each blade are always routed to bays 1 and 2 of the BladeCenter E chassis, and only Ethernet-capable modules can be used in those bays; that is, Ethernet switches or Optical Pass-Thru modules.

As another example, we are using a QLogic 4 Gb SFF FC Expansion card in our blades. The row in the table for this card shows that ports of this card are always routed to bays 3 and 4. Therefore, we have to use FC-capable modules in those bays; that is, FC switch or Optical Pass-Thru modules.

Refer to Table 1-3 on page 5 for a compatibility information for blades and supported expansion options.

2.1.2 BladeCenter H I/O topology

IBM BladeCenter H chassis has two types of fabrics inside:

- Standard fabric (which is almost the same as in the BladeCenter E chassis)

- High-speed fabric, which is capable of carrying 10 Gbps Ethernet 4X InfiniBand, and even higher-bandwidth technologies.

The high-speed fabric is only used when you install a high-speed expansion card into a blade server. This card has its own connectors to the midplane and a PCI Express socket on the blade itself. For information about servers that support these cards, see Table 1-3 on page 5.

The IBM BladeCenter H chassis has a total of ten I/O bays (as shown in Figure 1-8 on page 18). Each blade bay has a total of eight dedicated connection paths to the I/O modules (see Figure 2-2).

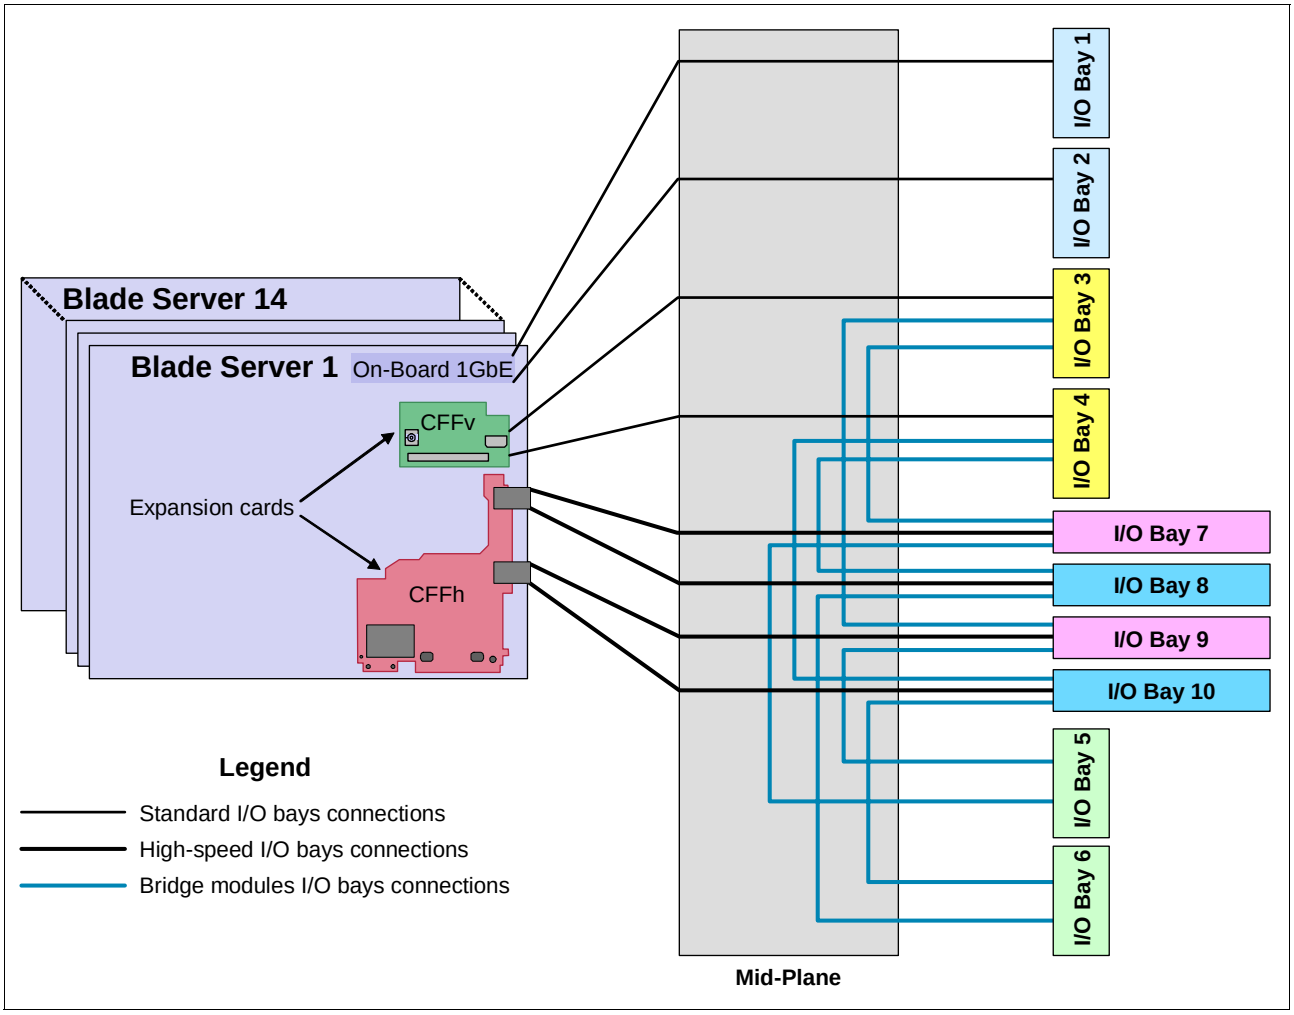

Figure 2-2 IBM BladeCenter H I/O topology

The bays are as follows:

- Bays 1 and 2 only support standard Ethernet-compatible I/O modules. These are routed internally to the onboard Ethernet controllers on the blades (and slot 1 in the SIO and BSE2 expansion blades).
- Bays 3 and 4 can be used either for standard switch or pass-thru modules (such as Fibre Channel connectivity or additional Ethernet ports) or for bridge modules. These are routed internally to the PCI-X connector on the blades.
- Bays 5 and 6 are dedicated for bridge modules only and do not directly connect to the blade bays. Bridge modules provide links to the I/O bays 7 through 10 and can be used as

additional outputs for I/O modules in those bays. In case I/O bays 3 and 4 are used for bridge modules, they are not directly connected to the blades, and bay 3 provides redundancy for bay 5, and bay 4 provides redundancy for bay 6.

- I/O bays 7 through 10 are used for high-speed switch modules such as a Cisco 4X InfiniBand Switch Module. These are routed internally through midplane connectors on certain expansion cards to the PCI Express connector on blades that have it.

Be sure that all installed I/O modules are compatible with interface ports present in the blades. For a list of supported combinations of standard I/O modules and I/O expansion cards for BladeCenter H, refer to Table 2-2 on page 65. For a list of supported combinations of high-speed modules and high-speed expansion cards for BladeCenter H, see Table 2-3 on page 65.

Table 2-2 Supported combinations of I/O modules and I/O expansion cards - IBM BladeCenter H

Part Number	Feature code	Expansion card	ESM, CPM, ICPM		FCSM	SASM[a]	OPM	
		I/O module bay number ♦	**1, 2**	**3, 4**	**3, 4**	**3, 4**	**1, 2**	**3, 4**
None	None	Integrated Gigabit Ethernet	Yes	No	No	No	Yes	No
39R8624	1548	SFF Gigabit Ethernet Expansion Card	Yes[b]	Yes	No	No	Yes[b]	Yes
39Y9310	2969	Ethernet Expansion Card (CFFv)	Yes[b]	Yes	No	No	Yes[b]	Yes
26K4841	None	IBM SFF FC Expansion Card	No	No	Yes	No	No	Yes
26R0890	1577	QLogic 4 Gb SFF FC Expansion Card	No	No	Yes	No	No	Yes
26R0884	None	QLogic 4 Gb StFF FC Expansion Card	No	No	Yes	No	No	Yes
41Y8527	2970	QLogic 4 Gb FC Expansion Card (CFFv)	No	No	Yes	No	No	Yes
39Y9186	2925	Emulex 4 Gb SFF FC Expansion Card	No	No	Yes	No	No	Yes
43W6859	2994	Emulex 4 Gb FC Expansion Card (CFFv)	No	No	Yes	No	No	Yes
32R1923	1458	QLogic iSCSI Expansion Card	No	Yes	No	No	No	Yes
39Y9190	2979	SAS Expansion Card (CFFv)	No	No	No	Yes	No	No
43W3974	1591	SAS Connectivity Card (CFFv)	No	No	No	Yes	No	No
46C7167	5490	ServeRAID-MR10ie (CIOv)	No	No	No	Yes[c]	No	No

a. Only the SAS Connectivity Module (39Y9195) is meant under the SASM abbreviation in this table because SAS RAID Controller Module (43W3584) is not supported in the BladeCenter H chassis.
b. Supported only if the expansion card is installed in slot 1 of either the BladeCenter SCSI Storage Expansion Unit II (39R8625) or BladeCenter Storage and I/O Expansion Unit (39R7563, feature code 4265).
c. This card only supports SAS Connectivity Module (39Y9195, feature code 2980). Currently, only HS12 (8028) supports this card, and HS12 requires SAS Connectivity Card (43W3974, feature code 1591) to be installed into it to support external attachments.

Table 2-3 Supported combinations of high-speed I/O modules and I/O expansion cards - BladeCenter H

Part Number	Feature Code	High-speed expansion card	ESM ICPM CPM[a,b]	FC SM[a]	OPM[a]	HS IBPM	HS IBSM	HS ESM
		I/O module bay number ♦	**7, 8, 9, 10[c]**	**8, 10[c]**	**7, 8, 9, 10[c]**	**7/8, 9/10[d]**	**7, 9**	**7, 8, 9, 10**
39Y9271	2967	NetXen 10 Gb Eth. Exp. Card (CFFh)	No	No	No	No	No	Yes[e]
39Y9306	2968	QLogic Eth. & 4 Gb FC Exp. Crd (CFFh)	Yes[e]	Yes	Yes	No	No	No
43W4423	2991	4X InfiniBand DDR Exp. Card (CFFh)	No	No	No	Yes	Yes	No
43W4420	2993	Voltaire 4X IB DDR Exp. Card (CFFh)	No	No	No	Yes	No	No
44W4479	5476	2/4 Port Ethernet Exp. Card (CFFh)	Yes[f]	No	Yes	No	No	No
44W4465	5479	Broadcom 10 Gb 4-port Ethernet (CFFh)	No	No	No	No	No	Yes[f,g]
44W4465	5489	Broadcom 10 Gb 2-port Ethernet (CFFh)	No	No	No	No	No	Yes[e]
44X1940	5485	QLogic Ethernet & 8 Gb FC (CFFh)	Yes[e,h]	Yes	Yes	No	No	No

a. Supported in high-speed I/O bays using MSIM (39Y9314, feature code 1465).
b. Nortel Layer 2/7 switch (32R1859, feature code 1494) is not supported in MSIM.

c. These high-speed I/O bays are converted to the standard I/O bays using MSIM. In this case, bay 7 represents the upper left bay of MSIM. Bay 8 represents the upper right bay. Bay 9 represents the lower left bay, and bay 10 represents the lower right bay.
d. One 4X InfiniBand Pass-thru module (HSIBPM) occupies two adjacent high-speed I/O bays.
e. This card has two Ethernet ports that are connected to the bays 7 and 9 of MSIM.
f. This card has four Ethernet ports that are connected to the bays 7, 8, 9, and 10 of MSIM.
g. When this card is installed into Storage and I/O Expansion Unit (39R7563, feature code 4265), then only two 10 Gb Ethernet ports can be utilized.
h. When this card is installed into Storage and I/O Expansion Unit (39R7563, feature code 4265), then only two 8 Gb FC ports can be utilized.

2.1.3 BladeCenter T I/O topology

IBM BladeCenter T chassis has eight blade bays and four I/O bays (as shown in Figure 1-11 on page 21). The architecture of I/O connections is the same as for IBM BladeCenter chassis, except that there are only eight blade bays. See Figure 2-3 for more information.

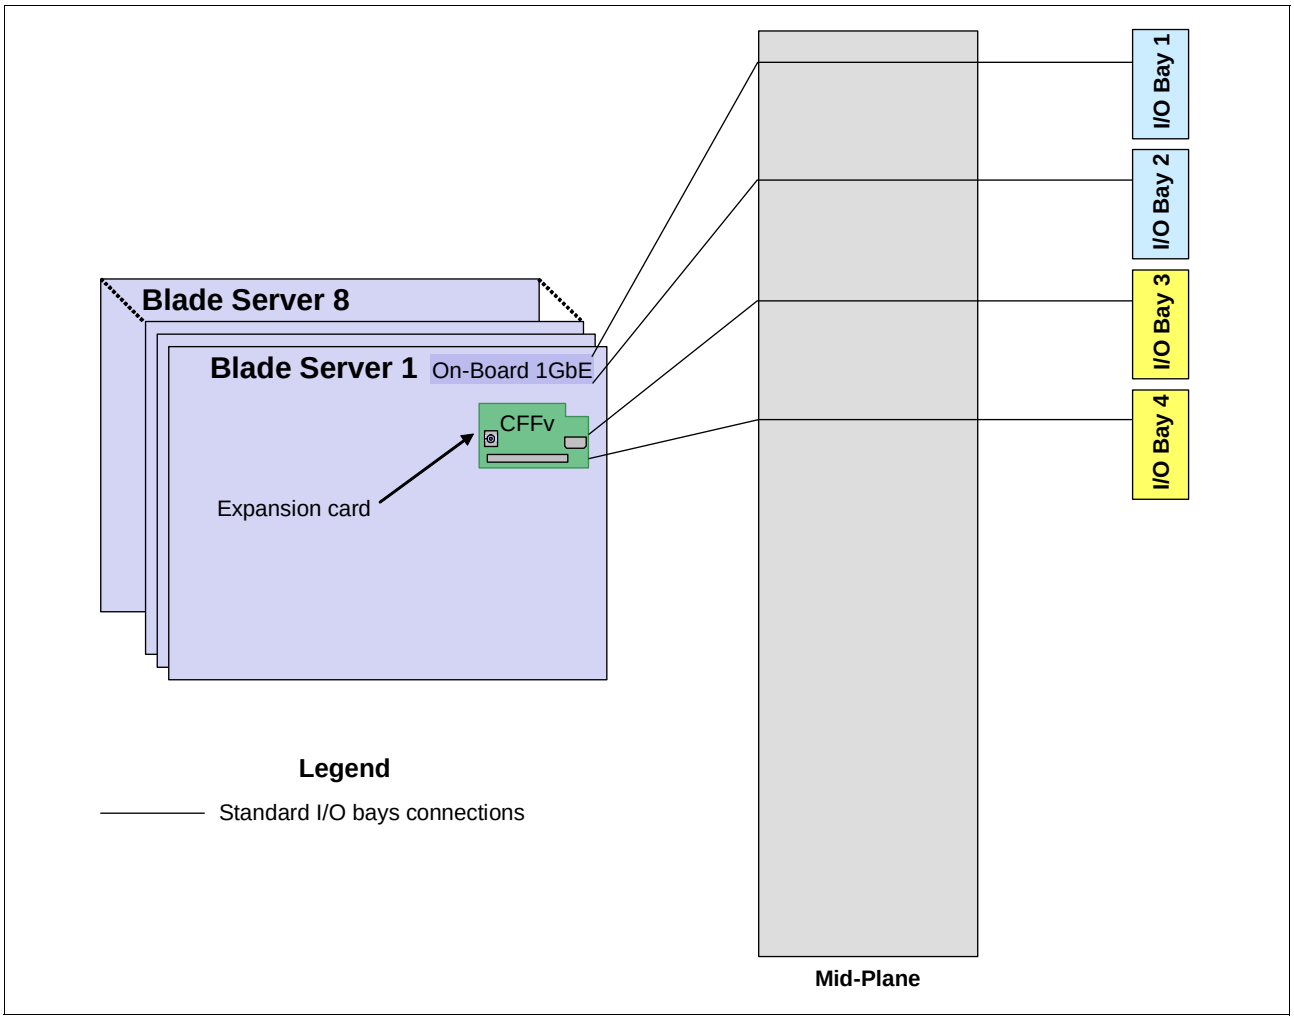

Figure 2-3 IBM BladeCenter T I/O topology

Be sure that all installed I/O modules are compatible with interface ports present in the blades. Refer to 2.1.1, “BladeCenter E I/O topology” on page 59 for a detailed description of path routing. See Table 2-4 on page 67 for information about supported configurations.

Note: If 4 Gb or 8 Gb Fibre Channel switch modules are used inside the BladeCenter T chassis, then internal connections between blades and switch modules operate at speeds up to 2 Gb, and external connections between switch modules and external FC switches or storage devices operate at speeds up to 4 Gb or 8 Gb respectively, depending on external FC devices capabilities.

Table 2-4 Supported combinations of I/O modules and I/O expansion cards - IBM BladeCenter T

Part Number	Feature code	Expansion card	ESM, CPM, ICPM		FCSM	SASM[a]	OPM	
		I/O module bay number ♦	1, 2	3, 4	3, 4	3, 4	1, 2	3, 4
None	None	Integrated Gigabit Ethernet	Yes	No	No	No	Yes	No
39R8624	1548	SFF Gigabit Ethernet Expansion Card	Yes[b]	Yes	No	No	Yes[b]	Yes
39Y9310	2969	Ethernet Expansion Card (CFFv)	Yes[b]	Yes	No	No	Yes[b]	Yes
26K4841	None	IBM SFF FC Expansion Card	No	No	Yes	No	No	Yes
26R0890	1577	QLogic 4 Gb SFF FC Expansion Card	No	No	Yes	No	No	Yes
26R0884	None	QLogic 4 Gb StFF FC Expansion Card	No	No	Yes	No	No	Yes
41Y8527	2970	QLogic 4 Gb FC Expansion Card (CFFv)	No	No	Yes	No	No	Yes
39Y9186	2925	Emulex 4 Gb SFF FC Expansion Card	No	No	Yes	No	No	Yes
43W6859	2994	Emulex 4 Gb FC Expansion Card (CFFv)	No	No	Yes	No	No	Yes
32R1923	1458	QLogic iSCSI Expansion Card	No	Yes	No	No	No	Yes
39Y9190	2979	SAS Expansion Card (CFFv)	No	No	No	Yes	No	No
43W3974	1591	SAS Connectivity Card (CFFv)	No	No	No	Yes	No	No
46C7167	5490	ServeRAID-MR10ie (CIOv)	No	No	No	Yes[c]	No	No

a. Only the SAS Connectivity Module (39Y9195) is meant under the SASM abbreviation in this table because SAS RAID Controller Module (43W3584) is not supported in the BladeCenter T chassis.
b. Supported only if the expansion card is installed in slot 1 of either the BladeCenter SCSI Storage Expansion Unit II (39R8625) or BladeCenter Storage and I/O Expansion Unit (39R7563, feature code 4265).
c. This card only supports SAS Connectivity Module (39Y9195, feature code 2980). Currently, only HS12 (8028) supports this card, and HS12 requires SAS Connectivity Card (43W3974, feature code 1591) to be installed into it to support external attachments.

2.1.4 BladeCenter HT I/O topology

IBM BladeCenter HT chassis like a BladeCenter H chassis has two types of fabrics inside:

- Standard fabric
- High-speed fabric

The high-speed fabric is only used when you install a high-speed expansion card into a blade server. This card has its own physical connectors to the midplane in addition to standard connectors located on the blade itself. For information about servers that support these cards, see Table 1-3 on page 5.

IBM BladeCenter HT chassis has a total of eight I/O bays (see Figure 1-8 on page 18). Each blade bay has a total of eight dedicated connection paths to the I/O modules (see Figure 2-4).

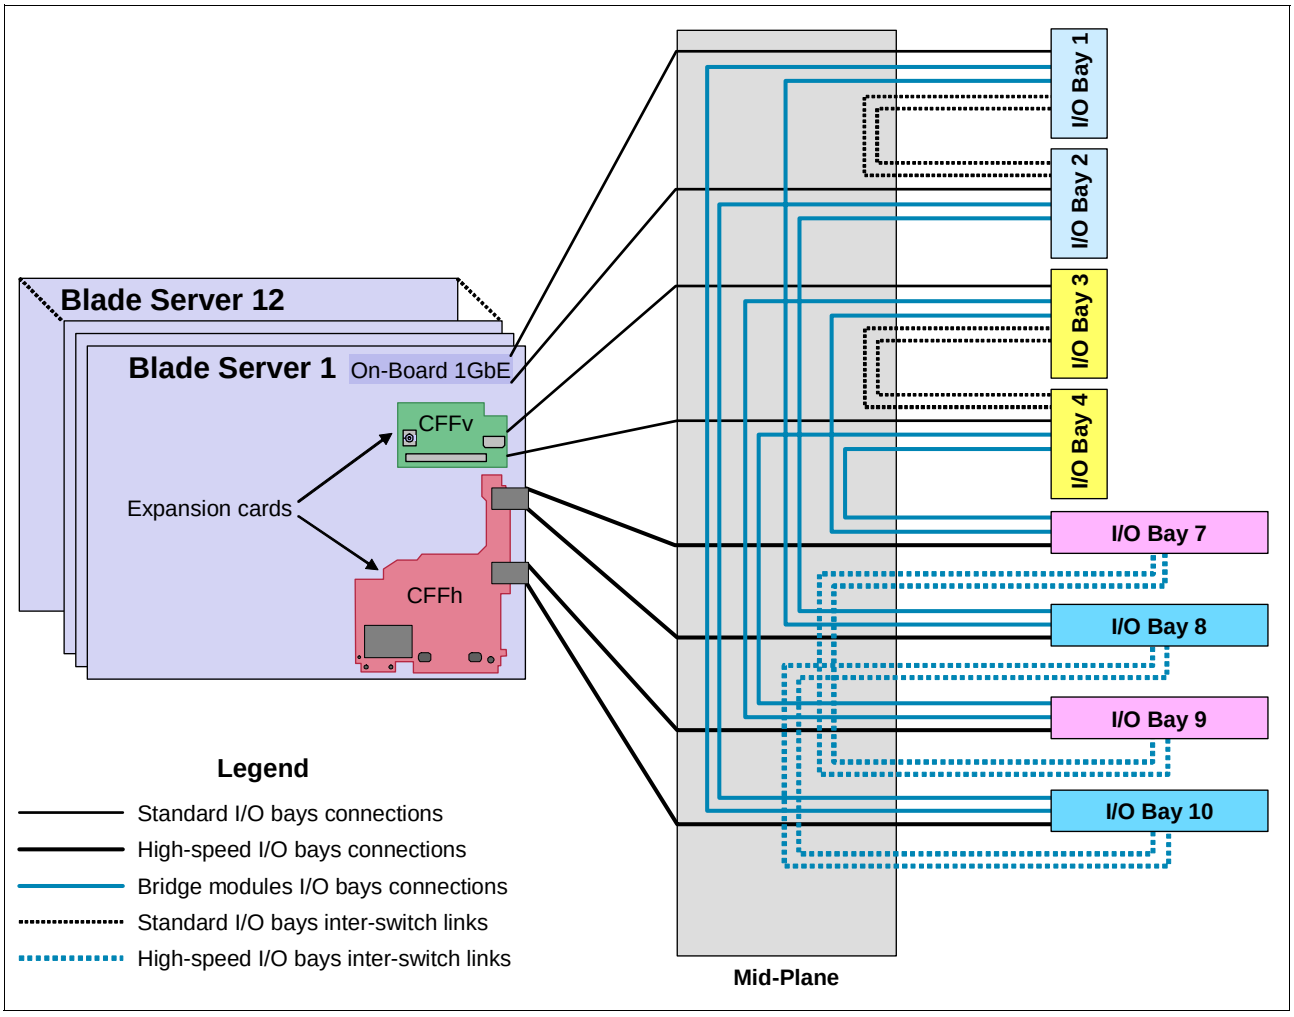

Figure 2-4 IBM BladeCenter HT I/O topology

The bays are as follows:

- Bays 1 and 2 support standard Ethernet-compatible I/O modules or bridge modules.
- Bays 3 and 4 can be used either for standard switch or pass-thru modules (such as Fibre Channel connectivity or additional Ethernet ports) or for bridge modules.
- I/O bays 7 through 10 are used for high speed switch modules such as the Cisco 4X InfiniBand Switch Module.
- I/O bays 1 and 2 are connected to high-speed I/O bays 8 and 10 in a redundant manner, and bays 3 and 4 are connected to bays 7 and 9 in a redundant manner (see Figure 2-4).
- I/O bay 1 is connected to I/O bay 2, bay 3 is connected to bay 4, bay 7 is connected to bay 9, and bay 8 is connected to bay 10. All of these connections have two links, and they can be used for interswitch communications if required (each switch module has 14 internal connections to the blades, while BladeCenter HT can hold up to 12 blades, so these two unused ports (ports 13 and 14) on the switch module can be utilized for interswitch links).

The use of any of I/O modules in the BladeCenter HT chassis requires additional interposers. See 2.11, “BladeCenter HT interposers” on page 209 for details.

The InterSwitch Links (ISLs) in BladeCenter HT take advantage of the unused internal ports 13 and 14 on I/O switch modules by internally connecting redundant switch pairs as follows:

- I/O module bay 1 port 13 to I/O module bay 2 port 13
- I/O module bay 1 port 14 to I/O module bay 2 port 14
- I/O module bay 3 port 13 to I/O module bay 4 port 13
- I/O module bay 3 port 14 to I/O module bay 4 port 14
- I/O module bay 7 port 13 to I/O module bay 9 port 13
- I/O module bay 7 port 14 to I/O module bay 9 port 14
- I/O module bay 8 port 13 to I/O module bay 10 port 13
- I/O module bay 8 port 14 to I/O module bay 10 port 14

The Inter-Switch Links allow the internal connection of pairs of Ethernet switches for the purposes of load balancing or network failover. The use of these internal ports frees up external ports that would otherwise be allocated to provide this functionality.

Inter-Switch Links are connected through the BladeCenter HT switch interposers to the backplane. High-speed I/O module bays 7-9 and 8-10 are always connected (only one type of interposer exists for high speed I/O bays, which provides ISL links). ISL connections for standard I/O module bays 1 and 2 and 3 and 4 are determined by the type of switch interposer installed in the bay behind the switch module. Refer to 2.11, “BladeCenter HT interposers” on page 209 for details about interposers.

Notes:

- Currently, only the Blade Network Technologies Ethernet switches support the InterSwitch Links in BladeCenter HT (also requires the appropriate interposer, part number 42C5301)
- The non-ISL interposer (part number 42C5300) should be used for all other switch modules, including Fibre Channel
- Both switch module bays in the redundant pair (bays 1 and 2 or bays 3 and 4) must have the same type of interposer (ISL or non-ISL)
- There is only one High Speed Switch Module interposer (part number 42C5302), which always connects ports 13 and 14 on HSSM bays.

See 2.11, “BladeCenter HT interposers” on page 209 for details about interposers.

Be sure that all installed I/O modules are compatible with interfaces present in the blades. For list of supported combinations of standard I/O modules and expansion cards for BladeCenter HT, refer to Table 2-5 on page 71. For a list of supported combinations of high-speed modules and high-speed expansion cards for BladeCenter HT, see Table 2-6 on page 72.

Fibre Channel speed restrictions:

- If 4 Gb Fibre Channel switch modules are installed into standard I/O bays (bays 3 and 4) of BladeCenter HT, then internal connections between blades and these switch modules operate at speeds up to 2 Gbps, and external connections between these switch modules and external FC switches or storage devices operate at speeds up to 4 Gbps, depending on the capabilities of the external FC devices connected.
- If 8 Gb Fibre Channel switch modules are installed into standard bays (bays 3 and 4) of a BladeCenter HT chassis, then internal connections between blades and these switch modules operate at speeds up to 4 Gbps, and external connections between these switch modules and external FC switches or storage devices operate at speeds up to 8 Gbps, depending on the capabilities of the external FC devices connected.
- If 4 Gb Fibre Channel switch modules are installed into MSIM-HT I/O bays (bays 8 and 10) of a BladeCenter HT chassis, then internal connections between blades and these switch modules operate at speeds up to 4 Gbps, and external connections between these switch modules and external FC switches or storage devices operate at speeds up to 4 Gbps, depending on the capabilities of the external FC devices connected.

Table 2-5 Supported combinations of I/O modules and I/O expansion cards - IBM BladeCenter HT

Part Number	Feature code	Expansion card	ESM, CPM, ICPM		FCSM	SASM[a]	OPM	
		I/O module bay number ♦	1, 2	3, 4	3, 4	3, 4	1, 2	3, 4
None	None	Integrated Gigabit Ethernet	Yes	No	No	No	Yes	No
39R8624	1548	SFF Gigabit Ethernet Expansion Card	Yes[b]	Yes	No	No	Yes[b]	Yes
39Y9310	2969	Ethernet Expansion Card (CFFv)	Yes[b]	Yes	No	No	Yes[b]	Yes
26K4841	None	IBM SFF FC Expansion Card	No	No	Yes	No	No	Yes
26R0890	1577	QLogic 4 Gb SFF FC Expansion Card	No	No	Yes	No	No	Yes
26R0884	None	QLogic 4 Gb StFF FC Expansion Card	No	No	Yes	No	No	Yes
41Y8527	2970	QLogic 4 Gb FC Expansion Card (CFFv)	No	No	Yes	No	No	Yes
39Y9186	2925	Emulex 4 Gb SFF FC Expansion Card	No	No	Yes	No	No	Yes
43W6859	2994	Emulex 4 Gb FC Expansion Card (CFFv)	No	No	Yes	No	No	Yes
32R1923	1458	QLogic iSCSI Expansion Card	No	Yes	No	No	No	Yes
39Y9190	2979	SAS Expansion Card (CFFv)	No	No	No	Yes	No	No
43W3974	1591	SAS Connectivity Card (CFFv)	No	No	No	Yes	No	No
46C7167	5490	ServeRAID-MR10ie (CIOv)	No	No	No	Yes[c]	No	No

a. Only the SAS Connectivity Module (39Y9195) is meant under the SASM abbreviation in this table because SAS RAID Controller Module (43W3584) is not supported in the BladeCenter E chassis.
b. Supported only if the expansion card is installed in slot 1 of either the BladeCenter SCSI Storage Expansion Unit II (39R8625) or BladeCenter Storage and I/O Expansion Unit (39R7563, feature code 4265).
c. This card only supports SAS Connectivity Module (39Y9195, feature code 2980). Currently, only HS12 (8028) supports this card, and HS12 requires SAS Connectivity Card (43W3974, feature code 1591) to be installed into it to support external attachments.

Note: The BladeCenter H and BladeCenter HT chassis have different multi-switch interconnect modules.

MSIM for BladeCenter HT is called MSIM-HT, and its part number is 44R5913, feature code 5491.

Table 2-6 Supported combinations of high-speed I/O modules and I/O expansion cards - BladeCenter HT

Part Number	Feature Code	High-speed expansion card	ESM[a,b]	FCSM[a,c]	HSIBPM	HSESM
		I/O module bay number ♦	**7, 8, 9, 10[d]**	**8, 10[c,d]**	**7/8, 9/10[e]**	**7, 9**
39Y9271	2967	NetXen 10 Gb Eth. Exp. Card (CFFh)	No	No	No	Yes
39Y9306	2968	QLogic Eth. & 4 Gb FC Exp. Crd (CFFh)	Yes[f]	Yes	No	No
43W4423	2991	4X InfiniBand DDR Exp. Card (CFFh)	No	No	Yes	No
43W4420	2993	Voltaire 4X IB DDR Exp. Card (CFFh)	No	No	Yes	No
44W4479	5476	2/4 Port Ethernet Exp. Card (CFFh)	Yes[g]	No	No	No
44W4465	5479	Broadcom 10 Gb 4-port Eth. Exp. Card (CFFh)	No	No	No	No
44W4465	5489	Broadcom 10 Gb 2-port Eth. Exp. Card (CFFh)	No	No	No	No
44X1940	5485	QLogic Eth. & 8 Gb FC Exp. Crd (CFFh)	No	No	No	No

a. Supported in high-speed I/O bays using MSIM-HT (44R5913, feature code 5491).
b. Only Nortel L2/3 Copper 1 Gb Ethernet (32R1860, feature code 1495), Cisco Intelligent Copper 1 Gb Ethernet (32R1892, feature code 1498), and Nortel 1/10 Gb Ethernet Uplink (44W4404, feature code 1590) ESMs are supported in MSIM-HT.
c. Only QLogic 10 port (43W6724, feature code 2986) and 20-port (43W6725, feature code 2987) FCSMs are supported in MSIM-HT.
d. These high-speed I/O bays are converted to the standard I/O bays using MSIM-HT. In this case, bay 7 represents the upper left bay of MSIM-HT. Bay 8 represents the upper right bay. Bay 9 represents the lower left bay, and bay 10 represents the lower right bay.
e. One 4X InfiniBand Pass-thru module (HSIBPM) occupies two adjacent high-speed I/O bays.
f. This card has two Ethernet ports connected to bays 7 and 9 of MSIM-HT.
g. This card has four Ethernet ports connected to bays 7, 8, 9, and 10 of MSIM-HT.

2.1.5 BladeCenter S I/O topology

The IBM BladeCenter S chassis provides connection paths between server blade bays and I/O or switch bays (see Figure 1-1 on page 12 and Figure 1-3 on page 13) through a hard-wired dual redundant midplane.

The I/O topology is shown in Figure 2-5 on page 73. Each of six blade bays have six dedicated I/O connections linked to four I/O bays. I/O bays 1 and 2 have two connections to each blade bay. Bays 3 and 4 have one connection to each blade bay.

Thus, I/O bays 1 and 2 have a total of 12 I/O connections (two per blade bay), and bays 3 and 4 have a total of 6 I/O connections (one per blade bay). Additionally, I/O bays 3 and 4 have special wiring to enable support of integrated SAS storage. This wiring consists of four x4 SAS links between these bays and DSMs (two links per I/O bay 3 or 4 connected to different DSMs), and two network links (1 Gb Ethernet) between these bays and I/O bay 1 (one link per I/O bay 3 or 4 connected to I/O bay 1) for RAID management purposes.

SAS links are only used when SAS-based integrated storage solution is implemented. Additionally, Ethernet links from bays 3 and 4 to bay 1 are only used when SAS RAID Controller Module is installed into bays 3 and 4.

If you plan to use integrated SAS storage for BladeCenter S, refer to 5.10, “IBM BladeCenter S integrated storage” on page 423 for more information.

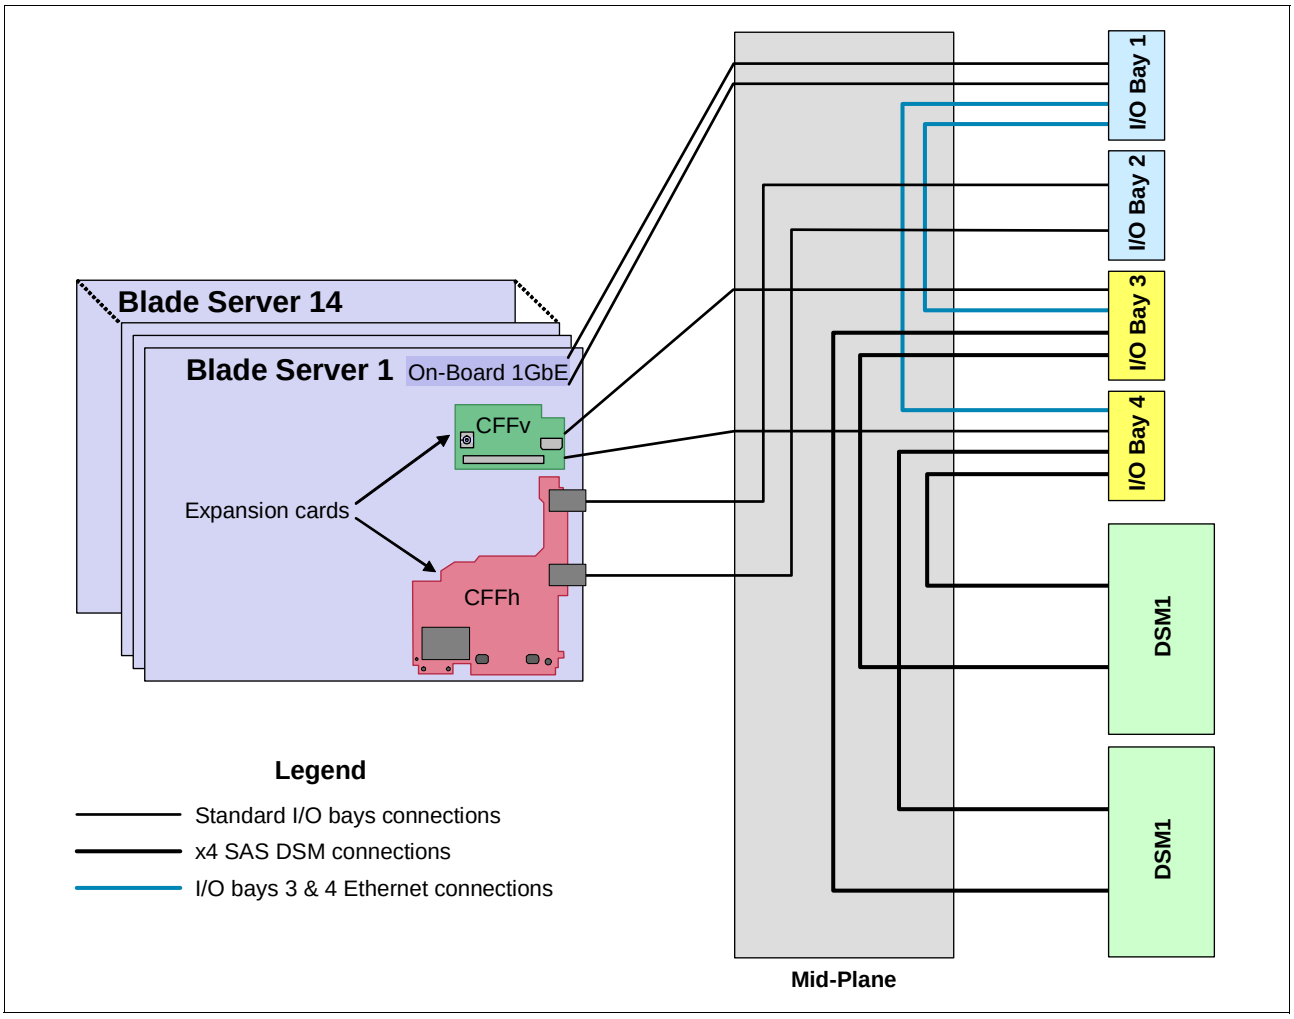

Figure 2-5 IBM BladeCenter S I/O connection paths

Each blade server has at least two Ethernet ports (Ethernet 1 and Ethernet 2). Both of them are routed to I/O bay 1 (and this is different from other BladeCenter chassis), which means that only Ethernet-capable I/O modules can be used in I/O bay 1. I/O expansion cards that you optionally install in blade servers have two ports which are routed to I/O bays 3 and 4, so depending on the expansion card installed, the compatible I/O modules must be chosen for those I/O bays. Unless an expansion card is installed in one or more processor blades, there is no need for I/O modules 3 and 4. I/O bay 2 is used when 2/4 Port Ethernet Expansion Card (CFFh) is installed into supported blade server; in this case, two Ethernet ports on this card are routed to the bay 2, and the other two ports on this card remain unused.

Note: I/O module bays 3 and 4 must both contain the same type of switch (either SAS connectivity modules, Ethernet switch modules, pass-thru modules, or Fibre Channel switch modules).

The mappings of onboard Ethernet ports to the ports of switch module in I/O bay 1 are listed in Table 2-7 on page 74.

Table 2-7 Onboard Ethernet port mappings for BladeCenter S

Blade Server Number	Onboard Ethernet Port Number	Switch Module Port Number
1	1	1
	2	8
2	1	2
	2	9
3	1	3
	2	10
4	1	4
	2	11
5	1	5
	2	12
6	1	6
	2	13

The I/O modules must be compatible with the I/O interfaces present in the blade servers. For example, when a Fibre Channel expansion card is installed in a blade server, I/O Modules 3 and 4 must also be Fibre Channel-based (that is, a FC switch module or Optical Pass-thru module, OPM), and vice versa. If you install FC switches in bays 3 and 4, then any expansion cards installed in all *other* blade servers in the same chassis must be Fibre Channel.

Table 2-8 on page 75 lists supported combinations of I/O modules and I/O expansion cards for the IBM BladeCenter S chassis.

Table 2-8 Supported combinations of I/O modules and I/O expansion cards - IBM BladeCenter S

Part Number	Feature code	Expansion card	ESM, CPM, ICPM			FC SM	SAS M	OPM		
		I/O module bay number ♦	1	2	3, 4	3, 4	3, 4	1	2	3, 4
None	None	Integrated Gigabit Ethernet	Yes	No	No	No	No	Yes	No	No
39R8624	1548	SFF Gigabit Ethernet Exp. Card	Yes[a]	No	Yes	No	No	Yes[a]	No	Yes
39Y9310	2969	Ethernet Expansion Card (CFFv)	Yes[a]	No	Yes	No	No	Yes[a]	No	Yes
26K4841	None	IBM SFF FC Expansion Card	No	No	No	Yes	No	No	No	Yes
26R0890	1577	QLogic 4 Gb SFF FC Exp. Card	No	No	No	Yes	No	No	No	Yes
26R0884	None	QLogic 4 Gb StFF FC Exp. Card	No	No	No	Yes	No	No	No	Yes
41Y8527	2970	QLogic 4 Gb FC Exp. Card (CFFv)	No	No	No	Yes	No	No	No	Yes
39Y9186	2925	Emulex 4 Gb SFF FC Exp. Card	No	No	No	Yes	No	No	No	Yes
43W6859	2994	Emulex 4 Gb FC Exp. Card (CFFv)	No	No	No	Yes	No	No	No	Yes
32R1923	1458	QLogic iSCSI Expansion Card	No	No	Yes	No	No	No	No	Yes
39Y9190	2979	SAS Expansion Card (CFFv)	No	No	No	No	Yes	No	No	No
43W3974	1591	SAS Connectivity Card (CFFv)	No	No	No	No	Yes[b]	No	No	No
46C7167	5490	ServeRAID-MR10ie (CIOv)	No	No	No	No	Yes[c]	No	No	No
44W4479	5476	2/4 Port Ethernet Exp. Card (CFFh)	No	Yes	No	No	No	No	Yes	No

a. Supported only if the expansion card is installed in slot 1 of either the BladeCenter SCSI Storage Expansion Unit II (39R8625) or BladeCenter Storage and I/O Expansion Unit (39R7563, feature code 4265).

b. This card only supports SAS Connectivity Module (39Y9195, feature code 2980). At the time of writing, SAS RAID Controller Module (43W3584, feature code 3734) is not supported by this card.

c. This card only supports SAS Connectivity Module (39Y9195, feature code 2980). Currently, only HS12 (8028) supports this card, and HS12 requires SAS Connectivity Card (43W3974, feature code 1591) to be installed into it to support external attachments.

Notes:

- Cisco Catalyst 3110G and 3110X ESMs are not supported in BladeCenter S.
- 20-port versions of FC switches (except Intelligent Optical Pass-thru module) are not supported in BladeCenter S.
- Due to specific wiring inside the BladeCenter S chassis, if you use an OPM, then you must have all four OPM cables to connect all ports of the installed blades.

As an example, examine the top data row, Integrated Gigabit Ethernet. Table 2-8 indicates that both onboard Ethernet ports on each blade are always routed to bay 1 of the BladeCenter S chassis, and only Ethernet-capable modules can be used in this bay; that is, Ethernet switches or Optical Pass-Thru modules.

As another example, we use QLogic 4 Gb SFF FC Expansion card in our blades. The row in the table for this card shows that ports of this card are always routed to bay 3 and 4, so we have to use FC-capable modules in those bays; that is, FC switch or Optical Pass-Thru modules.

Refer to Table 1-3 on page 5 for compatibility information about blades and their supported expansion options.

2.2 Ethernet switch modules

Various types of Ethernet switch modules from several vendors are available for BladeCenter, and they support different network layers and services, as shown in Table 2-9.

Table 2-9 Ethernet switch modules

Part number	Feature Code	Option description	Number x type of external ports	Network layers	Page
32R1888	1497	Cisco Systems Fiber Intelligent Gigabit ESM	4 x Gigabit Ethernet	Layer 2/3[a]	79
32R1892	1498	Cisco Systems Intelligent Gigabit ESM	4 x Gigabit Ethernet	Layer 2/3[a]	84
41Y8523	2989	Cisco Catalyst Switch Module 3110G	4 x Gigabit Ethernet, 2 x StackWise Plus	Layer 2/3	89
41Y8522	2988	Cisco Catalyst Switch Module 3110X	1 x 10 Gb Ethernet, 2 x StackWise Plus	Layer 2/3	97
43W4395	5450	Cisco Catalyst Switch Module 3012	4 x Gigabit Ethernet	Layer 2/3	106
39Y9324	1484	IBM Server Connectivity Module	6 x Gigabit Ethernet	Layer 2	113
32R1860	1495	Nortel Networks L2/3 Copper Gigabit ESM	6 x Gigabit Ethernet	Layer 2/3	115
32R1861	1496	Nortel Networks L2/3 Fiber Gigabit ESM	6 x Gigabit Ethernet	Layer 2/3	120
32R1859	1494	Nortel Networks Layer 2-7 Gigabit ESM	4 x Gigabit Ethernet	Layer 2/7	123
32R1783	1493	Nortel 10 Gb Uplink ESM	3 x 10 Gb Ethernet, 1 x 1 Gb Ethernet	Layer 2/3	126
44W4404	1590	Nortel 1/10 Gb Uplink ESM	3 x 10 Gb Ethernet, 6 x Gigabit Ethernet	Layer 2/3	130
39Y9267	2952	Nortel 10 Gb High Speed ESM	6 x 10 Gb Ethernet	Layer 2/3	134

a. Only Quality of Service (QoS) is supported by this switch for Layer 3.

Layer 2 support means that the Ethernet switch provides basic connectivity. It is capable of dealing with frame headers that contain physical addresses such as MAC addresses, for example, to forward network traffic based on a table of MAC addresses, or to distribute traffic across several physical links to form one logical higher-bandwidth link.

Layer 3 processing means that the switch provides advanced connectivity. It is capable of dealing with packet headers that contain logical addresses such as IP addresses, for example, to forward traffic based on an IP routing table, to build such table dynamically using routing protocols, or to control access to parts of the network by filtering traffic based on IP addresses.

If the switch supports layers above 3, that means it is able to inspect packet contents (such as TCP ports or even application protocols such as FTP, HTTP and so on), and not simply the header. This enables implementation of advanced features such as server load balancing or application security.

To meet the requirements of a client's network environment, the Ethernet switch modules must support certain network standards, protocols, and features that are required by the

network design. These features can be grouped into several categories which include VLAN support, performance, high-availability and redundancy, management, security, quality of service (QoS), and routing. Table 2-10 compares features and functions supported by different switches.

Note: Many software features and functions supported by the switch module depend on the firmware release loaded onto it. For example, old firmware supports 128 active VLANs only, while the newer version can support up to 1024 VLANs.

Refer to the Release Notes document for information about enhancements implemented in a specific firmware version. This document can be found on the firmware upgrade download page. Information covered in this section is based on the most current firmware releases available at the time of writing.

Table 2-10 Basic features supported by Ethernet switch modules

Feature	Cisco IGESM[a]	Cisco 3012	Cisco 3110G, 3110X	IBM SCM	Nortel L2/3 GbE SM[a]	Nortel L2/7 GbE SM	Nortel 10 Gb Uplink, 1/10 Gb Uplink ESMs	Nortel 10 Gb HSSM
Ports								
Number of external ports	4	4	4 (3110G) 1 (3110X)	6	6	4	4 (10G Uplink) 9 (1/10G Uplink)	6
Type of external ports	Ethernet							
Speed of external ports, Mbps	10/100/ 1000	10/100 /1000	3110G: 10/100/ 1000 3110X: 10000	10/100 /1000	10/100 / 1000	10/100 / 1000	10G Uplink: 3 x 10000, 1 x 1000 1/10G Uplink: 3 x 10000 6 x 10/100/1000	10000
Dedicated stacking ports	None	None	2 x StackWise Plus, 32 Gbps	None	None	None	None	None
Number of internal ports	14	14	14	14	14	14	14	14
Type of internal ports	Ethernet							
Speed of internal ports	1000	1000	1000	1000	1000	1000	1000	10000
VLAN support								
Number of VLANs supported	250	1005	1005	256	1024	1024	1024	1024
802.1Q VLAN Tagging	Yes	Yes	Yes	Yes[b]	Yes	Yes	Yes	Yes
Cisco ISL Trunking	No	Yes	Yes	No	No	No	No	No
Performance								
Link Aggregation	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Jumbo frames	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Feature	Cisco IGESM[a]	Cisco 3012	Cisco 3110G, 3110X	IBM SCM	Nortel L2/3 GbE SM[a]	Nortel L2/7 GbE SM	Nortel 10 Gb Uplink, 1/10 Gb Uplink ESMs	Nortel 10 Gb HSSM
IGMP Snooping	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Server Load Balancing	No	No	No	No	No	Yes	No	No
Content Switching	No	No	No	No	No	Yes	No	No
High availability and redundancy								
Spanning Tree Protocol	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Trunk Failover (Link State Tracking)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
VRRP	No	No	No	No	Yes	Yes	Yes	Yes
HSRP	No	Yes	Yes	No	No	No	No	No
Management								
Serial port	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Telnet/SSH	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SNMP protocol	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cisco CDP protocol	Yes	Yes	Yes	No	No	No	No	No
Cisco VTP protocol	Yes	Yes	Yes	No	No	No	No	No
Serial over LAN support[c]	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Switch stacking	No	No	Yes	No	No	No	Planned[d]	No
Security								
Port-based security	Yes	Yes	Yes	No	Yes	No	Yes	Yes
Access Control Lists (MAC-based)	Yes	Yes	Yes	No	Yes	No	Yes	Yes
Access Control Lists (IP-based)	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
TACACS+	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RADIUS	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
NAT	No	No	No	No	No	Yes	No	No
Quality of Service (QoS)								
802.1p	Yes	Yes	Yes	No	Yes	No	Yes	Yes
IP ToS/DSCP	Yes	Yes	Yes	No	Yes	No	Yes	Yes
Routing protocols								
RIP	No	Yes	Yes	No	Yes	Yes	Yes	Yes
OSPF	No	No	Yes[e]	No	Yes	Yes	Yes	Yes
BGP	No	No	Yes[e]	No	Yes	Yes	Yes	Yes

Feature	Cisco IGESM[a]	Cisco 3012	Cisco 3110G, 3110X	IBM SCM	Nortel L2/3 GbE SM[a]	Nortel L2/7 GbE SM	Nortel 10 Gb Uplink, 1/10 Gb Uplink ESMs	Nortel 10 Gb HSSM
EIGRP	No	No	Yes[e]	No	No	No	No	No
IPv6 support								
IPv6 host	No	Yes	Yes	No	No	No	Planned[f]	No
IPv6 routing	No	No	Yes[g]	No	No	No	Planned[f]	No
IPv6 Access Control Lists	No	No	Yes[g]	No	No	No	Planned[f]	No

a. Applied to both copper and fiber models except when exactly referenced
b. Supported for external ports only
c. Blade-specific feature, not all blades support Serial over LAN
d. Stacking support is planned for release for Nortel 1/10 Gb Uplink switch (44W4404, feature code 1590) in 2008
e. IP Services license is required for Cisco Catalyst 3110 family of switches to support this feature
f. IPv6 support is planned for release for Nortel 1/10 Gb Uplink switch (44W4404, feature code 1590) in 2008
g. Advanced IP Services license is required for Cisco Catalyst 3110 family of switches to support this feature

2.2.1 Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module

The Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module, part number 32R1888, feature code 1497, is an I/O module that is installed into a BladeCenter unit. The IGESM is a full wire-rated non-blocking switch for use with high performance servers.

The IGESM uses a custom chipset that permits its operation in the power and cooling requirements of the IBM BladeCenter, while still providing wire-rate, non-blocking performance to the attaching devices. The IOS in use on this switch is based on the Cisco Catalyst 2950 Enhanced Image code, but has certain enhancements and modifications to operate within the environment of the IBM BladeCenter.

Figure 2-6 on page 80 shows the Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module.

Note: This switch is shipped without SFP modules. SX SFP modules (part number GLC-SX-MM=) should be ordered separately from Cisco Systems or a Cisco Systems reseller.

Figure 2-6 Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module

The supported features and specifications for the Fiber Ethernet switch module are:

- Ports
  - RJ-45 serial port. This is the standard Cisco console port available on other Cisco equipment and offers a totally out-of-band management path if desired.
  - 14 internal full-duplex Gigabit ports, one connection to each of the blade servers in the BladeCenter unit.
  - Support for four small-form-factor fiber plug-in modules (SFP) for 1000 Mbps fiber connections to external networks (available separately from Cisco and Cisco resellers).
  - Two internal full-duplex 100 Mbps port connections to the management module.
- Scalability and performance features
  - Gigabit EtherChannel for enhanced fault tolerance and for providing up to 4 Gbps of bandwidth between switches, routers, and servers.
  - Support for standard frames with sizes from 64 to 1530 bytes and jumbo frames with a maximum size of 9216.
  - Per-port broadcast-storm control for preventing a faulty endstation from degrading overall system performance with broadcast storms.
  - Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) for automatic creation of EtherChannel links.
  - Internet Group Management Protocol (IGMP) snooping support to limit flooding of IP multicast traffic.
  - Multicast VLAN registration (MVR) to continuously send multicast streams in a multicast VLAN while isolating the streams from subscriber VLANs for bandwidth and security.
  - IGMP filtering for controlling the set of multicast groups to which hosts on a switch port can belong.

- Protected port (private VLAN edge port) option for restricting the forwarding of traffic to designated ports on the same switch.
- Dynamic address learning for enhanced security.
- Supports multiple EtherChannel load balance algorithms (SMAC or DMAC, SIP or DIP, XOR-SMAC/DMAC or XOR-SIP/DIP) to offer maximum performance on aggregated links.

- Manageability
  - Address Resolution Protocol (ARP) for identifying a switch through its IP address and its corresponding MAC address.
  - Cisco Discovery Protocol (CDP) Versions 1 and 2 to aid in troubleshooting and reporting on misconfiguration of ports connecting to other devices supporting CDP.
  - Network Time Protocol (NTP) for providing a consistent time stamp to all switches from an external source.
  - Directed unicast requests to a Trivial File Transfer Protocol (TFTP) server for obtaining software upgrades from a TFTP server.
  - Default configuration storage in flash memory to ensure that the switch can be connected to a network and can forward traffic with minimal user intervention.
  - In-band monitoring of the switch through built-in Cisco Device Manager Web-based tool.
  - In-band management access through up to 16 simultaneous Telnet connections for multiple command-line-interface (CLI)-based sessions over the network.
  - In-band management access through up to five simultaneous, encrypted Secure Shell (SSH) connections for multiple CLI-based sessions over the network is available only in the cryptographic software image.
  - In-band management access through SNMP versions 1, 2c, and 3 get and set requests.
  - Supported by CiscoWorks management software.
  - Protected Mode feature, to isolate switch management from Advanced Management Module, for increased security of the switch.
  - Cisco Network Assistance (CNA), free GUI-based application tool to configure most features of this switch. For more information and a download of CNA, go to:

    `http://www.cisco.com/go/cna`
  - Extensive debugging options to aid in troubleshooting and diagnosing issues.
  - Support for multiple management interfaces.

- Redundancy
  - Link state tracking to mirror the state of the external ports on the internal Ethernet links and to allow the failover of the processor blade traffic to an operational external link on a separate Cisco Ethernet switch.
  - Configurable unidirectional link detection (UDLD) for detecting and disabling unidirectional links. This feature will prevent a larger network failure in the event a unidirectional link is detected, thus reducing downtime in these situations.
  - IEEE 802.1D Spanning Tree Protocol (STP) for redundant backbone connections and loop-free networks.
  - IEEE 802.1s Multiple STP (MSTP) for grouping VLANs into a spanning-tree instance, and provided for multiple forwarding paths for data traffic and load balancing, up to 64 STP instances per switch.

  - IEEE 802.1w Rapid STP (RSTP) for rapid convergence of the spanning tree by immediately transitioning root and designated ports to the converting state.
  - Optional spanning-tree features available in the PVST+, rapid PVST+, and MSTP modes.
- VLAN support
  - Support for 250 total VLANs. These VLANs can be any VLAN ID from 1–4094, except 1001–1005, which are reserved by Cisco.
  - IEEE 802.1Q trunking protocol on all ports for network moves, adds, and changes; management and control of broadcast and multicast traffic; and network security by establishing VLAN groups for high-security users and network resources.
  - VLAN Membership Policy Server (VMPS) for dynamic VLAN membership.
  - VLAN Trunking Protocol (VTP) pruning for reducing network traffic by restricting flooded traffic to links destined for stations receiving the traffic.
  - Dynamic Trunking Protocol (DTP) for negotiating trunking on a link between two devices and for negotiating the type of trunking encapsulation (802.1Q) to be used.
  - VLAN 1 minimization to reduce the risk of spanning-tree loops or storms by enabling VLAN 1 to be disabled on any individual VLAN trunk link. With this feature enabled, no user traffic is sent or received. The microprocessor continues to send and receive control protocol frames.
  - Support for up to 64 instances of spanning tree per switch.
  - Support for nonstandard VLAN 4095 for serial-over-LAN traffic.
- Security
  - Bridge protocol data unit (BPDU) guard for shutting down a Port Fast-configured port when an invalid configuration occurs.
  - Protected port option for restricting the forwarding of traffic to designated ports on the same switch.
  - Password-protected access (read-only and write-only access) to management interfaces (the device manager and CLI) for protection against unauthorized configuration changes.
  - Port security option for limiting and identifying MAC addresses of the station allowed to access the port.
  - Port security aging to set the aging time for secure addresses on a port.
  - Multilevel security for a choice of security level, notification, and resulting actions.
  - MAC-based port-level security for restricting the use of a switch port to a specific group of source addresses and preventing switch access from unauthorized stations.
  - Terminal Access Controller Access Control System Plus (TACACS+), a proprietary feature for managing network security through a TACACS server.
  - IEEE 802.1X port-based authentication to prevent unauthorized devices from gaining access to the network.
  - IEEE 802.1X port-based authentication with VLAN assignment for restricting 802.1X-authenticated users to a specified VLAN.
  - IEEE 802.1X port-based authentication with port security for authenticating the port and managing network access for all MAC addresses, including that of the client.
  - IEEE 802.1X port-based authentication with voice VLAN to allow an IP phone access to the voice VLAN irrespective of the authorized or unauthorized state of the port.

  - IEEE 802.1X port-based authentication with guest VLAN to provided limited services to non-802.1X-compliant users.
  - Standard and extended IP access control lists (ACLs) for defining security policies.
  - Support for RADIUS and TACACS+ user authentication.
- Quality of Service (QoS) and Class of Service (CoS)
  - Classification - IEEE 802.1p CoS with eight priority queues on the Fiber Ethernet ports for prioritizing mission-critical and time-sensitive traffic from data, voice, and telephony applications.
  - IP Differentiated Services Code Point (IP DSCP) and CoS marking priorities on a per-port basis for protecting the performance of mission-critical applications.
  - Flow-based packet classification (classification based on information in the MAC, IP, and TCP/UDP headers) for high-performance QoS at the network edge, allowing for differentiated service levels for different types of network traffic and for prioritizing mission-critical traffic in the network.
  - Support for IEEE 802.1p CoS scheduling for classification and preferential treatment of high-priority voice traffic.
  - Trusted boundary (detect the presence of a Cisco IP Phone, trust the CoS value received, and ensure port security).
  - Traffic-shaping policies on the switch port for allocating the amount of the port bandwidth to a specific traffic flow.
  - Out-of-profile markdown for packets that exceed bandwidth utilization limits.
  - Egress Policing and Scheduling of Egress Queues - four egress queues on all Ethernet switch-module ports; support for strict priority and weighted round-robin (WRR) CoS policies.
- Monitoring
  - Switch LEDs that provide visual port and Fiber Ethernet switch module status.
  - SPAN/RSPAN support for local and remote monitoring of the network.
  - Four groups (history, statistics, alarms, and events) of embedded remote monitoring (RMON) agents for network monitoring and traffic analysis.
  - MAC address notification for tracking the MAC addresses that the Fiber Ethernet switch module has learned or removed.
  - Syslog facility for logging system messages about authentication or authorization errors, resource issues, and time- out events.
  - Layer 2 trace route to identify the physical path that a packet takes from a source device to a destination device.
- Network connectors for Small-Form-Factor Plug-in (SFP) modules (available separately) – multimode SX Fiber (500 meters maximum).

The Fiber IGESM supports the following IEEE standards:

- IEEE 802.1D Spanning Tree Protocol (STP)
- IEEE 802.1s Multiple STP
- IEEE 802.1w rapid STP
- IEEE 802.1p QoS prioritization
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1x Port-based Authentication
- IEEE 802.2 Logical Link Control
- IEEE 802.3ad Link Aggregation

- IEEE 802.3x Full Duplex Flow Control on external ports
- IEEE 802.3z
  - 1000BASE-SX Gb Fiber Ethernet
  - Link Negotiation

Refer to:

- *Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-63993
- Cisco *Systems Intelligent Gigabit Ethernet Switch Module Software Configuration Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55261

2.2.2 Cisco Systems Intelligent Gigabit Ethernet Switch Module

The Cisco Systems Intelligent Gigabit Ethernet Switch Module, part number 32R1892, feature code 1498, is an I/O module that is installed into a BladeCenter unit. This switch is a full wire-rated, non-blocking switch for use with high performance servers.

The IGESM uses a custom chipset that permits its operation in the power and cooling requirements of the IBM BladeCenter, while still providing wire-rate, non-blocking performance to the attaching devices. The IOS in use on this switch is based on the Cisco Catalyst 2950 Enhanced Image code, but has certain enhancements and modifications to operate within the environment of the IBM BladeCenter.

Figure 2-7 shows the Cisco Systems Intelligent Gigabit Ethernet Switch Module.

Figure 2-7 Cisco Systems Intelligent Gigibit Ethernet Switch Module

The supported features and specifications for the Gigabit Ethernet switch module are:

- Ports
  - RJ-45 serial port. This is the standard Cisco console port available on other Cisco equipment, and offers a totally out-of-band management path if desired.
  - Four external 1000BASE-T connectors for making 10/100/1000 Mbps connections to a backbone, end stations, and servers.
  - 14 internal full-duplex Gigabit ports, one connected to each of the blade servers in the BladeCenter unit.
  - Two internal full-duplex 100 Mbps ports connected to the management module.
- Scalability and performance features
  - Autosensing of speed on the 10/100/1000 ports and auto-negotiation of duplex mode on the ports for optimizing bandwidth.
  - Gigabit EtherChannel for enhanced fault-tolerance and for providing up to 4 Gbps of bandwidth between switches, routers, and servers.
  - Support for standard frames with sizes from 64 to 1530 bytes and jumbo frames with a maximum size of 9216.
  - Per-port broadcast-storm control for preventing a faulty endstation from degrading overall system performance with broadcast storms.
  - Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) for automatic creation of EtherChannel links.
  - Internet Group Management Protocol (IGMP) snooping support to limit flooding of IP multicast traffic.
  - Multicast Virtual Local Area Network (VLAN) registration (MVR) to continuously send multicast streams in a multicast VLAN while isolating the streams from subscriber VLANs for bandwidth and security.
  - IGMP filtering for controlling the set of multicast groups to which hosts on a switch port can belong.
  - Protected port (private VLAN edge port) option for restricting the forwarding of traffic to designated ports on the same switch.
  - Dynamic address learning for enhanced security.
  - Supports multiple EtherChannel load balance algorithms (SMAC or DMAC, SIP or DIP, XOR-SMAC/DMAC or XOR-SIP/DIP) to offer maximum performance on aggregated links.
- Manageability
  - Address Resolution Protocol (ARP) for identifying a switch through its IP address and its corresponding MAC address.
  - Cisco Discovery Protocol (CDP) Versions 1 and 2 to aid in troubleshooting and reporting on misconfiguration of ports connecting to other devices supporting CDP.
  - Network Time Protocol (NTP) for providing a consistent time stamp to all switches from an external source.
  - Directed unicast requests to a Trivial File Transfer Protocol (TFTP) server for obtaining software upgrades from a TFTP server.
  - Default configuration storage in flash memory to ensure that the switch can be connected to a network and can forward traffic with minimal user intervention.

- In-band monitoring of the switch through a built-in Cisco Device Manager Web-based tool.
- In-band management access through up to 16 simultaneous Telnet connections for multiple command line interface (CLI)-based sessions over the network.
- In-band management access through up to five simultaneous, encrypted Secure Shell (SSH) connections for multiple CLI-based sessions over the network. This option is available only in the cryptographic software image.
- In-band management access through SNMP versions 1, 2c, and 3 get and set requests.
- Supported by CiscoWorks management software.
- Protected Mode feature to isolate switch management from Advanced Management Module, for increased security of the switch.
- Cisco Network Assistance (CNA), free GUI-based application tool to configure most features of this switch. For more information and download of CNA, go to:

  http://www.cisco.com/go/cna
- Extensive debugging options to aid in troubleshooting and diagnosing issues.
- Support for multiple management interfaces.

► Redundancy
- Link state tracking to mirror the state of the external ports on the internal Ethernet links and to allow the failover of the processor blade traffic to an operational external link on a separate Cisco Ethernet switch.
- Configurable Unidirectional link detection (UDLD) for detecting and disabling unidirectional links. This feature prevents a larger network failure in the event that a unidirectional link is detected, thus reducing down time in these situations.
- IEEE 802.1D Spanning Tree Protocol (STP) for redundant backbone connections and loop-free networks.
- IEEE 802.1s Multiple STP (MSTP) for grouping VLANs into a spanning-tree instance, and provided for multiple forwarding paths for data traffic and load balancing, up to 64 STP instances per switch.
- IEEE 802.1w Rapid STP (RSTP) for rapid convergence of the spanning tree by immediately transitioning root and designated ports to the converting state.
- Optional spanning-tree features available in the PVST+, rapid PVST+, and MSTP modes.

► VLAN support
- Support for 250 total VLANs. These VLANs can be any VLAN ID from 1–4094, except 1001–1005, which are reserved by Cisco.
- IEEE 802.1Q trunking protocol on all ports for network moves, adds, and changes; management and control of broadcast and multicast traffic; and network security by establishing VLAN groups for high-security users and network resources.
- VLAN Membership Policy Server (VMPS) for dynamic VLAN membership.
- VLAN Trunking Protocol (VTP) pruning for reducing network traffic by restricting flooded traffic to links destined for stations receiving the traffic.
- Dynamic Trunking Protocol (DTP) for negotiating trunking on a link between two devices and for negotiating the type of trunking encapsulation (802.1Q) to be used.
- Voice VLAN for creating subnets for voice traffic from Cisco IP phones.

- VLAN 1 minimization to reduce the risk of spanning-tree loops or storms by enabling VLAN 1 to be disabled on any individual VLAN trunk link. With this feature enabled, no user traffic is sent or received. The switch CPU continues to send and receive control protocol frames.
- Support for up to 64 instances of spanning tree per switch.
- Support for nonstandard VLAN 4095 for Serial-over-LAN traffic.

- Security
  - Bridge protocol data unit (BPDU) guard for shutting down a Port Fast-configured port when an invalid configuration occurs.
  - Protected port option for restricting the forwarding of traffic to designated ports on the same switch.
  - Password-protected access (read-only and write-only access) to management interfaces (the device manager and CLI) for protection against unauthorized configuration changes.
  - Port security option for limiting and identifying MAC addresses of the station allowed to access the port.
  - Port security aging to set the aging time for secure addresses on a port.
  - Multilevel security for a choice of security level, notification, and resulting actions.
  - MAC-based port-level security for restricting the use of a switch port to a specific group of source addresses and preventing switch access from unauthorized stations.
  - Terminal Access Controller Access Control System Plus (TACACS+), a proprietary feature for managing network security through a TACACS server.
  - IEEE 802.1X port-based authentication to prevent unauthorized devices from gaining access to the network.
  - IEEE 802.1X port-based authentication with VLAN assignment for restricting 802.1X-authenticated users to a specified VLAN.
  - IEEE 802.1X port-based authentication with port security for authenticating the port and managing network access for all MAC addresses, including that of the client.
  - IEEE 802.1X port-based authentication with voice VLAN to allow an IP phone access to the voice VLAN irrespective of the authorized or unauthorized state of the port.
  - IEEE 802.1X port-based authentication with guest VLAN to provided limited services to non-802.1X-compliant users.
  - Standard and extended IP access control lists (ACLs) for defining security policies.
  - Support for RADIUS and TACACS+ user authentication.
- Quality of Service (QoS) and Class of Service (CoS)
  - Classification
    - IEEE 802.1p CoS with eight priority queues on the Gigabit Ethernet switch-module ports for prioritizing mission-critical and time-sensitive traffic from data, voice, and telephony applications.
    - IP Differentiated Services Code Point (IP DSCP) and CoS marking priorities on a per-port basis for protecting the performance of mission-critical applications.
    - Flow-based packet classification (classification based on information in the MAC, IP, and TCP/UDP headers) for high-performance QoS at the network edge, allowing for differentiated service levels for different types of network traffic and for prioritizing mission-critical traffic in the network.

        - Support for IEEE 802.1p CoS scheduling for classification and preferential treatment of high-priority voice traffic
    - Trusted boundary (detect the presence of a Cisco IP Phone, trust the CoS value received, and ensure port security.
    - Traffic-shaping policies on the switch port for allocating the amount of the port bandwidth to a specific traffic flow.
    - Out-of-profile markdown for packets that exceed bandwidth utilization limits.
    - Egress policing and scheduling of egress queues - four egress queues on all switch ports; support for strict priority and weighted round-robin (WRR) CoS policies.
- Monitoring
    - Switch LEDs that provide visual port and switch status.
    - SPAN/RSPAN support for local and remote monitoring of the network.
    - Four groups (history, statistics, alarms, and events) of embedded remote monitoring (RMON) agents for network monitoring and traffic analysis.
    - MAC address notification for tracking the MAC addresses that the switch has learned or removed.
    - Syslog facility for logging system messages about authentication or authorization errors, resource issues, and time-out events.
    - Layer 2 trace route to identify the physical path that a packet takes from a source device to a destination device.
- Network cables:
    - 10BASE-T:
        - UTP Category 3, 4, 5 (100 meters (328 feet) maximum)
        - 100-ohm STP (100 meters maximum)
    - 100BASE-TX:
        - UTP Category 5 (100 meters maximum)
        - EIA/TIA-568 100-ohm STP (100 meters maximum)
    - 1000BASE-T:
        - UTP Category 6
        - UTP Category 5e (100 meters maximum)
        - UTP Category 5 (100 meters maximum)
        - EIA/TIA-568B 100-ohm STP (100 meters maximum)

The Gigabit Ethernet switch module supports the following IEEE standards:

- IEEE 802.1D Spanning Tree Protocol (STP)
- IEEE 802.1s Multiple STP
- IEEE 802.1w Rapid STP
- IEEE 802.1p CoS Prioritization
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1x Port-based authentication
- IEEE 802.2 Logical Link Control
- IEEE 802.3 10BASE-T Ethernet
- IEEE 802.3u 100BASE-TX Fast Ethernet
- IEEE 802.3ab 1000BASE-TX Gigabit Ethernet
- IEEE 802.3ad Link Aggregation
- IEEE 802.3x Full-duplex Flow Control for external ports only
- IEEE 802.3z 1000BASE-X

Refer to the following documentation for more information:

- *Release Notes*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55262
- *Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-57858
- *Software Configuration Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55261
- *Command Reference*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55259
- *System Message Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55259

2.2.3 Cisco Catalyst Switch Module 3110G

The Cisco Catalyst Switch Module 3110G, part number 41Y8523, feature code 2989, is an I/O module that is installed into a BladeCenter unit. This switch is a full wire-rated, non-blocking switch for use with high performance servers.

This switch provides a next generation networking solution for blade server environments. Built upon Cisco's market-leading hardware and IOS software, the switches are engineered with technologies specifically designed to help meet the rigors of blade server-based application infrastructure. Specifically, the switches are designed to deliver scalable, high performance, highly resilient connectivity while supporting ongoing IT initiatives around reducing server infrastructure complexity and TCO by seamlessly integrating into existing Cisco management networks. The Cisco Catalyst Switch Module 3110G is shown in Figure 2-8 on page 90.

Note: In addition to the IBM sales channel, this switch module can also be ordered from Cisco Systems resellers or directly from Cisco Systems (part number WS-CBS3110G-S-I).

Figure 2-8 Cisco Catalyst Switch Module 3110G

The Cisco Catalyst Switch Module 3110G comes standard with IP Base feature set software. Additional features, such as IPv6 forwarding and routing support, as well as advanced routing protocols support (EIGRP, OSPF, BGP, PIM, and so on) for standard IP (IPv4), require additional licenses, as listed in Table 2-11.

Table 2-11 Additional feature sets for Cisco Catalyst Switch Module 3110G

Part number	Feature Code	Cisco Systems part number	Description
43W4434	4901	3110-IPS-LIC-I	IP Services S/W Upgrade License for Cisco Catalyst 3110[a]
None	None	3110-AISK9-LIC-I	S/W Upgrade License for Cisco Catalyst 3110[b] to Advanced IP Services

a. Provides support for advanced routing protocols including EIGRP, OSPF, BGP, PIM, and so on. Can be ordered through standard IBM sales channels or from Cisco Systems sales channels.
b. Provides support for IPv6 forwarding and routing. Available through Cisco Systems sales channels only.

The supported features and specifications for the Cisco Catalyst 3110G switch module are:

- Ports
  - USB-style serial port. This is the Cisco console port, and offers a totally out-of-band management path if desired. USB-to-DB-9 cable is used to connect the switch module to PC. This cable is shipped with the switch.
  - Four external RJ-45 1000BASE-T connectors for making 10/100/1000 Mbps connections to a backbone, end stations, and servers.
  - Two external high-speed StackWise Plus ports for switch module stacking to support Virtual Blade Switch (VBS) technology. Each 3110G switch module ships with one

1-meter StackWise Plus cable. Other cable lengths are available for order from Cisco Systems or Cisco Systems resellers if required:

- CAB-STK-E-0.5M= (0.5-meter cable)
- CAB-STK-E-1M= (1-meter cable)
- CAB-STK-E-3M= (3-meter cable)

Refer to 4.6, “Virtual Blade Switch technology” on page 384 for more information about VBS.

– 14 internal full-duplex Gigabit ports, one connected to each of the blade servers in the BladeCenter unit.

– One internal full-duplex 100 Mbps port connected to the management module.

► Performance features

– Autosensing of speed on the 10/100/1000 ports and auto-negotiation of duplex mode on the ports for optimizing bandwidth.

– Up to 64 Gbps of throughput in a switch stack.

– Gigabit EtherChannel for enhanced fault tolerance and for providing up to 8 Gbps of bandwidth between switches, routers, and servers.

– Support for standard frames with sizes from 64 to 1530 bytes and jumbo frames with a maximum size of 9216.

– Forwarding of Layer 2 frames and Layer 3 packets at 1 Gbps line rate across switches in stack.

– Per-port broadcast-storm control for preventing a faulty end station from degrading overall system performance with broadcast storms.

– Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) for automatic creation of EtherChannel links.

– Internet Group Management Protocol (IGMP) snooping support to limit flooding of IP multicast traffic.

– Multicast Virtual Local Area Network (VLAN) registration (MVR) to continuously send multicast streams in a multicast VLAN while isolating the streams from subscriber VLANs for bandwidth and security.

– IGMP filtering for controlling the set of multicast groups to which hosts on a switch port can belong.

– Dynamic address learning for enhanced security.

– Supports multiple EtherChannel load balance algorithms (SMAC or DMAC, SIP or DIP, XOR-SMAC/DMAC or XOR-SIP/DIP) to offer maximum performance on aggregated links.

– Web Cache Communication Protocol (WCCP) for redirecting traffic to wide area application engines, for enabling content requests to be fulfilled locally, and for localizing Web traffic patterns in the network (supported by IP Services feature set only)

► Manageability

– Address Resolution Protocol (ARP) for identifying a switch through its IP address and its corresponding MAC address.

– Cisco Discovery Protocol (CDP) Versions 1 and 2 to aid in troubleshooting and reporting on misconfiguration of ports connecting to other devices supporting CDP.

– Link Layer Discovery Protocol (LLDP) and LLDP Media Endpoint Discovery (LLDP-MED) for interoperability with third-party IP phones.

- Network Time Protocol (NTP) for providing a consistent time stamp to all switches from an external source.
- Directed unicast requests to a Trivial File Transfer Protocol (TFTP) server for obtaining software upgrades from a TFTP server.
- Default configuration storage in flash memory to ensure that the switch can be connected to a network and can forward traffic with minimal user intervention.
- In-band monitoring of the switch through built-in Cisco Device Manager Web-based tool.
- In-band management access through up to 16 simultaneous Telnet connections for multiple command line interface (CLI)-based sessions over the network.
- In-band management access through up to five simultaneous, encrypted Secure Shell (SSH) connections for multiple CLI-based sessions over the network. This option is available only in the cryptographic software image.
- In-band management access through SNMP versions 1, 2c, and 3 get and set requests.
- Out-of-band management (CLI) with switch module's console port.
- Supported by CiscoWorks management software.
- Protected Mode feature to isolate switch management from Advanced Management Module, for increased security of the switch.
- Cisco Network Services (CNS) embedded agents for automating switch management, configuration store, and delivery.
- Cisco Network Assistance (CNA), free GUI-based application tool to configure most features of this switch. For more information and download of CNA, go to:

  http://www.cisco.com/go/cna
- Extensive debugging options to aid in troubleshooting and diagnosing issues.
- Support for multiple management interfaces.

► Availability and redundancy

- Hot Standby Routing Protocol (HSRP) for Layer 3 router redundancy.
- Automatic stack master failover for replacing failed stack masters.
- Cross-stack EtherChannel for providing redundant links across switch stack.
- Link state tracking to mirror the state of the external ports on the internal Ethernet links and to allow the failover of the processor blade traffic to an operational external link on a separate Cisco Ethernet switch.
- Configurable Unidirectional link detection (UDLD) for detecting and disabling unidirectional links. This feature prevents a larger network failure in the event that a unidirectional link is detected, thus reducing downtime in these situations.
- IEEE 802.1D Spanning Tree Protocol (STP) for redundant backbone connections and loop-free networks.
- IEEE 802.1s Multiple STP (MSTP) for grouping VLANs into a spanning-tree instance, and provided for multiple forwarding paths for data traffic and load balancing.
- IEEE 802.1w Rapid STP (RSTP) for rapid convergence of the spanning tree by immediately transitioning root and designated ports to the converting state.
- Optional spanning-tree features available in the PVST+, rapid PVST+, and MSTP modes.

- Flex Link Layer 2 interfaces to back up one another as an alternative to STP for basic link redundancy.

- VLAN support
  - Support for 1005 total VLANs. These VLANs can be any VLAN ID from 1–4094, except 1001–1005, which are reserved by Cisco.
  - Cisco Inter-Switch Link (ISL) and IEEE 802.1Q trunking protocol on all ports for network moves, adds, and changes; management and control of broadcast and multicast traffic; and network security by establishing VLAN groups for high-security users and network resources.
  - VLAN Query Protocol (VQP) for dynamic VLAN membership.
  - VLAN Trunking Protocol (VTP) pruning for reducing network traffic by restricting flooded traffic to links destined for stations receiving the traffic.
  - Dynamic Trunking Protocol (DTP) for negotiating trunking on a link between two devices and for negotiating the type of trunking encapsulation (802.1Q) to be used.
  - Voice VLAN for creating subnets for voice traffic from Cisco IP phones.
  - VLAN 1 minimization to reduce the risk of spanning-tree loops or storms by enabling VLAN 1 to be disabled on any individual VLAN trunk link. With this feature enabled, no user traffic is sent or received. The switch CPU continues to send and receive control protocol frames.
  - Private VLANs to address VLAN scalability issues.
  - VLAN Flex Link Load Balancing to provide Layer 2 link redundancy without STP.
  - Support for up to 128 instances of spanning tree per switch or per switch stack.

- Security
  - Bridge protocol data unit (BPDU) guard for shutting down a Port Fast-configured port when an invalid configuration occurs.
  - Protected port option for restricting the forwarding of traffic to designated ports on the same switch.
  - Password-protected access (read-only and write-only access) to management interfaces (the device manager and CLI) for protection against unauthorized configuration changes.
  - Port security option for limiting and identifying MAC addresses of the station allowed to access the port.
  - Port security aging to set the aging time for secure addresses on a port.
  - Multilevel security for a choice of security level, notification, and resulting actions.
  - MAC-based port-level security for restricting the use of a switch port to a specific group of source addresses and preventing switch access from unauthorized stations.
  - MAC-based access control lists (ACLs).
  - Standard and extended IP access control lists (ACLs) for defining security policies on Layer 3 (router ACLs) and Layer 2 (port ACLs) interfaces.
  - Terminal Access Controller Access Control System Plus (TACACS+), a proprietary feature for managing network security through a TACACS server.
  - RADIUS for verifying the identity of, granting access to, and tracking activities of remote users.

- IEEE 802.1X port-based authentication to prevent unauthorized devices from gaining access to the network
  - IEEE 802.1X port-based authentication with VLAN assignment for restricting 802.1X-authenticated users to a specified VLAN.
  - IEEE 802.1X port-based authentication with port security for authenticating the port and managing network access for all MAC addresses, including that of the client.
  - IEEE 802.1X port-based authentication with voice VLAN to allow an IP phone access to the voice VLAN irrespective of the authorized or unauthorized state of the port.
  - IEEE 802.1X port-based authentication with guest VLAN to provided limited services to non-802.1X-compliant users.
  - IEEE 802.1X accounting to track network usage.

- Quality of Service (QoS) and Class of Service (CoS)
  - Automatic QoS (auto-QoS) to simplify the deployment of existing QoS features by classifying traffic and configuring egress queues.
  - Cross-stack QoS for configuring QoS features to all switches in a switch stack rather than on an individual-switch basis.
  - Classification
    - IP Type of Service/Differentiated Services Code Point (IP ToS/DSCP) and IEEE 802.1p CoS marking priorities on a per-port basis for protecting the performance of mission-critical applications.
    - IP ToS/DSCP and IEEE 802.1p CoS marking for flow-based packet classification (classification based on information in the MAC, IP, and TCP/UDP headers) for high-performance QoS at the network edge, allowing for differentiated service levels for different types of network traffic and prioritizing mission-critical traffic in the network.
    - Trusted port states (CoS, DSCP, and IP precedence) within a QoS domain and with a port bordering another QoS domain.
    - Trusted boundary for detecting the presence of a Cisco IP Phone, trusting the CoS value received, and ensuring port security.
  - Policing
    - Traffic-shaping policies on the switch port for managing how much of the port bandwidth should be allocated to a specific traffic flow.
    - Out-of-profile markdown for packets that exceed bandwidth utilization limits.
  - Ingress queuing and scheduling
    - Two configurable ingress queues for user traffic (one queue can be the priority queue).
    - Weighted tail drop (WTD) as the congestion-avoidance mechanism for managing the queue lengths and providing drop precedences for different traffic classifications.
    - Shaped round robin (SRR) as the scheduling service for specifying the rate at which packets are sent to the stack or internal ring (sharing is the only supported mode on ingress queues).
  - Egress queues and scheduling
    - Four egress queues per port.
    - WTD as the congestion-avoidance mechanism for managing the queue lengths and providing drop precedences for different traffic classifications.
    - SRR as the scheduling service for specifying the rate at which packets are dequeued to the egress interface (shaping or sharing is supported on egress queues).

- Automatic quality of service (QoS) voice over IP (VoIP) enhancement for port-based trust of DSCP and priority queuing for egress traffic.
- Egress policing and scheduling of egress queues - four egress queues on all switch ports; support for strict priority and weighted round-robin (WRR) CoS policies.

► Layer 3 features

– HSRP for Layer 3 router redundancy.

– IP routing protocols for load balancing and for constructing scalable, routed backbones:

- RIP Versions 1 and 2.
- OSPF (IP services feature set is required).
- Enhanced IGRP (EIGRP) (IP services feature set is required).
- Border Gateway Protocol (BGP) Version 4 (IP services feature set is required).

– IP routing between VLANs (inter-VLAN routing) for full Layer 3 routing between two or more VLANs, allowing each VLAN to maintain its own autonomous data-link domain.

– Policy-based routing (PBR) for configuring defined policies for traffic flows (IP services feature set is required).

– VPNs (IP services feature set is required).

– Fallback bridging for forwarding non-IP traffic between two or more VLANs (IP services feature set is required).

– Static IP routing for manually building a routing table of network path information.

– Equal-cost routing for load-balancing and redundancy.

– Internet Control Message Protocol (ICMP) and ICMP Router Discovery Protocol (IRDP) for using router advertisement and router solicitation messages to discover the addresses of routers on directly-attached subnets.

– Protocol-Independent Multicast (PIM) for multicast routing within the network, allowing for devices in the network to receive the multicast feed requested and for switches not participating in the multicast to be pruned. Includes support for PIM sparse mode (PIM-SM), PIM dense mode (PIM-DM), and PIM sparse-dense mode (IP services feature set is required).

– Multicast Source Discovery Protocol (MSDP) for connecting multiple PIM-SM domains (IP services feature set is required).

– Distance Vector Multicast Routing Protocol (DVMRP) tunneling for interconnecting two multicast-enabled networks across nonmulticast networks (an IP services feature set is required).

– DHCP relay for forwarding UDP broadcasts, including IP address requests, from DHCP clients.

– IPv6 support

- IPv6 host support (IPv6 unicast addressing, IPv6 traffic processing, IPv6 applications support including DNS, ping, traceroute, telnet, ftp, tftp, http, and ssh). IPv6 traffic forwarding is not supported. IPv6 host support is incorporated into an IP Base software feature set that comes standard with this switch module.
- IPv4 and IPv6 coexistence. The switch module supports dual IPv4 and IPv6 protocol stacks to provide seamless step-by-step migration to an IPv6 environment.
- IPv6 unicast routing capability (IPv6 traffic forwarding, static routes, RIP and OSPF) for forwarding IPv6 traffic through configured interfaces (an advanced IP services feature set is required).
- Support for EIGRP IPv6, which utilizes IPv6 transport, communicates with IPv6 peers, and advertises IPv6 routes (an advanced IP services feature set is required).

    - Support for IPv6 Access Control Lists (ACLs) (an advanced IP services feature set is required).
  - IP unicast reverse path forwarding (unicast RPF) for confirming source packet IP addresses.
  - Nonstop forwarding (NSF) awareness to enable the Layer 3 switch to continue forwarding packets from an NSF-capable neighboring router when the primary route processor (RP) is failing and the backup RP is taking over, or when the primary RP is manually reloaded for a nondisruptive software upgrade (an IP services feature set is required).
  - NSF-capable routing for OSPF and EIGRP that allows the switch to rebuild routing tables based on information from NSF-aware and NSF-capable neighbors (an IP services feature set is required).
- Monitoring
  - Switch LEDs that provide visual port, switch, and stack-level status.
  - SPAN/RSPAN support for local and remote monitoring of the network.
  - Four groups (history, statistics, alarms, and events) of embedded remote monitoring (RMON) agents for network monitoring and traffic analysis.
  - MAC address notification for tracking the MAC addresses that the switch has learned or removed.
  - Syslog facility for logging system messages about authentication or authorization errors, resource issues, and time-out events.
  - Layer 2 trace route to identify the physical path that a packet takes from a source device to a destination device.
  - Time Domain Reflector (TDR) to diagnose and resolve cabling problems on 10/100 and 10/100/1000 copper Ethernet ports.
  - Online diagnostics to test the hardware functionality of the supervisor engine, modules, and switch while the switch is connected to a live network.
  - On-board failure logging (OBFL) to collect information about the switch and the power supplies connected to it.
  - Enhanced object tracking (EOT) for HSRP to determine the proportion of hosts in a LAN by tracking the routing table state or to trigger the standby router failover.
  - IP Service Level Agreements (IP SLAs) support to measure network performance by using active traffic monitoring (IP services feature set is required).
- Network cables:
  - 10BASE-T:
    - UTP Category 3, 4, 5 (100 meters (328 feet) maximum)
    - 100-ohm STP (100 meters maximum)
  - 100BASE-TX:
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568 100-ohm STP (100 meters maximum)
  - 1000BASE-T:
    - UTP Category 6
    - UTP Category 5e (100 meters maximum)
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568B 100-ohm STP (100 meters maximum)

The Cisco Catalyst Switch Module 3110G supports the following IEEE standards:

- IEEE 802.1d Spanning Tree Protocol (STP)
- IEEE 802.1s Multiple STP (MSTP)
- IEEE 802.1w Rapid STP (RSTP)
- IEEE 802.1p CoS prioritization
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1x port-based authentication
- IEEE 802.2 Logical Link Control
- IEEE 802.3 10BASE-T Ethernet
- IEEE 802.3u 100BASE-TX Fast Ethernet
- IEEE 802.3ab 1000BASE-T Gigabit Ethernet
- IEEE 802.3z 1000BASE-X
- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3x Full-duplex Flow Control on all ports

Refer to the following documentation for more information:

- *Cisco Catalyst Switch Module 3110 and 3012 System Message Guide*
- *Cisco Catalyst Switch Module 3110 and 3012 Software Configuration Guide*
- *Cisco Catalyst Switch Module 3110G, 3110X, and 3012 Hardware Installation Guide*
- *Cisco Catalyst Switch Module 3110G, 3110X, and 3012 Getting Started Guide*
- *Cisco Catalyst Switch Module 3110 and 3012 Command Reference*

These publications can be downloaded from:

`http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5075938`

2.2.4 Cisco Catalyst Switch Module 3110X

The Cisco Catalyst Switch Module 3110X, part number 41Y8522, feature code 2988, is an I/O module that is installed into a BladeCenter unit. This switch is a full wire-rated, non-blocking switch for use with high performance servers.

This switch provides a next generation networking solution for blade server environments. Built upon Cisco's market leading hardware and IOS software, the switches are engineered with technologies specifically designed to help meet the rigors of blade server-based application infrastructure. Specifically, the switches are designed to deliver scalable, high performance, highly resilient connectivity while supporting ongoing IT initiatives around reducing server infrastructure complexity and TCO by seamlessly integrating into existing Cisco management networks. The Cisco Catalyst Switch Module 3110X is shown in Figure 2-8 on page 90.

Note: In addition to the IBM sales channel, this switch module can also be ordered from Cisco Systems resellers or directly from Cisco Systems (part number WS-CBS3110X-S-I).

Figure 2-9 Cisco Catalyst Switch Module 3110X with optional X2 transceiver module installed

The Cisco Catalyst Switch Module 3110X comes standard with IP Base feature set software. Additional features, such as IPv6 forwarding and routing support, as well as advanced routing protocols support (EIGRP, OSPF, BGP, PIM, and so on) for standard IP (IPv4), require additional licenses as listed in Table 2-12.

Table 2-12 Additional feature sets for Cisco Catalyst Switch Module 3110X

Part number	Feature code	Cisco Systems part number	Description
43W4434	4901	3110-IPS-LIC-I	IP Services S/W Upgrade License for Cisco Catalyst 3110[a]
None	None	3110-AISK9-LIC-I	S/W Upgrade License for Cisco Catalyst 3110 to Advanced IP Services[b]

a. Provides support for advanced routing protocols including EIGRP, OSPF, BGP, PIM, and so on. Can be ordered through standard IBM sales channels or from Cisco Systems sales channels.
b. Provides support for IPv6 forwarding and routing. Available through Cisco Systems sales channels only.

The supported features and specifications for the Cisco Catalyst 3110X switch module are:

- Ports
  - USB-style serial port. This is the Cisco console port, and offers a totally out-of-band management path if desired. USB-to-DB-9 cable is used to connect the switch module to PC. This cable is shipped with the switch.
  - One external 10 Gb Ethernet Module slot for making 10000 Mbps uplinks to a backbone switches or routers. This module slot operates at full-duplex and uses

hot-swappable Cisco X2 transceiver modules. The transceiver module is not included and must be ordered from a Cisco Systems reseller or directly from Cisco Systems.

The choice of X2 transceiver modules are:

- 10GBASE-CX4 X2 transceiver module for CX4 cable, copper, InfiniBand 4X connector (Cisco part number is X2-10GB-CX4=)
- 10GBASE-SR X2 transceiver module for MMF, 850-nm wavelength, SC duplex connector (Cisco part number is X2-10GB-SR=)
- 10GBASE-LRM X2 transceiver module for MMF, 1310-nm wavelength, SC duplex connector (Cisco part number is X2-10GB-LRM=)

Only one X2 transceiver module can be installed in the 10 Gb Ethernet Module slot.

– Two external high-speed StackWise Plus ports for switch module stacking to support Virtual Blade Switch (VBS) technology. Each 3110X switch module ships with one 1-meter StackWise Plus cable. Other cable lengths are available for order from Cisco Systems or Cisco Systems resellers if required:

- CAB-STK-E-0.5M= (0.5-meter cable)
- CAB-STK-E-1M= (1-meter cable)
- CAB-STK-E-3M= (3-meter cable)

Refer to 4.6, “Virtual Blade Switch technology” on page 384 for more information about VBS.

– 14 internal full-duplex Gigabit Ethernet ports (1000Base-X), one connected to each of the blade servers in the BladeCenter unit.

– One internal full-duplex 100 Mbps port connected to the management module.

► Performance features

– Fixed 10 Gbps speed on external 10 Gb Ethernet port for maximum uplink bandwidth.

– Up to 64 Gbps of throughput in a switch stack.

– 10-Gigabit EtherChannel for enhanced fault tolerance and for providing up to 80 Gbps of bandwidth between switches, routers, and servers.

– Support for standard frames with sizes from 64 to 1530 bytes and jumbo frames with a maximum size of 9216.

– Forwarding of Layer 2 frames and Layer 3 packets at 1 Gbps line rate across switches in the stack.

– Per-port broadcast-storm control for preventing a faulty end station from degrading overall system performance with broadcast storms.

– Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) for automatic creation of EtherChannel links.

– Internet Group Management Protocol (IGMP) snooping support to limit flooding of IP multicast traffic.

– Multicast Virtual Local Area Network (VLAN) registration (MVR) to continuously send multicast streams in a multicast VLAN while isolating the streams from subscriber VLANs for bandwidth and security.

– IGMP filtering for controlling the set of multicast groups to which hosts on a switch port can belong.

– Dynamic address learning for enhanced security.

– Supports multiple EtherChannel load balance algorithms (SMAC or DMAC, SIP or DIP, XOR-SMAC/DMAC or XOR-SIP/DIP) to offer maximum performance on aggregated links.

- Web Cache Communication Protocol (WCCP) for redirecting traffic to wide area application engines, for enabling content requests to be fulfilled locally, and for localizing Web traffic patterns in the network (supported by the IP Services feature set only).

- Manageability
  - Address Resolution Protocol (ARP) for identifying a switch through its IP address and its corresponding MAC address.
  - Cisco Discovery Protocol (CDP) Versions 1 and 2 to aid in troubleshooting and reporting on misconfiguration of ports connecting to other devices supporting CDP.
  - Link Layer Discovery Protocol (LLDP) and LLDP Media Endpoint Discovery (LLDP-MED) for interoperability with third-party IP phones.
  - Network Time Protocol (NTP) for providing a consistent time stamp to all switches from an external source.
  - Directed unicast requests to a Trivial File Transfer Protocol (TFTP) server for obtaining software upgrades from a TFTP server.
  - Default configuration storage in flash memory to ensure that the switch can be connected to a network and can forward traffic with minimal user intervention.
  - In-band monitoring of the switch through a built-in Cisco Device Manager Web-based tool.
  - In-band management access through up to 16 simultaneous Telnet connections for multiple command-line interface (CLI)-based sessions over the network.
  - In-band management access through up to five simultaneous, encrypted Secure Shell (SSH) connections for multiple CLI-based sessions over the network. This option is available only in the cryptographic software image.
  - In-band management access through SNMP versions 1, 2c, and 3 get and set requests.
  - Out-of-band management (CLI) with switch module's console port.
  - Supported by CiscoWorks management software.
  - Protected Mode feature to isolate switch management from Advanced Management Module, for increased security of the switch.
  - Cisco Network Services (CNS) embedded agents for automating switch management, configuration store, and delivery.
  - Cisco Network Assistance (CNA), a free GUI-based application tool to configure most features of this switch. For more information and to download CNA, go to:

    `http://www.cisco.com/go/cna`
  - Extensive debugging options to aid in troubleshooting and diagnosing issues.
  - Support for multiple management interfaces.

- Availability and redundancy
  - Hot Standby Routing Protocol (HSRP) for Layer 3 router redundancy.
  - Automatic stack master failover for replacing failed stack masters.
  - Cross-stack EtherChannel for providing redundant links across switch stack.
  - Link state tracking to mirror the state of the external ports on the internal Ethernet links and to allow the failover of the processor blade traffic to an operational external link on a separate Cisco Ethernet switch.

  - Configurable Unidirectional link detection (UDLD) for detecting and disabling unidirectional links. This feature prevents a larger network failure in the event that a unidirectional link is detected, thus reducing down time in these situations.
  - IEEE 802.1D Spanning Tree Protocol (STP) for redundant backbone connections and loop-free networks.
  - IEEE 802.1s Multiple STP (MSTP) for grouping VLANs into a spanning-tree instance, and provided for multiple forwarding paths for data traffic and load balancing.
  - IEEE 802.1w Rapid STP (RSTP) for rapid convergence of the spanning tree by immediately transitioning root and designated ports to the converting state.
  - Optional spanning-tree features available in the PVST+, rapid PVST+, and MSTP modes.
  - Flex Link Layer 2 interfaces to back up one another as an alternative to STP for basic link redundancy.
- VLAN support
  - Support for 1005 total VLANs. These VLANs can be any VLAN ID from 1–4094, except 1001–1005, which are reserved by Cisco.
  - Cisco Inter-Switch Link (ISL) and IEEE 802.1Q trunking protocol on all ports for network moves, adds, and changes; management and control of broadcast and multicast traffic; and network security by establishing VLAN groups for high-security users and network resources.
  - VLAN Query Protocol (VQP) for dynamic VLAN membership.
  - VLAN Trunking Protocol (VTP) pruning for reducing network traffic by restricting flooded traffic to links destined for stations receiving the traffic.
  - Dynamic Trunking Protocol (DTP) for negotiating trunking on a link between two devices and for negotiating the type of trunking encapsulation (802.1Q) to be used.
  - Voice VLAN for creating subnets for voice traffic from Cisco IP phones.
  - VLAN 1 minimization to reduce the risk of spanning-tree loops or storms by enabling VLAN 1 to be disabled on any individual VLAN trunk link. With this feature enabled, no user traffic is sent or received. The switch CPU continues to send and receive control protocol frames.
  - Private VLANs to address VLAN scalability issues.
  - VLAN Flex Link Load Balancing to provide Layer 2 link redundancy without STP.
  - Support for up to 128 instances of spanning tree per switch or per switch stack.
- Security
  - Bridge protocol data unit (BPDU) guard for shutting down a Port Fast-configured port when an invalid configuration occurs.
  - Protected port option for restricting the forwarding of traffic to designated ports on the same switch.
  - Password-protected access (read-only and write-only access) to management interfaces (the device manager and CLI) for protection against unauthorized configuration changes.
  - Port security option for limiting and identifying MAC addresses of the station allowed to access the port.
  - Port security aging to set the aging time for secure addresses on a port.
  - Multilevel security for a choice of security level, notification, and resulting actions.

- MAC-based port-level security for restricting the use of a switch port to a specific group of source addresses and preventing switch access from unauthorized stations.
- MAC-based access control lists (ACLs).
- Standard and extended IP access control lists (ACLs) for defining security policies on Layer 3 (router ACLs) and Layer 2 (port ACLs) interfaces.
- Terminal Access Controller Access Control System Plus (TACACS+), a proprietary feature for managing network security through a TACACS server.
- RADIUS for verifying the identity of, granting access to, and tracking activities of remote users.
- IEEE 802.1X port-based authentication to prevent unauthorized devices from gaining access to the network.
  - IEEE 802.1X port-based authentication with VLAN assignment for restricting 802.1X-authenticated users to a specified VLAN.
  - IEEE 802.1X port-based authentication with port security for authenticating the port and managing network access for all MAC addresses, including that of the client.
  - IEEE 802.1X port-based authentication with voice VLAN to allow an IP phone access to the voice VLAN irrespective of the authorized or unauthorized state of the port.
  - IEEE 802.1X port-based authentication with guest VLAN to provided limited services to non-802.1X-compliant users.
  - IEEE 802.1X accounting to track network usage.

- Quality of Service (QoS) and Class of Service (CoS)
  - Automatic QoS (auto-QoS) to simplify the deployment of existing QoS features by classifying traffic and configuring egress queues.
  - Cross-stack QoS for configuring QoS features to all switches in a switch stack rather than on an individual-switch basis.
  - Classification
    - IP Type of Service/Differentiated Services Code Point (IP ToS/DSCP) and IEEE 802.1p CoS marking priorities on a per-port basis for protecting the performance of mission-critical applications.
    - IP ToS/DSCP and IEEE 802.1p CoS marking for flow-based packet classification (classification based on information in the MAC, IP, and TCP/UDP headers) for high-performance QoS at the network edge, allowing for differentiated service levels for different types of network traffic and for prioritizing mission-critical traffic in the network.
    - Trusted port states (CoS, DSCP, and IP precedence) within a QoS domain and with a port bordering another QoS domain.
    - Trusted boundary for detecting the presence of a Cisco IP Phone, trusting the CoS value received, and ensuring port security.
  - Policing
    - Traffic-shaping policies on the switch port for managing how much of the port bandwidth should be allocated to a specific traffic flow.
    - Out-of-profile markdown for packets that exceed bandwidth utilization limits.
  - Ingress queuing and scheduling
    - Two configurable ingress queues for user traffic (one queue can be the priority queue).
    - Weighted tail drop (WTD) as the congestion-avoidance mechanism for managing the queue lengths and providing drop precedences for different traffic classifications.

- Shaped round robin (SRR) as the scheduling service for specifying the rate at which packets are sent to the stack or internal ring (sharing is the only supported mode on ingress queues).

– Egress queues and scheduling

- Four egress queues per port.
- WTD as the congestion-avoidance mechanism for managing the queue lengths and providing drop precedences for different traffic classifications.
- SRR as the scheduling service for specifying the rate at which packets are dequeued to the egress interface (shaping or sharing is supported on egress queues).
- Automatic quality of service (QoS) voice over IP (VoIP) enhancement for port-based trust of DSCP and priority queuing for egress traffic.
- Egress policing and scheduling of egress queues - four egress queues on all switch ports; support for strict priority and weighted round-robin (WRR) CoS policies.

► Layer 3 features

– HSRP for Layer 3 router redundancy.

– IP routing protocols for load balancing and for constructing scalable, routed backbones:

- RIP Versions 1 and 2
- OSPF (IP services feature set is required)
- Enhanced IGRP (EIGRP) (IP services feature set is required)
- Border Gateway Protocol (BGP) Version 4 (IP services feature set is required)

– IP routing between VLANs (inter-VLAN routing) for full Layer 3 routing between two or more VLANs, allowing each VLAN to maintain its own autonomous data-link domain.

– Policy-based routing (PBR) for configuring defined policies for traffic flows (IP services feature set is required).

– VPNs (IP services feature set is required).

– Fallback bridging for forwarding non-IP traffic between two or more VLANs (IP services feature set is required).

– Static IP routing for manually building a routing table of network path information.

– Equal-cost routing for load-balancing and redundancy.

– Internet Control Message Protocol (ICMP) and ICMP Router Discovery Protocol (IRDP) for using router advertisement and router solicitation messages to discover the addresses of routers on directly attached subnets.

– Protocol-Independent Multicast (PIM) for multicast routing within the network, allowing for devices in the network to receive the multicast feed requested and for switches not participating in the multicast to be pruned. Includes support for PIM sparse mode (PIM-SM), PIM dense mode (PIM-DM), and PIM sparse-dense mode (IP services feature set is required).

– Multicast Source Discovery Protocol (MSDP) for connecting multiple PIM-SM domains (IP services feature set is required).

– Distance Vector Multicast Routing Protocol (DVMRP) tunneling for interconnecting two multicast-enabled networks across nonmulticast networks (IP services feature set is required).

– DHCP relay for forwarding UDP broadcasts, including IP address requests, from DHCP clients.

- IPv6 support
  - IPv6 host support (IPv6 unicast addressing, IPv6 traffic processing, IPv6 applications support including DNS, ping, traceroute, telnet, ftp, tftp, http, and ssh). IPv6 traffic forwarding is not supported. IPv6 host support is incorporated into IP Base software feature set that comes standard with this switch module.
  - IPv4 and IPv6 coexistence. The switch module supports dual IPv4 and IPv6 protocol stacks to provide seamless step-by-step migration to an IPv6 environment.
  - IPv6 unicast routing capability (IPv6 traffic forwarding, static routes, RIP and OSPF) for forwarding IPv6 traffic through configured interfaces (advanced IP services feature set is required).
  - Support for EIGRP IPv6, which utilizes IPv6 transport, communicates with IPv6 peers, and advertises IPv6 routes (advanced IP services feature set is required).
  - Support for IPv6 Access Control Lists (ACLs) (an advanced IP services feature set is required).
- IP unicast reverse path forwarding (unicast RPF) for confirming source packet IP addresses.
- Nonstop forwarding (NSF) awareness to enable the Layer 3 switch to continue forwarding packets from an NSF-capable neighboring router when the primary route processor (RP) is failing and the backup RP is taking over, or when the primary RP is manually reloaded for a nondisruptive software upgrade (an IP services feature set is required).
- NSF-capable routing for OSPF and EIGRP that allows the switch to rebuild routing tables based on information from NSF-aware and NSF-capable neighbors (an IP services feature set is required).

► Monitoring

- Switch LEDs that provide visual port, switch, and stack-level status.
- SPAN/RSPAN support for local and remote monitoring of the network.
- Four groups (history, statistics, alarms, and events) of embedded remote monitoring (RMON) agents for network monitoring and traffic analysis.
- MAC address notification for tracking the MAC addresses that the switch has learned or removed.
- Syslog facility for logging system messages about authentication or authorization errors, resource issues, and time-out events.
- Layer 2 trace route to identify the physical path that a packet takes from a source device to a destination device.
- Online diagnostics to test the hardware functionality of the supervisor engine, modules, and switch while the switch is connected to a live network.
- On-board failure logging (OBFL) to collect information about the switch and the power supplies connected to it.
- Enhanced object tracking (EOT) for HSRP to determine the proportion of hosts in a LAN by tracking the routing table state or to trigger the standby router failover.
- IP Service Level Agreements (IP SLAs) support to measure network performance by using active traffic monitoring (an IP services feature set is required).

Note: Time Domain Reflector (TDR) is not supported on 10 Gb Ethernet ports.

- Network cables:
  - 10GBASE-SR, as described in 2.12, “BladeCenter S Office Enablement Kit” on page 209.

  Table 2-13 lists the 10GBASE-SR cabling specifications.

  Table 2-13 10GBASE-SR cabling specifications

Wavelength	Cable type	Core size (microns)	Modal bandwidth (MHz/km)	Maximum cable length
850 nm	MMF	62.5	160	85 feet (26 m)
850 nm	MMF	62.5	200	108 feet (33 m)
850 nm	MMF	50	400	217 feet (66 m)
850 nm	MMF	50	500	269 feet (82 m)
850 nm	MMF	50	2000	984 feet (300 m)

  - 10GBASE-LRM, as described in 2.13, “BladeCenter S storage modules” on page 211.

  Table 2-14 lists the 10GBASE-LRM cabling specifications.

  Table 2-14 10GBASE-LRM cabling specifications

Wavelength	Cable type	Core size (microns)	Modal bandwidth (MHz/km)	Maximum cable length
1310 nm	MMF	62.5	500	984 feet (300 m)
1310 nm	MMF	50	400	787 feet (240 m)
1310 nm	MMF	50	500	984 feet (300 m)

  - 10GBASE-CX4.

    The InfiniBand copper cable with 4X InfiniBand connector has a maximum cable length of 49 feet (15 m).

The Cisco Catalyst Switch Module 3110X supports the following IEEE standards:

- IEEE 802.1d Spanning Tree Protocol (STP)
- IEEE 802.1s Multiple STP (MSTP)
- IEEE 802.1w Rapid STP (RSTP)
- IEEE 802.1p CoS Prioritization
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1x port-based authentication
- IEEE 802.2 Logical Link Control
- IEEE 802.3 10BASE-T Ethernet
- IEEE 802.3u 100BASE-TX Fast Ethernet
- IEEE 802.3ab 1000BASE-T Gigabit Ethernet
- IEEE 802.3z 1000BASE-X Gigabit Ethernet
- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3x Full-duplex Flow Control on all ports
- IEEE 802.3ae 10GBASE-SR 10 Gb Ethernet
- IEEE 802.3ak 10GBASE-CX4 10 Gb Ethernet
- IEEE 802.3aq 10GBASE-LRM 10 Gb Ethernet

Refer to the following documentation for more information:

- *Cisco Catalyst Switch Module 3110 and 3012 System Message Guide*
- *Cisco Catalyst Switch Module 3110 and 3012 Software Configuration Guide*
- *Cisco Catalyst Switch Module 3110G, 3110X, and 3012 Hardware Installation Guide*
- *Cisco Catalyst Switch Module 3110G, 3110X, and 3012 Getting Started Guide*
- *Cisco Catalyst Switch Module 3110 and 3012 Command Reference*

These publications can be downloaded from:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5075938

2.2.5 Cisco Catalyst Switch Module 3012

The Cisco Catalyst Switch Module 3012, part number 43W4395, feature code 5450, is an I/O module that is installed into a BladeCenter unit. This switch is a full wire-rated, non-blocking switch for use with high performance servers.

This switch provides a next generation networking solution for blade server environments. Built upon Cisco's market-leading hardware and IOS software, the switches are engineered with technologies specifically designed to help meet the rigors of a blade server-based application infrastructure. Specifically, the switches are designed to deliver scalable, high performance, highly resilient connectivity while supporting ongoing IT initiatives around reducing server infrastructure complexity and TCO by seamlessly integrating into existing Cisco management networks.The Cisco Catalyst Switch Module 3012 is shown in Figure 2-10 on page 107.

Note: In addition to the IBM sales channel, this switch module can also be ordered from Cisco Systems resellers or directly from Cisco Systems (part number WS-CBS3012-IBM-I).

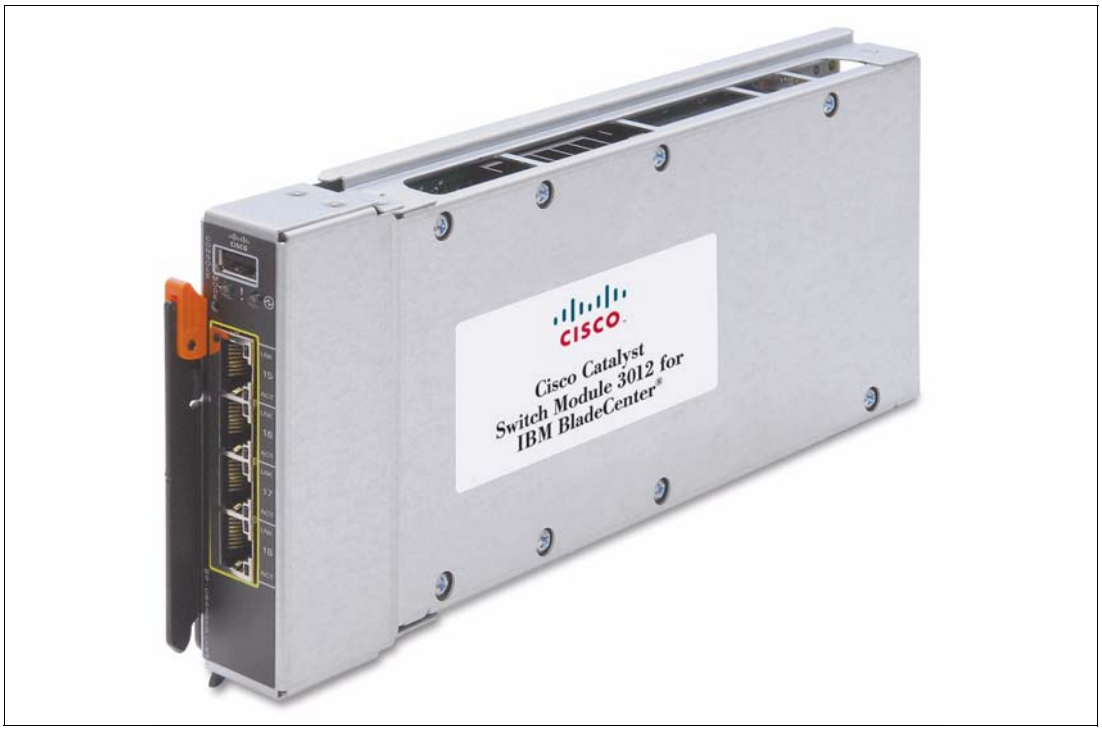

Figure 2-10 Cisco Catalyst Switch Module 3012

Note: Unlike the Cisco Catalyst 3110 family described previously in this document, which supports license upgrades, the Cisco Catalyst Switch Module 3012 supports the IP Base software feature set only.

The supported features and specifications for the Cisco Catalyst 3012 switch module are:

- Ports
  - USB-style serial port. This is the Cisco console port, and offers a totally out-of-band management path if desired. A USB-to-DB-9 cable is used to connect the switch module to PC. This cable is shipped with the switch.
  - Four external RJ-45 1000BASE-T connectors for making 10/100/1000 Mbps connections to a backbone, end stations, and servers.
  - 14 internal full-duplex Gigabit ports, one connected to each of the blade servers in the BladeCenter unit.
  - One internal full-duplex 100 Mbps port connected to the management module.
- Performance features
  - Autosensing of speed on the 10/100/1000 ports and auto-negotiation of duplex mode on the ports for optimizing bandwidth.
  - Gigabit EtherChannel for enhanced fault tolerance and for providing up to 4 Gbps of bandwidth between switches, routers, and servers.
  - Support for standard frames with sizes from 64 to 1530 bytes and jumbo frames with a maximum size of 9216.
  - Forwarding of Layer 2 frames and Layer 3 packets at 1 Gbps line rate.
  - Per-port broadcast-storm control for preventing a faulty end station from degrading overall system performance with broadcast storms.

- Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) for automatic creation of EtherChannel links.
- Internet Group Management Protocol (IGMP) snooping support to limit flooding of IP multicast traffic.
- Multicast Virtual Local Area Network (VLAN) registration (MVR) to continuously send multicast streams in a multicast VLAN, while isolating the streams from subscriber VLANs for bandwidth and security.
- IGMP filtering for controlling the set of multicast groups to which hosts on a switch port can belong.
- Dynamic address learning for enhanced security.
- Supports multiple EtherChannel load balance algorithms (SMAC or DMAC, SIP or DIP, XOR-SMAC/DMAC or XOR-SIP/DIP) to offer maximum performance on aggregated links.

► Manageability

- Address Resolution Protocol (ARP) for identifying a switch through its IP address and its corresponding MAC address.
- Cisco Discovery Protocol (CDP) Versions 1 and 2 to aid in troubleshooting and reporting on misconfiguration of ports connecting to other devices supporting CDP.
- Link Layer Discovery Protocol (LLDP) and LLDP Media Endpoint Discovery (LLDP-MED) for interoperability with third-party IP phones.
- Network Time Protocol (NTP) for providing a consistent time stamp to all switches from an external source.
- Directed unicast requests to a Trivial File Transfer Protocol (TFTP) server for obtaining software upgrades from a TFTP server.
- Default configuration storage in flash memory to ensure that the switch can be connected to a network and can forward traffic with minimal user intervention.
- In-band monitoring of the switch through the built-in Cisco Device Manager Web-based tool.
- In-band management access through up to 16 simultaneous Telnet connections for multiple command line interface (CLI)-based sessions over the network.
- In-band management access through up to five simultaneous, encrypted Secure Shell (SSH) connections for multiple CLI-based sessions over the network. This option is available only in the cryptographic software image.
- In-band management access through SNMP versions 1, 2c, and 3 get and set requests.
- Out-of-band management (CLI) with the switch module's console port.
- Supported by CiscoWorks management software.
- Protected Mode feature to isolate switch management from the Advanced Management Module, for increased security of the switch.
- Cisco Network Services (CNS) embedded agents for automating switch management, configuration store, and delivery.
- Cisco Network Assistance (CNA), a free GUI-based application tool to configure most features of this switch. For more information and to download CNA, go to:

  http://www.cisco.com/go/cna
- Extensive debugging options to aid in troubleshooting and diagnosing issues.

- Support for multiple management interfaces.

- Availability and redundancy
  - Hot Standby Routing Protocol (HSRP) for Layer 3 router redundancy.
  - Link state tracking to mirror the state of the external ports on the internal Ethernet links and to allow the failover of the processor blade traffic to an operational external link on a separate Cisco Ethernet switch.
  - Configurable Unidirectional link detection (UDLD) for detecting and disabling unidirectional links. This feature prevents a larger network failure in the event that a unidirectional link is detected, thus reducing downtime in these situations.
  - IEEE 802.1D Spanning Tree Protocol (STP) for redundant backbone connections and loop-free networks.
  - IEEE 802.1s Multiple STP (MSTP) for grouping VLANs into a spanning-tree instance, and provided for multiple forwarding paths for data traffic and load balancing.
  - IEEE 802.1w Rapid STP (RSTP) for rapid convergence of the spanning tree by immediately transitioning root and designated ports to the converting state.
  - Optional spanning-tree features available in the PVST+, rapid PVST+, and MSTP modes.
  - Flex Link Layer 2 interfaces to back up one another as an alternative to STP for basic link redundancy.

- VLAN support
  - Support for 1005 total VLANs. These VLANs can be any VLAN ID from 1–4094, except 1001–1005, which are reserved by Cisco.
  - Cisco Inter-Switch Link (ISL) and IEEE 802.1Q trunking protocol on all ports for network moves, adds, and changes; management and control of broadcast and multicast traffic; and network security by establishing VLAN groups for high-security users and network resources.
  - VLAN Query Protocol (VQP) for dynamic VLAN membership.
  - VLAN Trunking Protocol (VTP) pruning for reducing network traffic by restricting flooded traffic to links destined for stations receiving the traffic.
  - Dynamic Trunking Protocol (DTP) for negotiating trunking on a link between two devices and for negotiating the type of trunking encapsulation (802.1Q) to be used.
  - Voice VLAN for creating subnets for voice traffic from Cisco IP phones.
  - VLAN 1 minimization to reduce the risk of spanning-tree loops or storms by enabling VLAN 1 to be disabled on any individual VLAN trunk link. With this feature enabled, no user traffic is sent or received. The switch CPU continues to send and receive control protocol frames.
  - Private VLANs to address VLAN scalability issues.
  - VLAN Flex Link Load Balancing to provide Layer 2 link redundancy without STP.
  - Support for up to 128 instances of spanning tree per switch or per switch stack.

- Security
  - Bridge protocol data unit (BPDU) guard for shutting down a Port Fast-configured port when an invalid configuration occur.
  - Protected port option for restricting the forwarding of traffic to designated ports on the same switch.

- Password-protected access (read-only and write-only access) to management interfaces (the device manager and CLI) for protection against unauthorized configuration changes.
- Port security option for limiting and identifying MAC addresses of the station allowed to access the port.
- Port security aging to set the aging time for secure addresses on a port.
- Multilevel security for a choice of security level, notification, and resulting actions.
- MAC-based port-level security for restricting the use of a switch port to a specific group of source addresses and preventing switch access from unauthorized stations.
- MAC-based access control lists (ACLs).
- Standard and extended IP access control lists (ACLs) for defining security policies on Layer 3 (router ACLs) and Layer 2 (port ACLs) interfaces.
- Terminal Access Controller Access Control System Plus (TACACS+), a proprietary feature for managing network security through a TACACS server.
- RADIUS for verifying the identity of, granting access to, and tracking activities of remote users.
- IEEE 802.1X port-based authentication to prevent unauthorized devices from gaining access to the network.
  - IEEE 802.1X port-based authentication with VLAN assignment for restricting 802.1X-authenticated users to a specified VLAN.
  - IEEE 802.1X port-based authentication with port security for authenticating the port and managing network access for all MAC addresses, including that of the client.
  - IEEE 802.1X port-based authentication with voice VLAN to allow an IP phone access to the voice VLAN irrespective of the authorized or unauthorized state of the port.
  - IEEE 802.1X port-based authentication with guest VLAN to provided limited services to non-802.1X-compliant users.
  - IEEE 802.1X accounting to track network usage.

► Quality of Service (QoS) and Class of Service (CoS)

- Automatic QoS (auto-QoS) to simplify the deployment of existing QoS features by classifying traffic and configuring egress queues.
- Cross-stack QoS for configuring QoS features to all switches in a switch stack rather than on an individual-switch basis.
- Classification
  - IP Type of Service/Differentiated Services Code Point (IP ToS/DSCP) and IEEE 802.1p CoS marking priorities on a per-port basis for protecting the performance of mission-critical applications.
  - IP ToS/DSCP and IEEE 802.1p CoS marking for flow-based packet classification (classification based on information in the MAC, IP, and TCP/UDP headers) for high-performance QoS at the network edge, allowing for differentiated service levels for different types of network traffic and prioritizing mission-critical traffic in the network.
  - Trusted port states (CoS, DSCP, and IP precedence) within a QoS domain and with a port bordering another QoS domain.
  - Trusted boundary for detecting the presence of a Cisco IP Phone, trusting the CoS value received, and ensuring port security.

- Policing
  - Traffic-shaping policies on the switch port for managing how much of the port bandwidth should be allocated to a specific traffic flow.
  - Out-of-profile markdown for packets that exceed bandwidth utilization limits.
- Ingress queuing and scheduling
  - Two configurable ingress queues for user traffic (one queue can be the priority queue).
  - Weighted tail drop (WTD) as the congestion-avoidance mechanism for managing the queue lengths and providing drop precedences for different traffic classifications.
  - Shaped round robin (SRR) as the scheduling service for specifying the rate at which packets are sent to the stack or internal ring (sharing is the only supported mode on ingress queues).
- Egress queues and scheduling
  - Four egress queues per port.
  - WTD as the congestion-avoidance mechanism for managing the queue lengths and providing drop precedences for different traffic classifications.
  - SRR as the scheduling service for specifying the rate at which packets are dequeued to the egress interface (shaping or sharing is supported on egress queues).
  - Automatic quality of service (QoS) voice over IP (VoIP) enhancement for port-based trust of DSCP and priority queuing for egress traffic.
  - Egress policing and scheduling of egress queues - four egress queues on all switch ports; support for strict priority and weighted round-robin (WRR) CoS policies.

- Layer 3 features
  - HSRP for Layer 3 router redundancy.
  - IP routing protocols for load balancing and for constructing scalable, routed backbones:
    - RIP Versions 1 and 2

      Note: Unlike the Cisco Catalyst 3110 family of switches covered previously in this document, the Cisco Catalyst 3012 module does *not* support OSPF, BGP, and EIGRP routing protocols, policy-based routing, VPNs, and multicast routing,

  - IP routing between VLANs (inter-VLAN routing) for full Layer 3 routing between two or more VLANs, allowing each VLAN to maintain its own autonomous data-link domain.
  - Static IP routing for manually building a routing table of network path information.
  - Equal-cost routing for load-balancing and redundancy.
  - Internet Control Message Protocol (ICMP) and ICMP Router Discovery Protocol (IRDP) for using router advertisement and router solicitation messages to discover the addresses of routers on directly-attached subnets.
  - DHCP relay for forwarding UDP broadcasts, including IP address requests, from DHCP clients.
  - IPv6 support
    - IPv6 host support (IPv6 unicast addressing, IPv6 traffic processing, IPv6 applications support including DNS, ping, traceroute, telnet, ftp, tftp, http, and ssh). IPv6 traffic forwarding is not supported. IPv6 host support is incorporated into the IP Base software feature set that comes standard with this switch module.
    - IPv4 and IPv6 coexistence. The switch module supports dual IPv4 and IPv6 protocol stacks to provide seamless step-by-step migration to an IPv6 environment.

Note: Unlike the Cisco Catalyst 3110 family of switches covered previously in this document, the Cisco Catalyst 3012 module only supports IPv6 host functions, and does *not* support IPv6 routing (IPv6 traffic forwarding, IPv6 static routes, and OSPF, RIP, or EIGRP for IPv6), and IPv6 ACLs.

  - IP unicast reverse path forwarding (unicast RPF) for confirming source packet IP addresses.

- Monitoring
  - Switch LEDs that provide visual port, switch, and stack-level status.
  - SPAN/RSPAN support for local and remote monitoring of the network.
  - Four groups (history, statistics, alarms, and events) of embedded remote monitoring (RMON) agents for network monitoring and traffic analysis.
  - MAC address notification for tracking the MAC addresses that the switch has learned or removed.
  - Syslog facility for logging system messages about authentication or authorization errors, resource issues, and time out events.
  - Layer 2 trace route to identify the physical path that a packet takes from a source device to a destination device.
  - Time Domain Reflector (TDR) to diagnose and resolve cabling problems on 10/100 and 10/100/1000 copper Ethernet ports.
  - Online diagnostics to test the hardware functionality of the supervisor engine, modules, and switch while the switch is connected to a live network.
  - On-board failure logging (OBFL) to collect information about the switch and the power supplies connected to it.
  - Enhanced object tracking (EOT) for HSRP to determine the proportion of hosts in a LAN by tracking the routing table state or to trigger the standby router failover.
- Network cables:
  - 10BASE-T:
    - UTP Category 3, 4, 5 (100 meters (328 feet) maximum)
    - 100-ohm STP (100 meters maximum)
  - 100BASE-TX:
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568 100-ohm STP (100 meters maximum)
  - 1000BASE-T:
    - UTP Category 6
    - UTP Category 5e (100 meters maximum)
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568B 100-ohm STP (100 meters maximum)

The Cisco Catalyst Switch Module 3110G supports the following IEEE standards:

- IEEE 802.1d Spanning Tree Protocol (STP)
- IEEE 802.1s Multiple STP (MSTP)
- IEEE 802.1w Rapid STP (RSTP)
- IEEE 802.1p Tagged Packets
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1x port-based authentication

- IEEE 802.2 Logical Link Control
- IEEE 802.3 10BASE-T Ethernet
- IEEE 802.3u 100BASE-TX Fast Ethernet
- IEEE 802.3ab 1000BASE-T Gigabit Ethernet
- IEEE 802.3z 1000BASE-X Gigabit Ethernet
- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3x Full-duplex Flow Control on all ports

Refer to the following documentation for more information:

- *Cisco Catalyst Switch Module 3110 and 3012 System Message Guide*
- *Cisco Catalyst Switch Module 3110 and 3012 Software Configuration Guide*
- *Cisco Catalyst Switch Module 3110G, 3110X, and 3012 Hardware Installation Guide*
- *Cisco Catalyst Switch Module 3110G, 3110X, and 3012 Getting Started Guide*
- *Cisco Catalyst Switch Module 3110 and 3012 Command Reference*

These publications can be downloaded from:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5075938

2.2.6 IBM Server Connectivity Module

The IBM Server Connectivity Module for IBM BladeCenter, part number 39Y9324, feature code 1484, is a switch that provides basic Layer 2 functionality. This device is configurable by a non-networking system administrator through a graphical user interface (GUI) or a command line interface (CLI).

Typical networking concepts, such as spanning tree redundancy, virtual local area networks (VLANs), port filters, link aggregation, port trunks, and remote monitoring (RMON) statistics are not typically used or understood by this class of administrator.

The default operation will be to hide the underlying networking function and configuration from the user. Only a few simple networking parameters are available to the user to configure and control the Server Connectivity Module.

Figure 2-11 IBM Server Connectivity Module

Standard features and functions include:

- Internal ports
  - 14 internal full-duplex Gigabit ports, one connected to each of the blade servers in the BladeCenter unit
  - Two internal full-duplex 100 Mbps ports connected to the management module in slots 1 and 2
- External copper ports
  - Six external ports are provided using standard RJ-45 copper cable connectors
  - 10/100/1000 Mbps interfaces: the external ports connect at 10 Mbps Full Duplex, 100 Mbps Full Duplex, or 1 Gbps Full Duplex
  - Auto-negotiation capability
  - Copper connection for making connections to a backbone, end stations, and servers
- Internal Switching
  - Packet switching between the blade servers and management modules within the chassis to support Serial over LAN (SOL)
  - Blade-server-to-blade-server switching is supported for those blade server ports that are assigned to a common group (up to six groups with dedicated uplinks)
- Configuration of the Server Connectivity Module is only through a connection to a management-module port and not through the external switch ports
- Support for tagged VLANs: the administrator can define VLAN identifiers (IDs) to match specific server application requirements
- Cisco EtherChannel compatible static link aggregation
- Management-module control that allows for external ports to be enabled or disabled
- Support for Ethernet jumbo frame formats (up to 9 K bytes per frame)

- Two independent internal 100 Mbps FDX links connected to each of two management modules (jumbo frame support is not required on these links)
- Management-module I2C interface that provides VPD and register access
- Level 2 (L2) switching per current industry standards and practice
- “Non-networking” device appearance allows administration and installation of this device by the system manager rather than a network administrator
- Port aggregation (only for external ports)
- Port group failover (triggered by external ports)
- IEEE 802.3x Flow Control
- Internet group multicast protocol (IGMP) snooping
- IEEE 802.1Q support for SOL and optional user VLAN tagging of external ports
- RADIUS or TACACS+ user authentication
- Network cables
  - 10BASE-T:
    - UTP Category 3, 4, 5 (100 meters or 328.1 ft maximum)
    - 100-ohm STP (100 meters or 328.1 ft maximum)
  - 100BASE-TX:
    - UTP Category 5 (100 meters or 328.1 ft maximum)
    - EIA/TIA-568 100-ohm STP (100 meters or 328.1 ft maximum)
  - 1000BASE-T:
    - UTP Category 6
    - UTP Category 5e (100 meters or 328.1 ft maximum)
    - UTP Category 5 (100 meters or 328.1 ft maximum)
    - EIA/TIA-568B 100-ohm STP (100 meters or 328.1 ft maximum)

The Server Connectivity Module supports the following IEEE standards:

- IEEE 802.3x
- IEEE 802.1Q (tagged) VLANs
- IEEE 802.3ab 1000BASE-T Gigabit Ethernet

Refer to the following documentation for more information:

- *IBM Server Connectivity Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-64155
- IBM *Server Connectivity Module User's Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-64190

Note: If you plan to use this switch with HC10 workstation blades, then ensure you disable IGMP snooping. See RETAIN® tip H192042 for details:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5073011

2.2.7 Nortel Networks L2/3 Copper GbE Switch Module

The Nortel Networks Layer 2/3 Copper Gigabit Ethernet Switch Module, part number 32R1860, feature code 1495, is a switch option that enable administrators to consolidate full Layer 2-3 LAN switching and routing capabilities into a BladeCenter chassis. Consolidation

flattens the topology of the datacenter infrastructure and reduces the number of discrete devices, management consoles, and different management systems.

Figure 2-12 Nortel Networks L2/3 Copper GbE Switch Module

The switch module supports two unique operational modes which are determined by the firmware edition loaded onto the switch:

- Original or Standard Layer 2/3 operations
- BladeCenter Open Fabric Manager (OFM) with Nortel/BNT Extensions operations

Standard Layer 2/3 mode provides full implementation of networking services and technologies supported by the switch, and OFM for Nortel/BNT Extensions provides a preconfigured subset of functions to enable basic Layer 2 connectivity only.

Note: BladeCenter Open Fabric Manager for Nortel/BNT Extensions is a firmware upgrade to this Nortel switch and is different from the BladeCenter Open Fabric Manager offering from IBM.

Features and functions of the Nortel L2/3 GbE switch module for standard mode of operations include:

- Internal ports
  - 14 internal full-duplex Gigabit ports, one connected to each of the blade servers in the BladeCenter unit
  - Two internal full-duplex 10/100 Mbps ports connected to the management module
- External ports
  - Six 1000BASE-T copper RJ-45 connections for making 10/100/1000 Mbps connections to a backbone, end stations, and servers
  - An RS-232 serial port that provides an additional means to install software and configure the switch module

- Scalability and performance
  - Autosensing 10/1000/1000 Mbps external Ethernet ports for bandwidth optimization
  - Non-blocking architecture with wire-speed forwarding of traffic
  - Media access control (MAC) address learning: automatic update, supports up to 16 K MAC addresses
  - Up to 128 IP interfaces per switch
  - Static and LACP (IEEE 802.3ad) link aggregation, up to 6 Gb of total bandwidth per switch, up to three trunk groups, up to six ports per group
  - Support for jumbo frames (up to 9216 bytes)
  - Broadcast/multicast storm control
  - IGMP snooping for limit flooding of IP multicast traffic
  - IGMP filtering to control multicast traffic for hosts participating in multicast groups
  - Configurable traffic distribution schemes over trunk links based on source/destination IP or MAC addresses or both
  - Fast port forwarding and fast uplink convergence for rapid STP convergence
- Availability and redundancy
  - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy
  - IEEE 802.1D Spanning Tree Protocol (STP) for providing Layer 2 redundancy
  - IEEE 802.1s Multiple STP (MSTP) for topology optimization, up to 128 STP instances are supported by single switch
  - IEEE 802.1w Rapid STP (RSTP) provides rapid STP convergence for critical delay-sensitive, traffic-like voice or video
  - Layer 2 Trunk Failover to support active/standby configurations of network adapter teaming on blades
  - Interchassis redundancy (Layer 2 and Layer 3)
- VLAN support
  - Up to 1024 VLANs are supported per switch, VLAN numbers range is from 1 to 4095 (4095 is used for the management module's connection only)
  - 802.1Q VLAN tagging support on all ports
  - Private VLANs
- Security
  - VLAN-based, MAC-based, and IP-based Access Control Lists (ACLs)
  - 802.1X port-based authentication
  - Multiple user IDs and passwords
  - User access control
  - Radius/TACACS+
- Quality of Service (QoS)
  - Support for IEEE 802.1p, IP ToS/DSCP, and ACL-based (MAC/IP source and destination addresses, VLANs) traffic classification and processing
  - Traffic shaping and re-marking based on defined policies
  - Eight Weighted Round Robin (WRR) priority queues per port for processing qualified traffic

- Layer 3 functions
  - IP forwarding
  - IP filtering with ACLs, up to 4096 ACLs are supported
  - VRRP for router redundancy
  - Support for up to 128 static routes
  - Routing protocol support (RIP v1, RIP v2, OSPF v2, BGP-4), up to 1024 entries in routing table
  - Support for DHCP Relay
- Manageability
  - SNMP (V1 and V3)
  - HTTP browser graphical user interface (GUI)
  - Telnet interface for CLI
  - SSH
  - Serial interface for CLI
  - Scriptable command line interface (CLI)
  - Firmware image update (TFTP and FTP)
  - Network Time Protocol (NTP) for switch clock synchronization
  - BNT BLADEHarmony Manager support
- Monitoring
  - Switch LEDs for external port status and switch module status indication
  - Port mirroring for analyzing network traffic passing through switch
  - Change tracking and remote logging with syslog feature
  - POST diagnostics
- Special functions
  - Serial over LAN (SOL)
- Network cables
  - 10BASE-T:
    - UTP Category 3, 4, 5 (100 meters (328 feet) maximum)
    - 100-ohm STP (100 meters maximum)
  - 100BASE-TX:
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568 100-ohm STP (100 meters maximum)
  - 1000BASE-T:
    - UTP Category 6
    - UTP Category 5e (100 meters maximum)
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568B 100-ohm STP (100 meters maximum)
- RS-232 serial cable: 3-meter console cable DB-9 to USB connector (nonstandard use of USB connector) that comes with the GbE switch module

The GbE switch module supports the following IEEE standards:

- IEEE 802.1D Spanning Tree Protocol (STP)

- IEEE 802.1s Multiple STP (MSTP)
- IEEE 802.1w Rapid STP (RSTP)
- IEEE 802.1p Tagged Packets
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1x port-based authentication
- IEEE 802.2 Logical Link Control
- IEEE 802.3 10BASE-T Ethernet
- IEEE 802.3u 100BASE-TX Fast Ethernet
- IEEE 802.3ab 1000BASE-T Gigabit Ethernet
- IEEE 802.3z 1000BASE-X Gigabit Ethernet
- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3x Full-duplex Flow Control

OFM with Nortel/BNT Extensions mode is used for less complex environments where only basic network connectivity is required and there is no need for advanced networking services. OFM with Extensions mode's reduced set of features and easy-to-use GUI help to minimize administrative efforts and skills required to deploy this solution, as well as the impact (both from management and functional point of views) on upstream network devices.

One of the features of OFM with Nortel/BNT Extensions is ServerMobility, which enables you to assign IP addresses to servers based on their physical location in the BladeCenter chassis. This may help to keep server identity based on IP address in basic failover scenarios. For example, in you replace a failed server with a new unit, this new unit will have the same IP address as the failed server. Or, if the server fails, then the other server in the chassis gets the IP address of the failed server. A DHCP server is required for ServerMobility operations.

Specifically, the following functions are supported in OFM with Nortel/BNT Extensions mode:

- Port groups
- VLAN and VLAN tagging
- Trunking (link aggregation)
- Layer 2 Failover (Switch Failover)
- IGMP Snooping
- Web-based or CLI-based management interface
- Support for SSH
- RADIUS/TACACS+ access control to the switch
- Port monitoring
- ServerMobility

Notes:

- When the switch operates in OFM Extensions mode, do not connect the host devices such as workstations or servers to the external ports because this will lead to problems with network traffic flow.
- Spanning Tree Protocol is not supported in OFM Extensions mode. Layer 2 loop topologies *must* be avoided.

Refer to the following documentation for more information:

- *Nortel Networks Layer 2-3 GbE Switch Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58334
- *Nortel Networks Layer 2-3 GbE Switch Module Application Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58325
- *BladeCenter OFM User's Guide for Nortel L2-3 GbESM*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5074970

2.2.8 Nortel Networks L2/3 Fiber GbE Switch Module

The Nortel Networks Layer 2/3 Fiber Gigabit Ethernet Switch Module, part number 32R1861, feature code 1496, is a switch option that enable administrators to consolidate full Layer2-3 LAN switching and routing capabilities into a BladeCenter chassis. Consolidation flattens the topology of the datacenter infrastructure and reduces the number of discrete devices, management consoles, and different management systems.

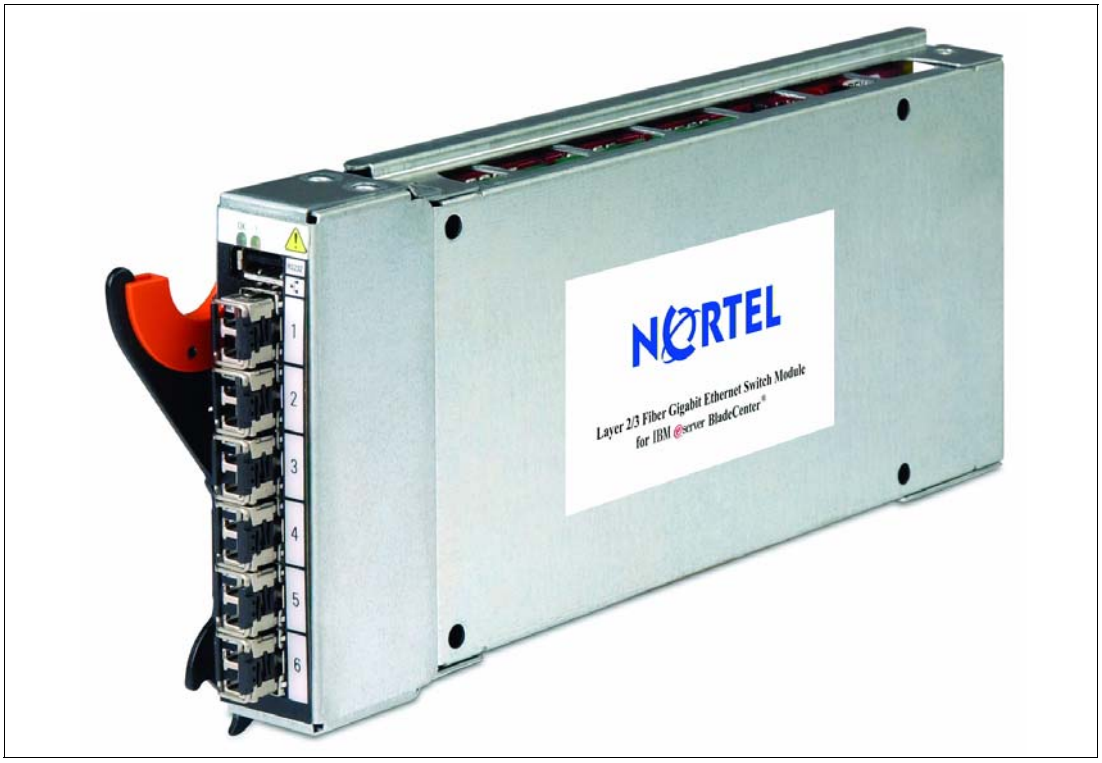

Figure 2-13 Nortel Networks L2/3 Fiber GbE Switch Module

Note: Six SFP modules for SX connections come standard with the Nortel L2/3 Fiber GbESM.

Standard features and functions include:

- Internal ports
  - 14 internal full-duplex Gigabit ports, one connected to each of the blade servers in the BladeCenter unit
  - Two internal full-duplex 10/100 Mbps ports connected to the management module
- External ports
  - Six 1000BASE-SX SFP tranceiver-based LC fiber connections for making 10/100/1000 Mbps connections to a backbone, end stations, and servers
  - RS-232 serial port that provides an additional means to install software and configure the switch module
- Performance
  - Autosensing 10/1000/1000 Mbps external Ethernet ports for bandwidth optimization
  - Non-blocking architecture with wire-speed forwarding of traffic

- Media access control (MAC) address learning: automatic update, supports up to 16 K MAC addresses
- Up to 128 IP interfaces per switch
- Static and LACP (IEEE 802.3ad) link aggregation, up to 6 Gb of total bandwidth per switch, up to three trunk groups, up to six ports per group
- Support for jumbo frames (up to 9216 bytes)
- Broadcast/multicast storm control
- IGMP snooping for limit flooding of IP multicast traffic
- IGMP filtering to control multicast traffic for hosts participating in multicast groups
- Configurable traffic distribution schemes over trunk links based on source/destination IP or MAC addresses or both
- Fast port forwarding and fast uplink convergence for rapid STP convergence

► Availability and redundancy
- Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy
- IEEE 802.1D Spanning Tree Protocol (STP) for providing L2 redundancy
- IEEE 802.1s Multiple STP (MSTP) for topology optimization, up to 128 STP instances are supported by single switch
- IEEE 802.1w Rapid STP (RSTP) provides rapid STP convergence for critical delay-sensitive traffic like voice or video
- Layer 2 Trunk Failover to support active/standby configurations of network adapter teaming on blades
- Interchassis redundancy (L2 and L3)

► VLAN support
- Up to 1024 VLANs are supported per switch, VLAN numbers range is from 1 to 4095 (4095 is used for the management module's connection only)
- 802.1Q VLAN tagging support on all ports
- Private VLANs

► Security
- VLAN-based, MAC-based, and IP-based Access Control Lists (ACLs)
- 802.1X port-based authentication
- Multiple user IDs and passwords
- User access control
- Radius/TACACS+

► Quality of Service (QoS)
- Support for IEEE 802.1p, IP ToS/DSCP, and ACL-based (MAC/IP source and destination addresses, VLANs) traffic classification and processing
- Traffic shaping and re-marking based on defined policies
- Eight Weighted Round Robin (WRR) priority queues per port for processing qualified traffic

► Layer 3 functions
- IP forwarding
- IP filtering with ACLs

  - VRRP for router redundancy
  - Support for up to 128 static routes
  - Routing protocol support (RIP v1, RIP v2, OSPF v2, BGP-4), up to 1024 entries in routing table
  - Support for DHCP Relay
- Manageability
  - SNMP (V1 and V3)
  - HTTP browser graphical user interface (GUI)
  - Telnet interface for CLI
  - SSH
  - Serial interface for CLI
  - Scriptable command line interface (CLI)
  - Firmware image update (TFTP and FTP)
  - Network Time Protocol (NTP) for switch clock synchronization
  - BNT BLADEHarmony Manager support
- Monitoring
  - Switch LEDs for external port status and switch module status indication
  - Port mirroring for analyzing network traffic passing through switch
  - Change tracking and remote logging with syslog feature
  - POST diagnostics
- Special functions
  - Serial over LAN (SOL)
- Network cables
  - 1000BASE-SX: 850 Nm wavelength, multimode fiber, 50 µ or 62.5 µ (550 meters maximum), with LC duplex connector
- RS-232 serial cable: 3 meter console cable DB-9 to USB connector (nonstandard use of USB connector) that comes with the GbE switch module

The GbE switch module supports the following IEEE standards:

- IEEE 802.1D Spanning Tree Protocol (STP)
- IEEE 802.1s Multiple STP (MSTP)
- IEEE 802.1w Rapid STP (RSTP)
- IEEE 802.1p Tagged Packets
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1x port-based authentication
- IEEE 802.2 Logical Link Control
- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3x Full-duplex Flow Control
- IEEE 802.3z 1000BASE-X Gigabit Ethernet
  - 1000BASE-SX Gigabit Ethernet
  - Link Negotiation

Note: Fiber connections use small-form-factor pluggable (SFP) transceivers that provide 1000BASE-SX (850 nm wavelength) communications over multimode fiber cables (50µ or 62.5µ) for distances of up to 550 meters (1000BASE-LX is not supported).

Refer to the following documentation for more information:

- *Nortel Networks Layer 2-3 GbE Switch Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58334
- *Nortel Networks Layer 2-3 GbE Switch Module Application Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58325

2.2.9 Nortel Networks Layer 2-7 Gigabit Ethernet Switch Module

The Nortel Networks Layer 2-7 Gigabit ESM, part number 32R1859, feature code 1494, serves as a switching and routing fabric for the IBM BladeCenter server chassis. It also introduces Layer 4-7 functionality for application- and server-based load balancing, advanced filtering, content-aware intelligence, embedded security services, and persistence support.

Figure 2-14 Nortel Networks Layer 2-7 GbE Switch Module

The Nortel Networks Layer 2-7 GbE Switch Module has the following features:

- Ports
  - Four external RJ-45 1000BASE-T connectors for making 10/100/1000 Mbps connections to a backbone, end stations, and servers
  - 14 internal full-duplex gigabit ports, one connected to each of the blade servers in the BladeCenter unit
  - Two internal full-duplex 10/100 Mbps ports connected to the management module
  - One RS-232 serial port that provides an additional means to install software and configure the switch module

- Scalability and performance
  - Autosensing 10/1000/1000 Mbps external Ethernet ports for bandwidth optimization
  - Non-blocking architecture with wire-speed forwarding of traffic
    - 148800 packets per second (pps) per port (for 100 Mbps)
    - 1488100 pps per port (for 1000 Mbps)
  - Media access control (MAC) address learning: automatic update, supports 2048 MAC addresses
  - Up to 128 IP interfaces per switch
  - Static and LACP (IEEE 802.3ad) link aggregation, up to 4 Gb of total bandwidth per switch, up to 13 trunk groups, up to four ports per group
  - Support for jumbo frames (up to 9216 bytes)
  - Broadcast/multicast storm control
  - IGMP snooping for limit flooding of IP multicast traffic
  - IGMP filtering to control multicast traffic for hosts participating in multicast groups
  - Configurable traffic distribution schemes over trunk links based on source/destination IP or MAC addresses or both
  - Fast port forwarding and fast uplink convergence for rapid STP convergence
  - Up to 64 real servers and up to 256 real services, up to 64 virtual servers, and up to 256 virtual services per switch for Server Load Balancing (SLB) configurations to increase application performance
  - 300 000 simultaneous Layer 2 through Layer 7 (L2-L7) sessions
- Availability and redundancy
  - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy
  - IEEE 802.1D Spanning Tree Protocol (STP) for providing L2 redundancy
  - IEEE 802.1s Multiple STP (MSTP) for topology optimization; up to 32 STP instances are supported by single switch
  - Layer 2 Trunk Failover to support active/standby configurations of network adapter teaming on blades
  - Server Load Balancing provides load distribution among servers in farm and automatic failover of user sessions in case of server or application failure
  - Real server health check for server or application status and content availability
  - Interchassis redundancy (L2-L7)
- VLAN support
  - Up to 1024 VLANs are supported per switch, VLAN numbers range is from 1 to 4095 (4095 is used for management module's connection only)
  - 802.1Q VLAN tagging support on all ports
- Security
  - IP-based (source/destination IP addresses, protocols, source/destination ports) filtering
  - Network Address Translation (NAT)
  - Denial of Service (DoS) attack prevention
  - Multiple user IDs and passwords

  - User access control
  - Radius/TACACS+
- Layer 3 functions
  - IP forwarding
  - IP filtering with ACLs, up to 1024 filters
  - VRRP for router redundancy
  - Support for static routes, up to 128 route entries
  - Routing protocol support (RIP v1, RIP v2, OSPF v2, BGP-4), up to 2048 route entries
  - Support for DHCP Relay
- Layer 4-7 functions
  - Server Load Balancing (SLB) for increased performance, availability and fault tolerance
  - Global SLB for load balancing across multiple physical sites
  - Intelligent cache redirection for HTTP
  - IBM Enterprise Workload Manager™ support
  - URL and cookie content-based load balancing for HTTP requests
  - Content-based load balancing for DNS requests
  - HTTP cookie and secure sockets layer (SSL) session ID persistency
- Manageability
  - SNMP (V1 and V3)
  - HTTP browser graphical user interface (GUI)
  - Telnet interface for CLI
  - SSH
  - Serial interface for CLI
  - Scriptable command line interface (CLI)
  - Firmware image update (TFTP and FTP)
  - Network Time Protocol (NTP) for switch clock synchronization
  - BNT BLADEHarmony Manager support
- Monitoring
  - Switch LEDs for external port status and switch module status indication
  - Port mirroring for analyzing network traffic passing through switch
  - Change tracking and remote logging with syslog feature
  - RMON support
  - POST diagnostics
- Special functions
  - Serial over LAN (SOL)
- Network cables
  - 10BASE-T:
    - UTP Category 3, 4, 5 (100 meters (328 feet) maximum)
    - 100-ohm STP (100 meters maximum)

  - 100BASE-TX:
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568 100-ohm STP (100 meters maximum)
  - 1000BASE-T:
    - UTP Category 6
    - UTP Category 5e (100 meters maximum)
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568B 100-ohm STP (100 meters maximum)
- RS-232 serial cable: 3-meter console cable DB-9-to-USB connector (nonstandard use of USB connector) that comes with the GbE switch module

Note: An RS-232 Serial Cable is included with the switch.

The GbE switch module supports the following IEEE standards:

- IEEE 802.1D Spanning Tree Protocol (STP)
- IEEE 802.1s Multiple STP (MSTP)
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.2 Logical Link Control
- IEEE 802.3 10BASE-T Ethernet
- IEEE 802.3u 100BASE-TX Fast Ethernet
- IEEE 802.3ab 1000BASE-T Gigabit Ethernet
- IEEE 802.3z 1000BASE-X Gigabit Ethernet
- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3x Full-duplex Flow Control

Refer to the following documentation for more information:

- *Nortel Networks Layer 2-7 GbE Switch Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-53065
- *Nortel Networks Layer 2-7 GbE Switch Module Application Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-53098

2.2.10 Nortel 10 Gb Uplink Ethernet Switch Module

The Nortel 10 Gigabit Uplink Ethernet Switch Module, part number 32R1783, feature code 1493, is a switch option that enables administrators to provide full Layer 2-3 10 Gb LAN switching and routing capabilities in a BladeCenter chassis. Such consolidation flattens the topology of the data center infrastructure and reduces the number of discrete devices, management consoles, and different management systems. The 10 Gigabit Uplink Ethernet Switch Module provides port flexibility, efficient traffic management, increased uplink bandwidth, and strong Ethernet switching price/performance.

Figure 2-15 Nortel 10 Gb Uplink Ethernet Switch Module with optional XFP Module installed

Standard features and functions include:

- Internal ports
  - 14 internal full-duplex Gigabit ports, one connected to each of the blade servers in the BladeCenter unit
  - Two internal full-duplex 10/100 Mbps ports connected to the management module
- External ports
  - Two 10 Gb CX4 ports
  - One 10 Gb XFP port (SR or LR) - XFP module is optional
  - One 1 Gb Ethernet RG-45 port
  - One RS-232 serial port that provides an additional means to install software and configure the switch module
- Scalability and performance
  - Autosensing 10/1000/1000 Mbps external Ethernet ports for bandwidth optimization
  - Non-blocking architecture with wire-speed forwarding of traffic
  - Media access control (MAC) address learning: automatic update, supports up to 16 K MAC addresses
  - Static and LACP (IEEE 802.3ad) link aggregation, up to 30 Gb of total bandwidth per switch, up to 11 trunk groups, up to four ports per group
  - Support for jumbo frames (up to 9216 bytes)
  - Broadcast/multicast storm control
  - Up to 128 IP interfaces

  - IGMP snooping for limit flooding of IP multicast traffic
  - IGMP filtering to control multicast traffic for hosts participating in multicast groups
  - Configurable traffic distribution schemes over trunk links based on source/destination IP or MAC addresses or both
  - Fast port forwarding and fast uplink convergence for rapid STP convergence
- Availability and redundancy
  - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy
  - IEEE 802.1D Spanning Tree Protocol (STP) for providing L2 redundancy
  - IEEE 802.1s Multiple STP (MSTP) for topology optimization, up to 128 STP instances are supported by single switch
  - IEEE 802.1w Rapid STP (RSTP) provides rapid STP convergence for critical delay-sensitive traffic like voice or video
  - Layer 2 Trunk Failover to support active/standby configurations of network adapter teaming on blades
  - Interchassis redundancy (L2 and L3)
- VLAN support
  - Up to 1024 VLANs are supported per switch, VLAN numbers range is from 1 to 4095 (4095 is used for the management module's connection only)
  - 802.1Q VLAN tagging support on all ports
  - Private VLANs
- Security
  - VLAN-based, MAC-based, and IP-based Access Control Lists (ACLs)
  - 802.1X port-based authentication
  - Multiple user IDs and passwords
  - User access control
  - Radius/TACACS+
- Quality of Service (QoS)
  - Support for IEEE 802.1p, IP ToS/DSCP, and ACL-based (MAC/IP source and destination addresses, VLANs) traffic classification, and processing
  - Traffic shaping and re-marking based on defined policies
  - Eight Weighted Round Robin (WRR) priority queues per port for processing qualified traffic
- Layer 3 functions
  - IP forwarding
  - IP filtering with ACLs, up to 896 filters
  - VRRP for router redundancy
  - Support for static routes, up to 128 entries in routing table
  - Routing protocol support (RIP v1, RIP v2, OSPF v2, BGP-4), up to 2048 entries in routing table
  - Support for DHCP Relay

- Manageability
  - SNMP (V1 and V3)
  - HTTP browser graphical user interface (GUI)
  - Telnet interface for CLI
  - SSH
  - Serial interface for CLI
  - Scriptable command-line interface (CLI)
  - Firmware image update (TFTP and FTP)
  - Network Time Protocol (NTP) for switch clock synchronization
  - BNT BLADEHarmony Manager support
- Monitoring
  - Switch LEDs for external port status and switch module status indication
  - Port mirroring for analyzing network traffic passing through the switch
  - Change tracking and remote logging with syslog feature
  - POST diagnostics
- Special functions
  - Serial over LAN (SOL)
- Network cables
  - 10GBASE-SR: 850 nm communication using multimode fiber cable (50 µ or 62.5 µ) up to 300 m
  - 10GBase-LR: 1310 nm communication using single mode fiber cable (10 µ) up to 10 km
  - 1000BASE-T:
    - UTP Category 6
    - UTP Category 5e (100 meters maximum)
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568B 100-ohm STP (100 meters maximum)
- RS-232 serial cable: 3-meter console cable DB-9-to-USB connector (nonstandard use of USB connector) that comes with the GbE switch module

The GbE switch module supports the following IEEE standards:

- IEEE 802.1D Spanning Tree Protocol (STP)
- IEEE 802.1s Multiple STP (MSTP)
- IEEE 802.1w Rapid STP (RSTP)
- IEEE 802.1p Tagged Packets
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1x port-based authentication
- IEEE 802.2 Logical Link Control
- IEEE 802.3ab 1000BASE-T Gigabit Ethernet
- IEEE 802.3z 1000BASE-X Gigabit Ethernet
- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3x Full-duplex Flow Control
- IEEE 802.3ae
  - 10GBASE-SR 10 Gb Ethernet fiber optics short range
  - 10GBASE-LR 10 Gb Ethernet fiber optics long range
- IEEE 802.3ak 10GBASE-CX4 10 Gb Ethernet copper

The Nortel 10G Uplink Switch has three 10 Gb ports as shown in Figure 2-15 on page 127. Two of them are CX4 copper ports and an IBM BladeCenter CX4 cable is required to connect to them. The third port is an XFP connection and uses either short range (SR) or long range (LR) XFP transceivers.

Table 2-15 lists the XFP and cable options.

Table 2-15 XFP and cable options for Nortel 10 Gb Uplink Ethernet Switch Module

Part number	Feature code	Description
32R1877	4360	Short Range Fiber XFP Transceiver
32R1878	4361	Long Range Fiber XFP Transceiver
32R1937	4362	IBM BladeCenter CX4 1.5m cable
32R1941	4363	IBM BladeCenter CX4 3.0m cable

Figure 2-16 shows the XFP Module and CX4 cable.

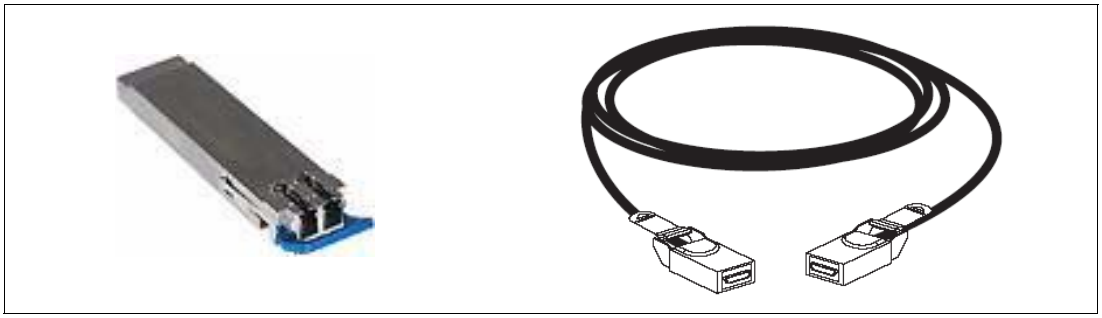

Figure 2-16 XFP Module and BladeCenter CX4 cable

Refer to the following documentation for more information:

- *Nortel 10 Gb Uplink Ethernet Switch Module Installation Guide*
  `http://www.ibm.com/support/docview.wss?uid=psg1MIGR-64158`
- *Nortel 10 Gb Uplink Ethernet Switch Module Application Guide*
  `http://www.ibm.com/support/docview.wss?uid=psg1MIGR-64687`

2.2.11 Nortel 1/10 Gb Uplink Ethernet Switch Module

The Nortel 1/10Gb Uplink Ethernet Switch Module for IBM BladeCenter, part number 44W4404, feature code 1590, is a switch option that enables administrators to offer full Layer 2 and 3 switching and routing capability with combined 1 Gb and 10 Gb uplinks in a BladeCenter chassis. Such consolidation simplifies the data center infrastructure and helps reduce the number of discrete devices, management consoles, and management systems. In addition, the next-generation switch module hardware is capable of supporting IPv6 Layer 3 frame forwarding protocols via a future firmware upgrade.

This ESM delivers port flexibility, efficient traffic management, increased uplink bandwidth, and strong Ethernet switching price/performance. This module is shown in Figure 2-17 on page 131.

Figure 2-17 Nortel 1/10 Gb Uplink Ethernet Switch Module

Standard features and functions include:

- Internal ports
  - 14 internal full-duplex Gigabit ports, one connected to each of the blade servers in the BladeCenter unit
  - Two internal full-duplex 10/100 Mbps ports connected to the management module
- External ports
  - Three slots for 10 Gb Ethernet SFP+ modules (support for 10GBASE-SR only at the time of writing). SFP+ modules are optional and must be ordered from IBM (10GbE 850 nm Fiber SFP+ Transceiver (SR), part number 44W4408, feature code 4942)
  - Six 10/100/1000 1000BASE-T Gigabit Ethernet ports with RJ-45 connectors
  - One RS-232 serial port that provides an additional means to install software and configure the switch module
- Scalability and performance
  - Fixed-speed external 10 Gb Ethernet ports for maximum uplink bandwidth
  - Autosensing 10/1000/1000 external Gigabit Ethernet ports for bandwidth optimization
  - Non-blocking architecture with wire-speed forwarding of traffic
  - Media access control (MAC) address learning: automatic update, supports up to 16 K MAC addresses
  - Up to 128 IP interfaces per switch
  - Static and LACP (IEEE 802.3ad) link aggregation, up to 36 Gb of total bandwidth per switch, up to 16 trunk groups, up to six ports per group
  - Support for jumbo frames (up to 9216 bytes)
  - Broadcast/multicast storm control

  - IGMP snooping for limit flooding of IP multicast traffic
  - IGMP filtering to control multicast traffic for hosts participating in multicast groups
  - Configurable traffic distribution schemes over trunk links based on source/destination IP or MAC addresses or both
  - Fast port forwarding and fast uplink convergence for rapid STP convergence
- Availability and redundancy
  - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy
  - IEEE 802.1D Spanning Tree Protocol (STP) for providing L2 redundancy
  - IEEE 802.1s Multiple STP (MSTP) for topology optimization, up to 128 STP instances are supported by single switch
  - IEEE 802.1w Rapid STP (RSTP) provides rapid STP convergence for critical delay-sensitive traffic like voice or video
  - Layer 2 Trunk Failover to support active/standby configurations of network adapter teaming on blades
  - Interchassis redundancy (L2 and L3)
- VLAN support
  - Up to 1024 VLANs are supported per switch, VLAN numbers range is from 1 to 4095 (4095 is used for management module's connection only)
  - 802.1Q VLAN tagging support on all ports
  - Private VLANs
- Security
  - VLAN-based, MAC-based, and IP-based Access Control Lists (ACLs)
  - 802.1X port-based authentication
  - Multiple user IDs and passwords
  - User access control
  - Radius/TACACS+
- Quality of Service (QoS)
  - Support for IEEE 802.1p, IP ToS/DSCP, and ACL-based (MAC/IP source and destination addresses, VLANs) traffic classification and processing
  - Traffic shaping and re-marking based on defined policies
  - Eight Weighted Round Robin (WRR) priority queues per port for processing qualified traffic
- Layer 3 functions
  - IP forwarding
  - IP filtering with ACLs, up to 896 ACLs are supported
  - VRRP for router redundancy
  - Support for up to 128 static routes
  - Routing protocol support (RIP v1, RIP v2, OSPF v2, BGP-4), up to 2048 entries in routing table
  - Support for DHCP Relay

- Manageability
  - SNMP (V1 and V3)
  - HTTP browser graphical user interface (GUI)
  - Telnet interface for CLI
  - SSH
  - Serial interface for CLI
  - Scriptable command line interface (CLI)
  - Firmware image update (TFTP and FTP)
  - Network Time Protocol (NTP) for switch clock synchronization
- Monitoring
  - Switch LEDs for external port status and switch module status indication
  - Port mirroring for analyzing network traffic passing through switch
  - Change tracking and remote logging with syslog feature
  - POST diagnostics
- Special functions
  - Serial over LAN (SOL)
- Network cables
  - 10GBASE-SR
    - 850 nm communication using multimode fiber cable (50µ or 62.5µ) up to 300 m
  - 10BASE-T:
    - UTP Category 3, 4, 5 (100 meters (328 feet) maximum)
    - 100-ohm STP (100 meters maximum)
  - 100BASE-TX:
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568 100-ohm STP (100 meters maximum)
  - 1000BASE-T:
    - UTP Category 6
    - UTP Category 5e (100 meters maximum)
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568B 100-ohm STP (100 meters maximum)
- RS-232 serial cable: 3-meter console cable DB-9-to-USB connector (nonstandard use of USB connector) that comes with the GbE switch module

The GbE switch module supports the following IEEE standards:

- IEEE 802.1D Spanning Tree Protocol (STP)
- IEEE 802.1s Multiple STP (MSTP)
- IEEE 802.1w Rapid STP (RSTP)
- IEEE 802.1p Tagged Packets
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1x port-based authentication
- IEEE 802.2 Logical Link Control
- IEEE 802.3 10BASE-T Ethernet
- IEEE 802.3u 100BASE-TX Fast Ethernet
- IEEE 802.3ab 1000BASE-T Gigabit Ethernet
- IEEE 802.3z 1000BASE-X Gigabit Ethernet

- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3x Full-duplex Flow Control
- IEEE 802.3ae
  - 10GBASE-SR

Refer to the following documentation for more information:

- *Nortel 1/10 Gb Uplink Ethernet Switch Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5076217
- *Nortel 1/10 Gb Uplink Ethernet Switch Module Application Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5076214
- *Nortel 1/10 Gb Uplink Ethernet Switch Module Command Reference*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5076525

2.2.12 Nortel 10 Gb Ethernet High Speed Switch Module

The Nortel 10 Gb Ethernet High Speed Switch Module, part number 39Y9267, feature code 2952, is a high-speed switch option that provides full high performance 10 Gigabit end-to-end connectivity from the blade server to the network in either a BladeCenter H or an HT chassis. The 10 Gb Ethernet High Speed Switch Module provides port flexibility, efficient traffic management, increased bandwidth, and exceptional 10 Gb Ethernet switching price/performance.

Figure 2-18 Nortel 10 Gb Ethernet High Speed Switch Module

Standard features and functions include:

- Internal ports
  - 14 internal 10 Gb ports, one connected to each of the blade servers in the BladeCenter unit
  - Two internal 100 Mbps ports connected to the management module
- External ports
  - Six external 10 Gb XFP ports (SR or LR) (XFP modules are optional.)
  - One 10/100/1000 Mbps Ethernet RJ-45 port for management use

  - One RS-232 serial port for management use
- Scalability and performance
  - Fixed-speed 10 Gbps external Ethernet ports for maximum network bandwidth
  - Non-blocking architecture with wire-speed forwarding of traffic
  - Media access control (MAC) address learning: automatic update, supports up to 16 K MAC addresses
  - Static and LACP (IEEE 802.3ad) link aggregation, up to 60 Gb of total bandwidth per switch, up to 13 trunk groups, up to six ports per group
  - Support for jumbo frames (up to 12288 bytes)
  - Broadcast/multicast storm control
  - Up to 250 IP interfaces
  - IGMP snooping for limit flooding of IP multicast traffic
  - IGMP filtering to control multicast traffic for hosts participating in multicast groups
  - Configurable traffic distribution schemes over trunk links based on source/destination IP or MAC addresses or both
  - Fast port forwarding and fast uplink convergence for rapid STP convergence
- Availability and redundancy
  - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy
  - IEEE 802.1D Spanning Tree Protocol (STP) for providing L2 redundancy
  - IEEE 802.1s Multiple STP (MSTP) for topology optimization, up to 128 STP instances are supported by single switch
  - IEEE 802.1w Rapid STP (RSTP) provides rapid STP convergence for critical delay-sensitive traffic like voice or video
  - Layer 2 Trunk Failover to support active/standby configurations of network adapter teaming on blades
  - Interchassis redundancy (L2 and L3)
- VLAN support
  - Up to 1024 VLANs are supported per switch, VLAN numbers range is from 1 to 4095 (4095 is used for management module's connection only)
  - 802.1Q VLAN tagging support on all ports
- Security
  - VLAN-based, MAC-based, and IP-based Access Control Lists (ACLs)
  - 802.1X port-based authentication
  - Multiple user IDs and passwords
  - User access control
  - Radius/TACACS+
- Quality of Service (QoS)
  - Support for IEEE 802.1p, IP ToS/DSCP, and ACL-based (MAC/IP source and destination addresses, VLANs) traffic classification and processing
  - Traffic shaping and re-marking based on defined policies
  - Eight Weighted Round Robin (WRR) priority queues per port for processing qualified traffic

- Layer 3 functions
  - IP forwarding
  - IP filtering with ACLs, up to 384 filters
  - VRRP for router redundancy
  - Support for up to 128 static routes
  - Routing protocol support (RIP v1, RIP v2, OSPF v2, BGP-4), up to 2048 entries in routing table
  - Support for DHCP Relay
- Manageability
  - SNMP (V1 and V3)
  - HTTP browser graphical user interface (GUI)
  - Telnet interface for CLI
  - SSH
  - Serial interface for CLI
  - Scriptable command line interface (CLI)
  - Firmware image update (TFTP and FTP)
  - Network Time Protocol (NTP) for switch clock synchronization
  - BNT BLADEHarmony Manager support
- Monitoring
  - Switch LEDs for external port status and switch module status indication
  - Port mirroring for analyzing network traffic passing through switch
  - Change tracking and remote logging with syslog feature
  - POST diagnostics
- Special functions
  - Serial over LAN (SOL)
- Network cables
  - 10GBASE-SR: 850 nm communication using multimode fiber cable (50 µ or 62.5 µ) up to 300 m
  - 10GBase-LR: 1310 nm communication using single mode fiber cable (10 µ) up to 10 km
  - 1000BASE-T:
    - UTP Category 6
    - UTP Category 5e (100 meters maximum)
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568B 100-ohm STP (100 meters maximum)
- RS-232 serial cable: 3-meter console cable DB-9-to-USB connector (nonstandard use of USB connector) that comes with the GbE switch module

The GbE switch module supports the following IEEE standards:

- IEEE 802.1D Spanning Tree Protocol (STP)
- IEEE 802.1s Multiple STP (MSTP)
- IEEE 802.1w Rapid STP (RSTP)
- IEEE 802.1p Tagged Packets

- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1x port-based authentication
- IEEE 802.2 Logical Link Control
- IEEE 802.3ab 1000BASE-T Gigabit Ethernet
- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3x Full-duplex Flow Control
- IEEE 802.3ae
  - 10GBASE-SR 10 Gb Ethernet fiber optics short range
  - 10GBASE-LR 10 Gb Ethernet fiber optics long range

The Nortel 10 Gb Switch has six external 10 Gb ports, as shown in Figure 2-18 on page 134. All of them are XFP ports and either use short range (SR) or long range (LR) XFP transceivers.

Table 2-16 lists the XFP and cable options.

Table 2-16 XFP and cable options for Nortel 10 Gb Ethernet High Speed Switch Module

Part number	Feature Code	Description
32R1877	4360	Short Range Fiber XFP Transceiver
32R1878	4361	Long Range Fiber XFP Transceiver

Figure 2-19 shows the XFP Module.

Figure 2-19 XFP Module

This switch uses the same XFPs as the Nortel 10 Gb Uplink Switch Module.

Refer to the following documentation for more information:

- *Installation Guide for Nortel 10 Gb Ethernet Switch Module - IBM BladeCenter*
  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5070408
- *Nortel 10 Gb Ethernet Switch Module Application Guide - IBM BladeCenter*
  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5070403

2.3 InfiniBand switch modules

InfiniBand is an advanced I/O technology that provides the following features:

- Scales-out data centers by interconnecting blades and chassis together with InfiniBand as the interconnect
- Takes advantage of high bandwidth and low latency characteristics of the InfiniBand standard with Remote Direct Memory Access (RDMA)

- Enables consolidation of LAN and SAN connectivity for an entire data cluster to a centralized location
- Virtualizes and shares I/O and storage across an entire BladeCenter or collection of BladeCenter chassis for cost savings and high availability

There are two type of InfiniBand modules: switch modules and bridge modules. Switch modules provide full native InfiniBand switching capability while bridges provide connectivity to traditional LANs or SANs using InfiniBand fabric.

There are several InfiniBand switch modules available for BladeCenter, as shown in Table 2-17.

Table 2-17 InfiniBand switch modules for IBM BladeCenter

Part Number	Feature Code	Description	Number of external ports	Type of external ports	Page
32R1756	1574	Cisco Systems 4X InfiniBand Switch Module	8	4X (10 Gbps)	138
39Y9207	2941	QLogic InfiniBand Ethernet Bridge Module	6	10/100/1000 Mbps	141
39Y9211	2942	QLogic InfiniBand Fibre Channel Bridge Module	6	1/2/4 Gbps	143

Note: The 4X InfiniBand Pass-thru Module is covered in 2.6.6, “4X InfiniBand Pass-thru Module” on page 182.

2.3.1 Cisco Systems 4X InfiniBand Switch Module

The Cisco 4X InfiniBand Switch Module, part number 32R1756, feature code 1574, adds InfiniBand (IB) switching capability to hosts in your IBM BladeCenter H or HT chassis. When you add one or two switch modules to your chassis and add HCA expansion cards to your blade servers, your servers can communicate to one another over IB within the chassis. When you connect the switch module to an outside IB fabric, servers can communicate with all nodes that connect to the IB network.

This switch delivers low-latency, high-bandwidth connectivity (up to 240 Gbps full duplex) between InfiniBand connected internal server blades, additional BladeCenter H chassis, stand-alone servers, and external gateways for connectivity to Ethernet LANs and Fibre Channel SANs.

Key features:

- Up to eight external 4X (10 Gbps) InfiniBand ports
- Sixteen internal 4X (10 Gbps) InfiniBand ports
- Up to 240 Gbps of aggregate bandwidth available per switch
- Dual switch configurations to provide additional bandwidth, and redundancy

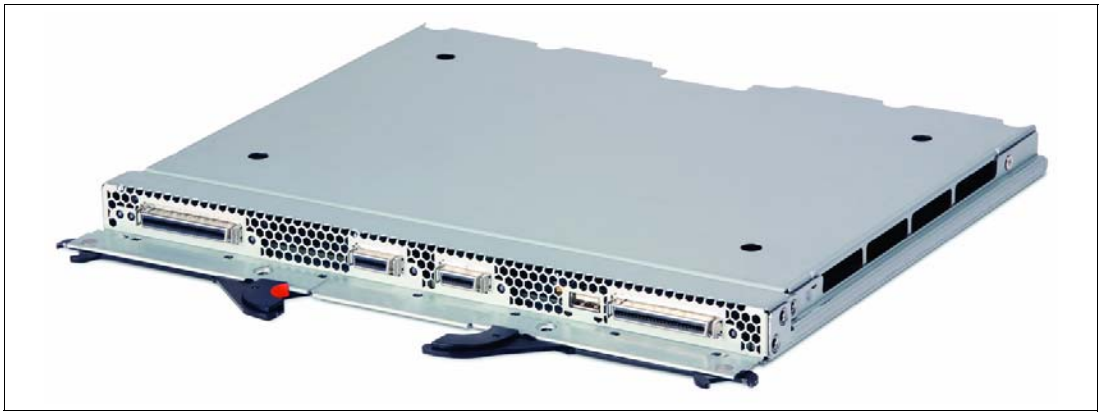

Figure 2-20 Cisco Systems 4X IB Switch Module

Each switch module includes 16 internal 4X ports to the backplane and eight 4X ports (in the form of two 4X connectors and two 12x connectors) on the front panel. The switch module provides fully non-blocking switching for all 24 ports. On the backplane, 14 of the 16 internal 4X ports provide 10 Gbps connections to the HCA expansion cards on server blades. The two remaining internal 4X ports provide connections to the chassis expansion modules. All external 4X connectors provide 10 Gbps connections to the outside IB network and can auto negotiate connection speed.

The Cisco 4X InfiniBand Switch Module transmit information to and from management modules over Ethernet (through an internal Ethernet switch) to facilitate setup and management. After you set up a switch module and bring it online, the on-board Cisco Subnet Manager brings distributed intelligence to the IB network.

Within the chassis, Cisco IB Switch Modules manage traffic to and from HCA expansion cards on the BladeCenter hosts. Each HCA expansion card adds two IB ports to a BladeCenter host. Each HCA port connects through the unit backplane to a particular switch module slot. The first IB port of each HCA card (ib0) connects to the Cisco IB Switch Module in I/O slot 7, and the second IB port of each HCA card (ib1) connects to I/O slot 9 (see 2.1.2, “BladeCenter H I/O topology” on page 62).

You can manage your IB switch module with any of the following interfaces:

- Simple Network Management Protocol (SNMP) versions 1, 2, and 3 with Cisco’s Management Information Base (MIBs)
- TopspinOS command line interface (CLI)
- Chassis Manager Web-based GUI
- Element Manager Java™-based GUI
- APIs (through SNMP)

Note: To implement the VFrame solution using the Cisco InfiniBand switch module, you need to purchase Cisco SFS switches and Cisco VFrame Server Fabric Virtualization software separately from Cisco Systems resellers. In addition, you currently cannot implement a VFrame environment if you are also using either of the QLogic Bridge Modules.

Table 2-18 lists Cisco 4X InfiniBand switch module-related options.

Table 2-18 Cisco 4X InfiniBand switch module options

Part number	Description
26R0813	InfiniBand 3 meter 4X Cable for IBM BladeCenter
26R0814	InfiniBand 4 meter 12X Cable for IBM BladeCenter
26R0816	InfiniBand 3 meter 12X to (3) 4X Cable for IBM BladeCenter
26R0847	InfiniBand 8 meter 4X Cable for IBM BladeCenter
26R0849	InfiniBand 8 meter 12X to (3) 4X Cable for IBM BladeCenter
43W4377	InfiniBand 10 meter 4X Cable for IBM BladeCenter
43W4378	InfiniBand 10 meter 12X Cable for IBM BladeCenter
43W4379	InfiniBand 11 meter 12X to (3) 4X Cable for IBM BladeCenter

For more information refer to *Cisco 4X InfiniBand Switch Module for IBM BladeCenter User Guide:*

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-65966

2.3.2 QLogic InfiniBand Ethernet Bridge Module

The QLogic InfiniBand Ethernet Bridge Module, part number 39Y9207, feature code 2941, is used in conjunction with the Cisco 4X InfiniBand Switch Module to provide seamless integration of BladeCenter's InfiniBand fabric into traditional LANs. It uses a concept of Virtual NICs to virtualize InfiniBand fabric to provide the blade server's access to the resources residing in a traditional network.

Figure 2-21 QLogic InfiniBand Ethernet Bridge Module

The Ethernet bridge module has the following features and specifications:

- Support for two internal InfiniBand double-data rate (DDR)-capable 4X (20 Gbps) links to the high-speed switch modules (HSSMs)
- Support for various types of transport data traffic:
  - Reliable
  - Unreliable
  - Connected
  - Unconnected
- Six external autosensing 10/100/1000 Mbps RJ-45 Ethernet (copper) ports
- 802.3ad link aggregation
- Support for jumbo frames
- Support of IEEE 802.1Q VLAN tagging
- Support for up to 1,150 Virtual NIC ports per module
- Automatic port and module failover
- TCP/UDP and IP header checksum offload and host checking

- 802.1p Priority Queuing/Scheduling
- Switched internal I2C Interface to the management modules
- Two internal 100 Mbps full-duplex Ethernet links to the management modules
- Power-on diagnostics and status reporting
- Support of simple network management protocol (SNMP) management information bases (MIBs) and traps through the Ethernet management ports

The Ethernet Bridge Module supports the following management methods:

- Command-line interface (CLI) via Telnet or SSH
- Web-based access with QLogic's Chassis Viewer
- SNMP/MIB

Refer to the following documentation for more information:

- *QLogic InfiniBand Ethernet Bridge Module and Fibre Channel Bridge Module Installation Guide - IBM BladeCenter H (Type 8852)*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5070108

- *InfiniBand Fibre Channel Bridge Module for IBM BladeCenter & InfiniBand Ethernet Bridge Module for IBM BladeCenter - Support*

  http://support.qlogic.com/support/oem_detail_all.asp?oemid=377

2.3.3 QLogic InfiniBand Fibre Channel Bridge Module

The QLogic InfiniBand Fibre Channel Bridge Module, part number 39Y9211, feature code 2942, is used in conjunction with Cisco 4X InfiniBand Switch Module to provide seamless integration of BladeCenter's InfiniBand fabric into traditional SAN infrastructure. It uses a concept of Virtual HBAs to virtualize the InfiniBand fabric to provide the blade server's access to the resources in traditional SANs.

Figure 2-22 QLogic InfiniBand Fibre Channel Bridge Module

The bridge module has the following features and specifications:

- Support for two internal InfiniBand double-data rate (DDR)-capable 4X (20 Gbps) links to the high-speed switch modules (HSSMs)
- Support for various types of transport data traffic:
  - Reliable
  - Unreliable
  - Connected
  - Unconnected
- Six external autosensing Fibre Channel 1/2/4 Gbps ports using standard small form-factor pluggable (SFP) connectors (SFPs not included)
- All common Fibre Channel topologies supported (direct attachment, switch, arbitrated loop)
- Supports up to 128 virtual host bus adapter (HBA) ports per module
- Automatic port and module failover
- Load balancing and port aggregation

- Logical unit number (LUN) mapping and masking
- SCSI-SRP, SCSI-FCP, and FC-PH-3 compliant
- Switched internal I2C Interface to the management modules
- Two internal 100 Mbps full-duplex Ethernet links to the management modules
- Power-on diagnostics and status reporting
- Support of simple network management protocol (SNMP) management information bases (MIBs) and traps through the Ethernet management ports

The bridge module supports the following management methods:

- Command-line interface (CLI) via Telnet or SSH
- Web-based access with QLogic's Chassis Viewer
- SNMP/MIB

This module ships standard without SFPs; they should be ordered additionally (see Table 2-19).

Table 2-19 Supported SFPs for QLogic InfiniBand Fibre Channel Bridge Module

Part number	Description
22R4897	4 Gbps SW SFP Transceiver 4 Pack
22R4902	4 Gbps SW SFP Transceiver
19K1271	2 Gb Fibre Channel Short Wave SFP Module
22R0483	2 Gb Fibre Channel Short Wave SFP 4 pack
19K1272	2 Gb Fibre Channel Long Wave SFP Module
22R0484	2 Gb Fibre Channel Long Wave SFP 4 pack

Refer to the following documentation for more information:

- *QLogic InfiniBand Ethernet Bridge Module and Fibre Channel Bridge Module Installation Guide - IBM BladeCenter H (Type 8852)*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5070108

- *InfiniBand Fibre Channel Bridge Module for IBM BladeCenter & InfiniBand Ethernet Bridge Module for IBM BladeCenter - Support*

  http://support.qlogic.com/support/oem_detail_all.asp?oemid=377

2.4 Fibre Channel switch modules

This section discusses different SAN storage options available for IBM BladeCenter. For information related to storage infrastructure planning and integration, see Chapter 5, "Storage integration" on page 395.

Note: Many software features and functions supported by the switch module depend on the firmware release that is loaded onto it. For example, old firmware does not support internal port reconfiguration (activation/deactivation) for 10-port switch versions, although the newer version does support it.

Refer to the Release Notes document for information about enhancements implemented in a specific firmware version. This document can be found on the firmware upgrade download page. Information covered in this section is based on the most current firmware releases available at the time of writing.

Table 2-20 lists the current Fibre Channel switch modules.

Table 2-20 Fibre Channel switch modules

Part number	Feature Code	Description	Number of external ports	Port interface bandwidth	Page
32R1812	1569	Brocade 20-port SAN Switch Module	6	4 Gbps	146
32R1813	1571	Brocade 10-port SAN Switch Module[a]	6	4 Gbps	146
39Y9280	2983	Cisco Systems 20-port 4 Gb Fibre Channel Switch Module	6	4 Gbps	149
39Y9284	2984	Cisco System s 10-port 4 Gb Fibre Channel Switch Module[a]	6	4 Gbps	149
26R0881	1560	QLogic 20-port 4 Gb Fibre Channel Switch Module	6	4 Gbps	151
43W6725	2987	QLogic 20-port 4 Gb SAN Switch Module	6	4 Gbps	153
43W6724	2986	QLogic 10-port 4 Gb SAN Switch Module[a]	6	4 Gbps	153
43W6723	2985	QLogic 4 Gb Intelligent Pass-thru Module	6	4 Gbps	156
44X1905	5478	QLogic 20-Port 8 Gb SAN Switch Module	6	8 Gbps	158
44X1907	5482	QLogic 8 Gb Intelligent Pass-thru Module	6	8 Gbps	161

a. Only 10 ports are activated on these switches. An optional upgrade to 20 ports (14 internal + 6 external) is available.

2.4.1 Brocade 20-port and 10-port SAN Switch Modules

The Brocade 20-port SAN Switch Module, part number 32R1812, feature code 1569, and Brocade 10-port SAN Switch Module, part number 32R1813, feature code1571, provide the integration of the Brocade fabric into the BladeCenter architecture. These SAN fabric switches deliver the advanced Brocade functions, performance, manageability, scalability and security required by the most demanding storage area networks (SANs). Integration of SAN switching capabilities reduces complexity and increases SAN manageability. The functions of these modules are equivalent to the Brocade SilkWorm SW200e switch. These switches are shown in Figure 2-23.

Note: When using the 4 Gb switch module with the BladeCenter T or HT chassis, the internal ports operate at 2 Gbps and the external ports operate at 4 Gbps.

Figure 2-23 Brocade 10-port and 20-port 4 Gb Fibre Channel Switch Modules for IBM BladeCenter

These two switches have almost the same features and functions except for the number of activated ports: the 10-port switch has ten activated ports, and the 20-port switch has 20 activated ports.

Prior to release 5.0.5a of the Brocade Fabric operating system, only the first seven internal ports and first three external ports were active on a 10-port switch, and this could not be changed. Starting from Release 5.0.5a, the Dynamic Ports on Demand (DPOD) licensing feature was introduced for 10-port Brocade switch. DPOD automatically enables switch ports as they are connected to either external devices or blade servers, and up to ten ports (any combination of existing internal and external physical ports) can be active at the same time. DPODs do this by detecting active links and assigning licenses to these ports. It is possible to exclude ports from DPOD process by permanently disabling them.

These switches have the following features:

- Six external autosensing Fibre Channel ports that operate at a maximum of 4 Gbps (gigabits per second)
- 14 internal fixed-speed Fibre Channel ports that operate at a maximum of 4 Gbps
- Two internal full-duplex 100 Mbps Ethernet interfaces, terminated at a single MAC
- Power-on diagnostics and status reporting
- Standard Ethernet and Fibre Channel management information base (MIB) support
- External ports that can be configured as F_ports (fabric ports), FL_ports (fabric loop ports), or E_ports (expansion ports)
- Internal ports that are configured as F_ports at 2 Gbps or 4 Gbps
- Frame-filtering technology that enables advanced zoning and monitoring capabilities
- Support up to 239 switches depending on configuration
- Optional small form-factor pluggable (SFP) modules

The following software features come with the switch module:

- Brocade Advanced Web Tools
- Brocade Advanced Zoning
- 20-port licensing (20-port version only)
- 10-port licensing (10-port version only)
- Support for Brocade Fabric Access API and API Scripting Toolkit

The switch supports the following fabric management (all management connections go through the management module):

- Management methods:
  - Web interface through Advanced Web Tools
  - Command-line interface (CLI) through Telnet program
  - Application program interface (API)
- Brocade Fabric Manager Advanced Feature Options - provides a Java-based application program to help manage the storage area network (SAN) fabric
- Management information base (MIB)
- Switch simple network management protocol (SNMP) agent - enables a network management workstation to receive configuration values, traffic information, and Fibre Channel failure data through SNMP and the Ethernet interface

Note: The switch comes without SFP modules, and they have to be ordered. Table 2-21 lists the part numbers that are supported.

Table 2-21 Supported SFPs for Brocade 4 Gb FC Switch Module

Part number	Description
22R4897	4 Gbps SW SFP Transceiver 4 Pack
22R4902	4 Gbps SW SFP Transceiver
19K1271	2 Gb Fibre Channel Short Wave SFP Module
22R0483	2 Gb Fibre Channel Short Wave SFP 4 pack
19K1272	2 Gb Fibre Channel Long Wave SFP Module
22R0484	2 Gb Fibre Channel Long Wave SFP 4 pack

Table 2-22 lists optional features available for these switches.

Table 2-22 Optional software features for Brocade 4 Gb FC Switch Module

Part Number	Description
26K5607	Brocade ISL Trunking for IBM BladeCenter
26K5612	Brocade Advanced Performance Monitoring for IBM BladeCenter
26K5613	Brocade Extended Fabrics for IBM BladeCenter
32R1919	Brocade Fabric Manager v5.x for IBM BladeCenter
26K5617	Brocade Advanced Security (Secure OS) for IBM BladeCenter
32R1855	Brocade 4 Gb Fabric Watch for IBM BladeCenter
32R1822	Brocade 4 Gb 10-Port Upgrade for IBM BladeCenter[a]

a. Available for 10-port switch only.

For additional information, refer to *Brocade 4 Gb SAN Switch Module Installation Guide:*

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-62707

2.4.2 Cisco Systems 4 Gb 20-port and 10-port Fibre Channel Switch Modules

Cisco Systems 4 Gb 20-port Fibre Channel Switch Module, part number 39Y9280, and Cisco Systems 4 Gb 10-port Fibre Channel switch Module, part number 39Y9284, provide high-performance end-to-end SAN solutions using 4 Gb Fibre Channel technology. These modules allow seamless integration of IBM BladeCenter into existing Cisco SANs, and have functions equivalent to the Cisco MDS 9124 switch. These switch module are shown in Figure 2-24.

Figure 2-24 Cisco Systems 4 Gb 20-port and 10-port Fibre Channel Switch Modules

These two switches have almost the same features and functions except for the number of activated ports: the 10-port switch has ten activated ports, and the 20-port switch has 20 activated ports.

Note: When using the 4 Gb switch module with the BladeCenter T or HT units, the internal ports operate at 2 Gbps and the external ports operate at 4 Gbps.

As of Cisco MDS SAN-OS Release 3.1 (1), On-Demand Port Activation (ODPA) licensing is used for the 10-port Cisco switch. ODPA automatically enables switch ports as they are connected to either blade servers (any seven of 14 internal ports can be enabled), or external devices (first three ports can be enabled consequently), and up to ten ports can be active at the same time. It is possible to make ports ineligible to get the license, or move licenses among ports.

These modules have the following features and functions:

- Six external autosensing Fibre Channel ports that operate at 4, 2, or 1 Gbps (Gigabits per second)
- 14 internal autosensing Fibre Channel ports that operate at 4 or 2 Gbps
- Two internal full-duplex 100 Mbps Ethernet interfaces
- Power-on diagnostics and status reporting
- External ports can be configured as F_ports (fabric ports), FL_ports (fabric loop ports), or E_ports (expansion ports)
- Internal ports are configured as F_ports at 2 Gbps or 4 Gbps
- N-port ID Virtualization (NPIV) support
- VSAN support (up to 16 VSANs per switch)
- Advanced Traffic Management capabilities include:
  - Virtual output queuing
  - PortChannels support up to six external physical ISL links aggregated into one logical bundle
  - Fabric-Shortest-Path-First (FSPF)-based multipathing for load balancing up to 16 equal-cost paths
  - QoS for bandwidth management and traffic prioritization
- Access Control Lists
- Nondisruptive software upgrade
- Ability to support up to 239 domain IDs depending on configuration
- Optional small form-factor pluggable (SFP) modules

The following software features come with the switch modules:

- Cisco Fabric Manager
- 20-port licensing (20-port version only)
- 10-port licensing (10-port version only)
- SMI-S and SNMP-based API

The switches support the following fabric management (all management connections go through the management module): Management methods:

- Web interface through Cisco Fabric Manager
- Cisco MDS9000 family command-line interface (CLI) through Telnet program
- Application program interface (API)

Note: The switch comes without SFP modules, and they have to be ordered. Table 2-23 on page 150 lists the part numbers that are supported.

Table 2-23 Supported SFPs for Cisco 4 Gb FC Switch Modules

Part number	Description
41Y8596	Cisco Systems 4 Gb Short-wave Length 4-pack SFP Module
41Y8598	Cisco Systems 4 Gb Short-wave Length SFP Module
41Y8600	Cisco Systems 4 Gb Long-wave Length SFP Module

Table 2-24 lists optional features available for these switches.

Table 2-24 Optional software features for Cisco 4 Gb FC Switch Module

Part number	Description
39Y9290	Cisco Systems 10 port License Upgrade[a]

a. Available for 10-port switch only.

For additional information refer to *Cisco Fibre Channel Switch Module set-up information - IBM BladeCenter:*

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5070625

2.4.3 QLogic 20-port 4 Gb Fibre Channel Switch Module

The QLogic 4 Gb 20-port FC switch, part number 26R0881, enables high-performing end-to-end SAN solutions using 4 Gb FC technology. These standards-based switches are affordable 4 Gb Fibre Channel Switches providing interoperability with FC-SW-2-compliant SANs. SAN fabric switch modules deliver the advanced functions, performance, manageability, and scalability required by the most demanding SANs. This switch has equivalent function to the QLogic SANbox® 5600. The QLogic switch module is shown in Figure 2-25.

Figure 2-25 QLogic 20-port 4 Gb Fibre Channel Switch Module

Note: When using the 4 Gb switch module with the BladeCenter T or HT units, the internal ports operate at 2 Gbps and the external ports operate at 4 Gbps.

The switch module has the following features:

- Six external autosensing Fibre Channel ports that operate at a maximum of 4 Gbps
- 14 internal fixed-speed Fibre Channel ports that operate at a maximum of 4 Gbps
- Two internal full-duplex 100 Mbps Ethernet interfaces
- External ports that can be configured as F_ports (fabric ports), FL_ports (fabric loop ports), or E_ports (expansion ports)
- Internal ports that are configured as F_ports at 2 Gbps or 4 Gbps
- Power-on diagnostics and status reporting
- Fabric security for controlling management access
- Support for Non-Disruptive Code Load Activation (NDCLA)
- Simple name server implementation
- Registered State Change Notification (RSCN)
- Support for standards-based FC-SW2 interoperability
- Error detection
- Frame bundling
- Configurable Fabric Address Notification (FAN)
- Support up to 239 switches depending on configuration
- Optional small form-factor pluggable (SFP) modules

The following software features come with the switch module:

- SANBrowser Web interface
- 20-port licensing (20-port version only)

The switch supports the following fabric management (all management connections go through the management module):

- Management methods:
  - Web interface through SANBrowser
  - Command line interface (CLI) through Telnet program
- SANsurfer® application
- Switch simple network management protocol (SNMP) agent - enables a network management workstation to receive configuration values, traffic information, and Fibre Channel failure data through SNMP and the Ethernet interface

Note: The switch comes without SFP modules, and they have to be ordered. Table 2-25 lists the part numbers that are supported.

Table 2-25 Supported SFPs for QLogic 4 Gb FC Switch Module

Part Number	Description
22R4897	4 Gbps SW SFP Transceiver 4 Pack
22R4902	4 Gbps SW SFP Transceiver
19K1271	2 Gb Fibre Channel Short Wave SFP Module
22R0483	2 Gb Fibre Channel Short Wave SFP 4 pack

Part Number	Description
19K1272	2 Gb Fibre Channel Long Wave SFP Module
22R0484	2 Gb Fibre Channel Long Wave SFP 4 pack

Table 2-26 lists optional features available for these switches.

Table 2-26 Optional software features for QLogic 4 Gb FC Switch Module

Part Number	Description
32R1795	McDATA Mode Firmware Upgrade for IBM BladeCenter QLogic Switch Modules
32R1797	McDATA SANtegrity Activation for IBM BladeCenter
32R1912	QLogic/McDATA 10-Port Upgrade for IBM BladeCenter[a]

a. Available for 10-port switch only.

For more information refer to *QLogic 4 Gb Fibre Channel Switch Module Installation Guide*:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58197

2.4.4 QLogic 20-port and 10-port 4 Gb SAN Switch Modules

QLogic 20-port 4 Gb FC SAN Switch Module, part number 43W6725, and QLogic 10-port 4 Gb SAN Switch Module, part number 43W6724, enable high performance 4 Gb SAN solutions, and allow BladeCenter integration with any open standard-based storage area networks utilizing full fabric switching functionality. These switches are shown in Figure 2-26.

Figure 2-26 QLogic 20-port and 10-port 4 Gb FC SAN Switch Modules

These two switches have almost the same features and functions except for the number of activated ports: the 10-port switch has ten activated ports, and the 20-port switch has 20 activated ports.

Notes:

- When using the 4 Gb switch module with the BladeCenter T unit, the internal ports operate at 2 Gbps and the external ports operate at 4 Gbps.
- When using the 4 Gb switch module with the BladeCenter HT unit in standard I/O bays, the internal ports operate at 2 Gbps and the external ports operate at 4 Gbps.
- When using the 4 Gb switch module with the BladeCenter HT unit in MSIM-HT, the internal ports operate at 4 Gbps and the external ports operate at 4 Gbps.

For the QLogic 10-Port 4 Gb SAN Switch Module, the default port license assignments are 1 through 7 for internal ports (the first seven blade bays), and 0, 15, and 16 for external ports (the first three external ports). Port license assignments can be changed to other available internal ports (8 through 14) or external ports (17, 18, or 19) by using the QLogic management tools.

The switch module has the following features:

- Six external autosensing (1, 2, or 4 Gbps) Fibre Channel ports that operate at a maximum of 4 Gbps
- 14 internal Fibre Channel ports that operate at a maximum of 4 Gbps
- Two internal full-duplex 100 Mbps Ethernet interfaces
- External ports that can be configured as F_ports (fabric ports), FL_ports (fabric loop ports), or E_ports (expansion ports)
- Internal ports that are configured as F_ports at 2 Gbps or 4 Gbps
- Power-on diagnostics and status reporting
- Fabric security for controlling management access (optional feature)
- Support for Non-Disruptive Code Load Activation (NDCLA)
- Registered State Change Notification (RSCN)
- Support for standards-based FC-SW2 interoperability
- Error detection
- Frame bundling
- Configurable Fabric Address Notification (FAN)
- Support up to 239 switches depending on configuration
- Optional small form-factor pluggable (SFP) modules

The following software features come with the switch module:

- QuickTools™ Web interface
- 20-port licensing (20-port version only)
- 10-port licensing (10-port version only)

The switch supports the following fabric management (all management connections go through the management module):

- Management methods:
  - Web interface through QuickTools
  - Command-line interface (CLI) through Telnet program
- Enterprise Fabric Suite™ 2007 application
- Switch simple network management protocol (SNMP) agent: enables a network management workstation to receive configuration values, traffic information, and Fibre Channel failure data through SNMP and the Ethernet interface

Note: The switch comes without SFP modules, and they have to be ordered additionally. Table 2-27 lists the part numbers that are supported.

Table 2-27 Supported SFPs for QLogic 4 Gb SAN Switch Modules

Part number	Description
39R6475	4 Gbps SW SFP Transceiver
26K7941	4 Gbps shortwave SFP transceiver pair
23R1703	4 Gbps 4 km longwave SFP Transceiver
23R1704	4 Gbps 4 km longwave SFP Transceiver 4 pack
22R4897	4 Gbps SW SFP Transceiver 4 Pack
22R4902	4 Gbps SW SFP Transceiver
22R4242	4 Gbps SW SFP Transceiver pair
19K1271	2 Gb Fibre Channel Short Wave SFP Module
22R0483	2 Gb Fibre Channel Short Wave SFP 4 pack
19K1272	2 Gb Fibre Channel Long Wave SFP Module
22R0484	2 Gb Fibre Channel Long Wave SFP 4 pack

Table 2-28 lists optional features available for these switches.

Table 2-28 Optional software features for QLogic 4 Gb SAN Switch Modules

Part number	Description
43W6836	QLogic SAN Doctor for IBM BladeCenter
43W6837	QLogic Fabric Security for IBM BladeCenter
43W6838	QLogic Enterprise Fabric Suite 2007 for IBM BladeCenter
43W4413	QLogic Enterprise Fabric Suite Major Release Upgrade for IBM BladeCenter
43W4414	QLogic Fabric Manager Support Renewal for IBM BladeCenter
32R1912	QLogic/McDATA 10-Port Upgrade for IBM BladeCenter[a]

a. Available for 10-port switch only.

Refer to the following documentation for more information:

- *QLogic 4 Gb Intelligent Pass-thru Module and 4 Gb SAN Switch Modules Getting Started Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5071313

- *QLogic 4 Gb Intelligent Pass-thru Module and 4 Gb SAN Switch Module for IBM Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5071303

2.4.5 QLogic 4 Gb Intelligent Pass-thru Module

QLogic 4 Gb Intelligent Pass-thru Module for IBM BladeCenter, part number 43W6723, provides seamless integration of IBM BladeCenter solution into existing Fibre Channel storage networks using N_Port ID Virtualization (NPIV) technology. This module is shown in Figure 2-27.

Figure 2-27 QLogic 4 Gb Intelligent Pass-thru Module for IBM BladeCenter

Note: When using the 4 Gb pass-thru module with the BladeCenter T or HT unit, the internal ports operate at 2 Gbps and the external ports operate at 4 Gbps.

The module concentrates multiple blade servers into the external ports. The external ports connect to external Fibre Channel switches that support NPIV. The internal ports connect directly to blade servers through the enclosure backplanes.

The pass-thru module presents one or more blade servers per port to the fabric. The module expands the fabric, but unlike a full fabric switch, it does not count against the fabric domain.

Important: This module *must* be connected to the external Fibre Channel switch (or switches) that support NPIV.

Direct connections from the module's external ports to storage devices or hosts must be avoided.

The Intelligent Pass-thru Module has the following features:

- Six external autosensing (1, 2, or 4 Gbps) Fibre Channel ports (TF_ports - Transparent fabric ports) that operate at a maximum of 4 Gbps
- 14 internal autosensing (2 or 4 Gbps) Fibre Channel ports (TH_ports - Transparent host ports) that operate at a maximum of 4 Gbps
- Two internal full-duplex 100 Mbps Ethernet interfaces for management
- N_Port ID Virtualization (NPIV) functionality
- Power-on diagnostics and status reporting
- Support for Non-Disruptive Code Load Activation (NDCLA)
- Registered State Change Notification (RSCN)
- Support for standards-based FC-SW2 interoperability
- Error detection
- Frame bundling
- Configurable Fabric Address Notification (FAN)
- Optional small form-factor pluggable (SFP) modules

The following software feature come with the switch module - QuickTools Web interface.

The module supports the following fabric management (all management connections go through the management module):

- Management methods:
  - Web interface through QuickTools
  - Command-line interface (CLI) through Telnet program
- Switch simple network management protocol (SNMP) agent - enables a network management workstation to receive configuration values, traffic information, and Fibre Channel failure data through SNMP and the Ethernet interface

Note: The switch comes without SFP modules, and they have to be ordered additionally. Table 2-27 lists the part numbers that are supported.

Table 2-29 Supported SFPs for QLogic 4 Gb Intelligent Pass-thru Module

Part number	Description
39R6475	4 Gbps SW SFP Transceiver
26K7941	4 Gbps shortwave SFP Transceiver pair
23R1703	4 Gbps 4 km longwave SFP Transceiver
23R1704	4 Gbps 4 km longwave SFP Transceiver 4 pack
22R4897	4 Gbps SW SFP Transceiver 4 Pack
22R4902	4 Gbps SW SFP Transceiver

Part number	Description
22R4242	4 Gbps SW SFP Transceiver pair
19K1271	2 Gb Fibre Channel Short Wave SFP Module
22R0483	2 Gb Fibre Channel Short Wave SFP 4 pack
19K1272	2 Gb Fibre Channel Long Wave SFP Module
22R0484	2 Gb Fibre Channel Long Wave SFP 4 pack

This module can be upgraded to the 20-port full fabric SAN switch by using a license key (see Table 2-30).

Table 2-30 Optional software features for QLogic 4 Gb Intelligent Pass-thru Module

Part number	Description
43W6839	QLogic 20-port Full Fabric License for IBM BladeCenter

Refer to the following documentation for more information:

- *QLogic 4 Gb Intelligent Pass-thru Module and 4 Gb SAN Switch Modules Getting Started Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5071313

- *QLogic 4 Gb Intelligent Pass-thru Module and 4 Gb SAN Switch Module for IBM Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5071303

2.4.6 QLogic 20-port 8 Gb SAN Switch Module

The QLogic 20-port 8Gb SAN Switch Module, part number 44X1905, feature code 5478, is a high-speed addition to the BladeCenter switch portfolio which helps enable scalability in SAN size and complexity while maintaining ease of management. The QLogic 8 Gb SAN Switch Modules are full-fabric Fibre Channel Switches and are available in a 20-port active configuration, which enables high-performances up to 2, 4 and 8 Gb per second SAN solutions. The switch is shown in Figure 2-28 on page 159.

Figure 2-28 QLogic 20-port 8 Gb FC SAN Switch Module

Note: When this switch module is used with the BladeCenter T unit, the internal ports operate at 2 Gbps and the external ports operate at speeds up to 8 Gbps. When this module is used with BladeCenter HT chassis in I/O bays 3 and 4, the internal ports operate at speeds up to 4 Gbps and the external ports operate at speeds up to 8 Gbps.

The switch module has the following features:

- Six external autosensing Fibre Channel ports (2, 4, or 8 Gbps for 8 Gb SFP+ transceivers, and 1, 2, or 4 Gbps for 4 Gb SFP transceivers) that operate at a maximum speed of 8 Gbps
- 14 internal Fibre Channel ports that operate as F_ports (fabric ports) at 2, 4, or 8 Gbps, depending on FC expansion cards installed into blades
- Two internal full-duplex 100 Mbps Ethernet interfaces
- External ports that can operate as F_ports (fabric ports), FL_ports (fabric loop ports), or E_ports (expansion ports)
- Power-on self-test diagnostics and status reporting
- Support for Non-Disruptive Code Load Activation (NDCLA)
- Registered State Change Notification (RSCN)
- Support for standards-based FC-SW2 interoperability
- Error detection
  - Cyclic redundancy check (CRC)
  - 8-byte and 10-byte conversion
  - Parity
  - Long frame and short frame
  - D_ID mismatch
  - S_ID mismatch

- Frame bundling
  - No frame bundling (frames are intermixed from different source ports)
  - Soft lockdown (the I/O module waits for the sequence to be completed or a gap in the frame traffic to occur before it services requests from a different port)
- Configurable Fabric Address Notification (FAN)
- Support up to 239 switches depending on configuration
- 8 Gb switch fabric aggregate bandwidth: 320 Gbps at full duplex
- Maximum frame size: 2148 bytes (2112 byte payload)
- Nonblocking architecture to prevent latency
- Support for Call Home function
- Support for Domain Name Service (DNS)
- Support for Internet protocol (IP) Version 6
- Support for Internet protocol security (IPsec)
- Support for separate trap community strings for each trap address
- Support for Simple Network Management Protocol (SNMP) Version 3
- Support for vital product data (VPD)
- Optional small form-factor pluggable plus (SFP+) modules

The following software features come with the switch module:

- QuickTools Web interface

The switch supports the following fabric management (all management connections go through the management module):

- Management methods:
  - Web interface through QuickTools
  - Command-line interface (CLI) through Telnet program
  - Enterprise Fabric Suite 2007 application
  - Switch simple network management protocol (SNMP) agent: enables a network management workstation to receive configuration values, traffic information, and Fibre Channel failure data through SNMP and the Ethernet interface

Note: The switch comes without SFP+ modules, and they have to be ordered additionally. Table 2-31 lists the part numbers that are supported.

Table 2-31 Supported SFPs and SFPs+ for QLogic 20-port 8 Gb SAN Switch Module

Part number	Feature Code	Description
44X1964	5075	IBM 8 Gb SFP+ SW Optical Transceiver
39R6475	4804	4 Gbps SW SFP Transceiver
26K7941	2410	4 Gbps shortwave SFP transceiver pair
22R4897	2414	4 Gbps SW SFP Transceiver 4 Pack
22R4902	2410	4 Gbps SW SFP Transceiver

Table 2-32 on page 161 lists optional features available for these switches.

Table 2-32 Optional software features for QLogic 8 Gb SAN Switch Modules

Part number	Feature Code	Description
43W6836	4874	QLogic SAN Doctor for IBM BladeCenter
43W6838	4876	QLogic Enterprise Fabric Suite 2007 for IBM BladeCenter
43W4413	4878	QLogic Enterprise Fabric Suite Major Release Upgrade
43W4414	4879	QLogic Fabric Manager Support Renewal for IBM BladeCenter

Refer to the following documentation for more information:

- *QLogic 8 Gb Intelligent Pass-thru Module and 20-port 8 Gb SAN Switch Module Installation and User's Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5079302

2.4.7 QLogic 8 Gb Intelligent Pass-thru Module

The QLogic 8 Gb Intelligent Pass-thru Module for IBM BladeCenter, part number 44X1907, feature code 5482, provides seamless integration of IBM BladeCenter solution into existing Fibre Channel storage networks using N_Port ID Virtualization (NPIV) technology. This module is shown in Figure 2-29.

Figure 2-29 QLogic 8 Gb Intelligent Pass-thru Module for IBM BladeCenter

Note: When this switch module is used with the BladeCenter T unit, the internal ports operate at 2 Gbps and the external ports operate at speeds up to 8 Gbps. When this module is used with BladeCenter HT chassis in I/O bays 3 and 4, the internal ports operate at speeds up to 4 Gbps and the external ports operate at speeds up to 8 Gbps.

The module concentrates multiple blade servers into the external ports. The external ports connect to external Fibre Channel switches that support NPIV. The internal ports connect directly to blade servers through the enclosure backplanes.

The pass-thru module presents one or more blade servers per port to the fabric. The module expands the fabric, but unlike a full fabric switch, it does not count against the fabric domain.

Note: This module *must* be connected to the external Fibre Channel switch (or switches) that support NPIV. Direct connections from the module's external ports to storage devices or hosts must be avoided.

The switch module has the following features:

- Six external autosensing Fibre Channel ports (2, 4, or 8 Gbps for 8 Gb SFP+ transceivers, and 1, 2, or 4 Gbps for 4 Gb SFP transceivers) that operate as TF_Ports (Transparent Fabric Ports) at a maximum speed of 8 Gbps
- 14 internal Fibre Channel ports that operate as TH_Ports (Transparent Host Ports) at 2, 4, or 8 Gbps depending on FC expansion cards installed into blades
- Two internal full-duplex 100 Mbps Ethernet interfaces
- N_Port ID Virtualization (NPIV) functionality
- Power-on self-test diagnostics and status reporting
- Support for Non-Disruptive Code Load Activation (NDCLA)
- Registered State Change Notification (RSCN)
- Support for standards-based FC-SW2 interoperability
- Error detection
  - Cyclic redundancy check (CRC)
  - 8-byte and 10-byte conversion
  - Parity
  - Long frame and short frame
  - D_ID mismatch
  - S_ID mismatch
- Frame bundling
  - No frame bundling (frames are intermixed from different source ports)
  - Soft lockdown (the I/O module waits for the sequence to be completed or a gap in the frame traffic to occur before it services requests from a different port)
- Configurable Fabric Address Notification (FAN)
- Support up to 239 switches depending on configuration
- 8 Gb switch fabric aggregate bandwidth: 320 Gbps at full duplex
- Maximum frame size: 2148 bytes (2112 byte payload)
- Nonblocking architecture to prevent latency
- Support for Call Home function
- Support for Domain Name Service (DNS)
- Support for Internet protocol (IP) Version 6
- Support for Internet protocol security (IPsec)
- Support for separate trap community strings for each trap address
- Support for Simple Network Management Protocol (SNMP) Version 3

- Support for vital product data (VPD)
- Optional small form-factor pluggable plus (SFP+) modules

The following software features come with the switch module:

- QuickTools Web interface

The switch supports the following fabric management (all management connections go through the management module):

- Management methods:
  - Web interface through QuickTools
  - Command-line interface (CLI) through Telnet program
  - Enterprise Fabric Suite 2007 application
  - Switch simple network management protocol (SNMP) agent: enables a network management workstation to receive configuration values, traffic information, and Fibre Channel failure data through SNMP and the Ethernet interface

Note: The switch comes without SFP+ modules, and they have to be ordered additionally. Table 2-31 lists the part numbers that are supported.

Table 2-33 Supported SFPs and SFPs+ for QLogic 8 Gb Intelligent Pass-thru Modules

Part number	Feature Code	Description
44X1964	5075	IBM 8 Gb SFP+ SW Optical Transceiver
39R6475	4804	4 Gbps SW SFP Transceiver
26K7941	2410	4 Gbps shortwave SFP transceiver pair
22R4897	2414	4 Gbps SW SFP Transceiver 4 Pack
22R4902	2410	4 Gbps SW SFP Transceiver

Table 2-32 on page 161 lists optional features available for these switches.

Table 2-34 Optional software features for QLogic 8 Gb Intelligent Pass-thru Modules

Part number	Feature Code	Description
43W6836	4874	QLogic SAN Doctor for IBM BladeCenter
43W6838	4876	QLogic Enterprise Fabric Suite 2007 for IBM BladeCenter
43W4413	4878	QLogic Enterprise Fabric Suite Major Release Upgrade
43W4414	4879	QLogic Fabric Manager Support Renewal for IBM BladeCenter

Refer to the following documentation for more information:

- *QLogic 8 Gb Intelligent Pass-thru Module and 20-port 8 Gb SAN Switch Module Installation and User's Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5079302

2.5 SAS I/O modules

SAS I/O modules provide affordable storage connectivity for BladeCenter chassis using SAS technology to create simple fabric for external shared or non-shared storage attachments, or to perform RAID controller functions inside the BladeCenter S chassis for HDDs installed into DSMs and external EXP3000 expansions. Table 2-35 lists the SAS I/O modules available.

Table 2-35 SAS I/O modules

Part number	Feature Code	Description	Number of external ports	External port interface bandwidth	Page
39Y9195	2980	SAS Connectivity Module	4	12 Gbps	164
43W3584	3734	SAS RAID Controller Module[a]	4	12 Gbps	166

a. Supported in BladeCenter S chassis only

2.5.1 SAS Connectivity Module

The SAS Connectivity Module, part number 39Y9195, feature code 2980, is used for affordable SAS-based storage attachments in IBM BladeCenter solutions.

At the time of writing only the following devices can be attached to this module:

- IBM System Storage DS3200 (both data storage and SAN boot are supported simultaneously)
- IBM BladeCenter Boot Disk System (SAN boot only, no data storage support)
- BladeCenter S integrated storage DSMs (by using an integrated RAID controller on the blade server itself, or by using a ServeRAID-MR10ie controller installed into supported blade servers)
- IBM System Storage EXP3000 (with BladeCenter S and ServeRAID-MR10ie RAID controllers)
- Selected tape devices

Refer to 5.9, “External SAS-based storage for IBM BladeCenter” on page 410, 5.10, “IBM BladeCenter S integrated storage” on page 423, and 5.11, “InfiniBand and Fibre Channel SANs” on page 432 for detailed compatibility information about supported combinations of chassis, blades, storage devices. and operating systems. The SAS Connectivity Module is shown in Figure 2-30 on page 165.

Figure 2-30 SAS Connectivity Module for IBM BladeCenter

The SAS Connectivity Module has the following characteristics:

- SAS expander
- Four external x4 links for storage servers with mini-SAS connector type (SFF-8088)
- Fourteen internal x1 links to blade servers
- Serial SCSI Protocol (SSP)
- Serial Management Protocol (SMP) as defined in the SAS specification
- Link error detection
- Power-on diagnostics and status reporting

The switch supports the following fabric management (all management connections go through the management module):

- Management methods:
  - Web interface
  - Advanced Management Module's web interface for integrated storage configuration
  - Command line interface (CLI) through Telnet program
  - Storage Configuration Manager application

- SAS connectivity module simple network management protocol (SNMP) agent - enables a network management workstation to receive configuration values and SAS link data through SNMP and the Ethernet interface.

Table 2-36 lists SAS cables that can be used with the SAS Connectivity Module. These cables are x4, and have mini-SAS (SFF-8088) male connectors on both sides.

Table 2-36 SAS cables supported by SAS Connectivity Module

Part number	Feature code	Description
39R6531	3707	IBM 3m SAS Cable
39R6529	3708	IBM 1m SAS Cable

Refer to the following documentation for additional information:

- *SAS Connectivity Module Installation and User Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5072374

2.5.2 SAS RAID Controller Module

The SAS RAID Controller Module for IBM BladeCenter S, part number 43W3584, feature code 3734, provides integrated RAID and SAS-based storage area networking capabilities for IBM BladeCenter S-based solutions.

At the time of writing only the following devices can be used with this module:

- BladeCenter S integrated storage DSMs
- SAS Expansion Card (39Y9190)

Refer to 5.10, “IBM BladeCenter S integrated storage” on page 423 for detailed compatibility information about supported combinations of chassis, blades, storage devices, and operating systems. The SAS RAID Controller Module is shown in Figure 2-31 on page 167.

Figure 2-31 SAS RAID Controller Module for IBM BladeCenter S

The SAS RAID Controller ships standard with a battery backup unit (BBU), which is shown in Figure 2-32.

Figure 2-32 Battery Backup Unit for SAS RAID Controller Module

The SAS RAID Controller Module consists of two subsystems:

- RAID Controller to perform RAID management (working with storage pools, volumes, RAID levels, and host mappings)
- SAS Switch to provide SAS connectivity for blade servers and disks.

The SAS RAID Controller Module has the following characteristics:

- Standard I/O module form-factor
- Four external x4 3 Gb SAS links with mini-SAS connector type (SFF-8088)
- Six internal x1 3 Gb SAS links to blade servers
- Two internal x4 3 Gb SAS links to two Disk Storage Modules (DSMs) in BladeCenter S chassis
- 1 Gb Ethernet interface for RAID management
- RAID levels supported: 0, 1, 5, 10
- Battery backup unit provides up to 72 hours of data storage in cache in case of complete power failure.
- 10/100 Mb Ethernet for I/O module management (monitoring status, updating firmware, IP addresses configuration etc)
- Serial SCSI Protocol (SSP)
- Serial Management Protocol (SMP) as defined in the SAS specification
- Link error detection
- Power-on diagnostics and status reporting

The module supports the following management tools and interfaces (all management connections go through the management module except RAID management, which goes through the Ethernet switch installed in I/O bay 1):

- Management methods:
  - IBM Start Now Advisor
  - Advanced Management Module's Web interface for initial module configuration and status monitoring
  - IBM Storage Configuration Manager application for GUI-based full-featured management
  - SAS RAID Controller Module Web interface to monitor status of module components
  - Command line interface (CLI) through Telnet program for CLI-based full-featured management
  - CLI-based firmware update tool
  - SAS connectivity module simple network management protocol (SNMP) agent - enables a network management workstation to receive configuration values and SAS link data through SNMP and the Ethernet interface.

Note: Each SAS RAID Controller Module requires two IP addresses: one for RAID management, and one for the management of I/O module itself.

Refer to the following documentation for additional information:

- *SAS RAID Controller Module Installation and User Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5078040

2.6 Pass-thru and interconnect modules

The following topics are discussed in this section:

Pass-thru modules are I/O options that can be used for direct connections of blades ports to the external infrastructure devices such as network of storage switches. The Pass-thru module is almost like a traditional network patch-panel, but it routes internal blade ports to the outside.

There are several types of standard pass-thru modules: Copper Pass-thru Module (CPM), Intelligent Copper Pass-thru Module (ICPM), and Optical Pass-thru Module (OPM). All are capable of carrying Ethernet signals, but only OPM can be used for Fibre Channel signals.

Standard OPM or CPM modules (used in non-high speed bays 1-4) use a unique design that links several internal blade ports to a single external port on the pass-thru module. Therefore, specific “hydra” or “octopus” pass-thru cables have to be used with pass-thru modules to provide standard external connections (one connector per each blade port).

ICPM uses a different design that provides 14 external RJ-45 connectors, one for every blade server. This means there is no need for specific octopus cables, and standard Ethernet patch cables can be used.

Additionally, IBM offers a high-speed pass-thru module (4X InfiniBand Pass-thru Module) which occupies two adjacent high-speed slots and is used to route InfiniBand signals from the 4X DDR InfiniBand blade’s high-speed expansion cards to the outside. Each port on this module represents a single connection from one port of an expansion card.

Note: The QLogic 4 Gb and 8 Gb Intelligent Pass-thru Modules are listed in Table 2-37, but it is more of a Fibre Channel switch module so it is discussed in 2.4, “Fibre Channel switch modules” on page 144.

Interconnect modules are represented by the Multi-Switch Interconnect Module (MSIM), which is used with the BladeCenter H and BladeCenter HT chassis. It provides a means for existing standard switch modules to communicate with blade servers using the high-speed fabric. This module is installed into two adjacent high-speed switch slots, which allows the possibility of increasing the number of standard ports per blade up to eight without density losses.

Table 2-37 lists available pass-thru and interconnect modules.

Table 2-37 Pass-thru and interconnect modules for IBM BladeCenter

Part Number	Feature code	Description	Page
43W6723	2985	QLogic 4 Gb Intelligent Pass-thru Module	156
44X1907	5482	QLogic 8 Gb Intelligent Pass-thru Module	161
39Y9320	2900	Copper Pass-thru Module	170
44W4483	5452	Intelligent Copper Pass-thru Module	171

Part Number	Feature code	Description	Page
39Y9316	1556	Optical Pass-thru Module	173
39Y9314	1465	Multi-switch Interconnect Module	174
44R5913	5491	Multi-switch Interconnect Module for BladeCenter HT	178
43W4419	2990	4X InfiniBand Pass-thru Module for IBM BladeCenter	182

2.6.1 IBM BladeCenter Copper Pass-thru Module

The IBM BladeCenter Copper Pass-thru Module (CPM), part number 39Y9320, feature code 2900, provides an unconfigured network connection that enables the blade servers in the BladeCenter unit to connect to an existing network infrastructure. No configuration of the Copper Pass-thru Module is required.

The CPM provides a single connection from each blade - one RJ-45 male connector goes directly from each of the 14 blades in the chassis for connection into an external switch or patch panel. The CPM comes with one cable, which provides external connections for one of three groups of adjacent blade servers in the chassis - blade bays 1-5, 6-10, or 11-14. If you have 14 blades in the chassis, then you will need two extra cables, which will give you 15 RJ-45 connectors (1 is left unused). If you install the CPM in bay 1 of the chassis, then it will route all Ethernet signals from controller 1 in each blade. Similarly, if you put a CPM in bay 2 it will route all Ethernet controller 2s from each blade server.

For supported I/O bays and technologies, refer to Figure 2-1 on page 60.

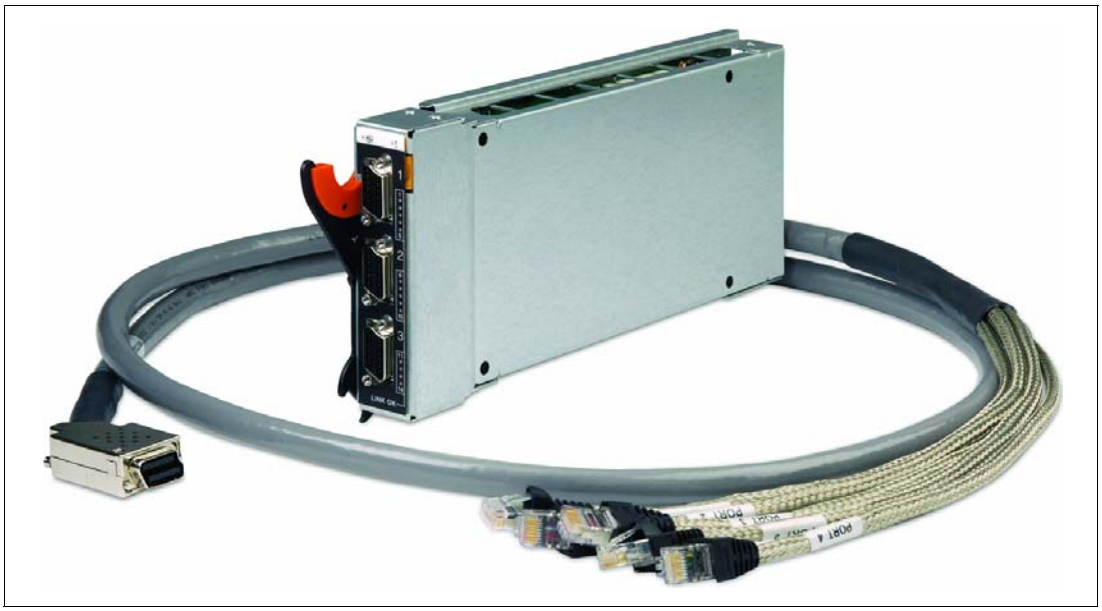

Figure 2-33 IBM BladeCenter Copper Pass-thru Module and cable

The features include:

- Ports
  - There are three external copper ports for making connections to the network infrastructure. Each port provides connections for up to five bi-directional copper channels. Each external port is hardwired to a specific set of blade bays in the chassis (from 1 to 5, from 6 to 10, and from 11 to 14), and this cannot be changed.

  - There are 14 internal bi-directional ports, connected through the Copper Pass-thru cables to each blade server.
- Status and error LEDs
  - Informational LEDs on the front panel indicate I/O module port status and errors.
- Cables:
  - Up to three Copper Pass-thru Module cables can be connected to the Copper Pass-thru Module (one Copper Pass-thru Module cable is provided).
  - The Copper Pass-thru Module cables are terminated with industry-standard bi-directional connectors.

Table 2-38 lists CPM-related options.

Table 2-38 CPM options

Part number	Description
39Y9320	IBM BladeCenter Copper Pass-thru Module (includes one cable)
39Y9170	IBM BladeCenter Copper Pass-thru Module Cable

The cable for the CPM has a multi-port copper connector on one end and a fanout to five RJ45 Ethernet connectors at the other. One Copper Pass-thru cable comes with the Copper Pass-thru Module, and you can purchase two additional cables separately if you require them.

The CPM cable length is 3.1 meters. If the client network environment requires you to extend its length, then use an Ethernet coupler or a patch panel.

Note: The CPM is a Gb Ethernet that only supports a 1000 Mbps connection and requires no configuration. Ensure that the external switch devices to which CPM connects provide 1000base-T ports.

The channel that is created is “neutral” to the protocol, enabling the deployment of external switch devices with direct access to blade servers.

Note: Serial over LAN is not supported on this module.

For more information about this topic, refer to *IBM BladeCenter Copper Pass-thru Module Installation Guide:*

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-59020

2.6.2 Intelligent Copper Pass-thru Module for IBM BladeCenter

The Intelligent Copper Pass-thru Module for IBM BladeCenter (ICPM), part number 44W4483, feature code 5452, provides an unswitched/unblocked network connection to each blade server that enables the blade servers in the BladeCenter unit to connect to an existing network infrastructure. No configuration of the ICPM is required.

The ICPM provides a single connection per blade for every blade’s internal Ethernet port. One RJ-45 female connector per blade is on the ICPM module itself. There are no special cables required for ICPM, because it uses standard Ethernet RJ-45 patch cords.

Note the following points:

- If you install the ICPM in bay 1 of the chassis, then it will route all Ethernet signals from controller 1 in each blade.
- Similarly, if you put the ICPM in bay 2, it will route all Ethernet controller 2s from each blade server.
- If you install additional Ethernet-capable expansion cards or MSIM or both, then each ICPM will route all signals from the Ethernet ports of all blades with the same port number.
- If you install ICPM in bay 1 of BladeCenter S, then both integrated Ethernet ports on blade will be routed to the same ICPM. ICPM is shown in Figure 2-34.

For supported I/O bays and technologies, refer to Table 2-1 on page 62.

Figure 2-34 Intelligent Copper Pass-thru Module for IBM BladeCenter

The features include:

- Ports
  - 14 external RJ-45 copper ports for making connections to the network infrastructure. The external ports can operate at 100 or 1000 Mbps.
  - 14 internal bi-directional ports, connected to each of the blade servers.
  - Two internal 100 Mb Ethernet links for management.
- Functions
  - One-to-one mapping of internal Ethernet ports on the blade to external copper RJ-45 ports with wire-speed processing.
  - If an external link goes down, then the corresponding internal link goes down as well. If an internal link goes down, then the corresponding external link goes down as well.

- Serial over LAN (SOL) and cKVM technologies are supported by this module.

- Monitoring
  - ICPM has informational LEDs on the front panel for indicating I/O module status. Each Ethernet port on ICPM has two LEDs to indicate link status and port activity.
  - The module supports self-test diagnostics.
- Management
  - This module is managed through the Advanced Management Module's management interface.
- Cables:
  - 100BASE-TX:
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568 100-ohm STP (100 meters maximum)
  - 1000BASE-T:
    - UTP Category 6
    - UTP Category 5e (100 meters maximum)
    - UTP Category 5 (100 meters maximum)
    - EIA/TIA-568B 100-ohm STP (100 meters maximum)

For more information, refer to *Intelligent Copper Pass-thru Module Installation Guide:*

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5076328

2.6.3 IBM BladeCenter Optical Pass-thru Module

The Optical Pass-thru Module, part number 39Y9316, provides an unconfigured network connection that enables the blade servers in the BladeCenter unit to connect to an existing network infrastructure. No configuration of the Optical Pass-thru Module is required. Refer to Figure 2-1 on page 60 for supported I/O configurations.

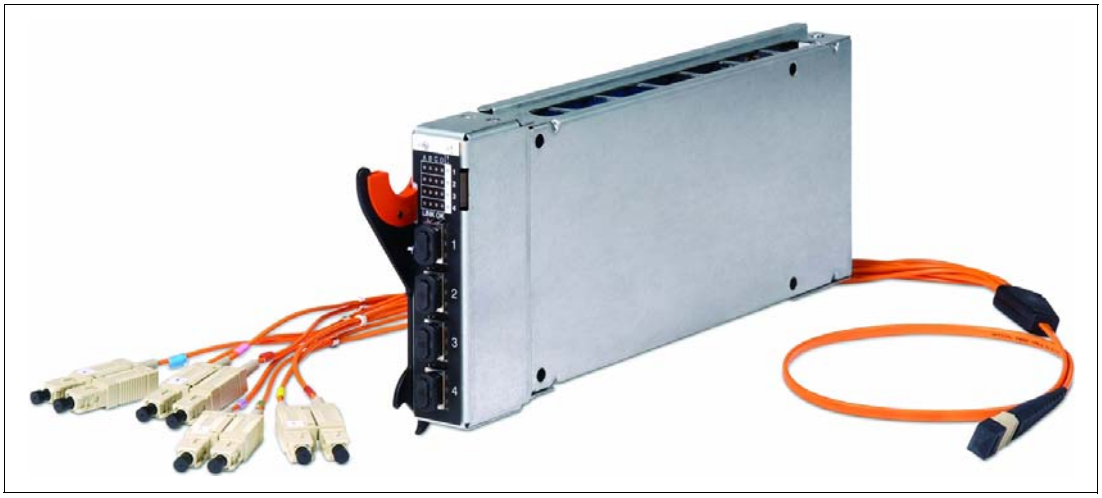

Figure 2-35 IBM BladeCenter Optical Pass-thru Module and SC cable

The IBM BladeCenter Optical Pass-thru Module features include:

- Ports
  - There are four external optical ports for making connections to the network infrastructure. Each port provides connections for up to four bi-directional optical

channels. Each external port is hardwired to the specific set of blade bays in chassis (from 1 to 4, from 5 to 8, from 9 to 12, and from 13 to 14), and this cannot be changed.

  - There are 14 internal full-duplex ports, one connected to each of the blade servers.

- Status and failure LEDs
  - Informational LEDs on the front panel indicate I/O module and port status and faults.
- Cables
  - Up to four Optical Pass-thru Module cables (not included) can be connected to the Optical Pass-thru Module.
  - The Optical Pass-thru Module cable is terminated with industry-standard duplex connectors.

Table 2-39 lists OPM-related options.

Table 2-39 OPM options

Option number	Feature code	Description
39Y9316	1556	IBM BladeCenter Optical Pass-thru Module (does *not* include any cables)
39Y9171	4234	IBM BladeCenter Optical Pass-thru Module SC Cable
39Y9172	4235	IBM BladeCenter Optical Pass-thru Module LC Cable

The Optical Pass-thru cable does not come with the OPM; you must purchase it as a separate item. The cable provides four connections between the Optical Pass-thru Module and an external Ethernet or Fibre Channel environment.

The Optical Pass-thru Module Cable is 1.5 meters in length, and optical couplers are supplied with the cables so that their length can be extended to 150 meters. The OPM cable uses 62.5 microns multimode fiber optics.

The Multi-Port Optical connector on the Optical Pass-thru cable is attached to the OPM fiber optic port. The other end of the cable has four SC or LC duplex optical connectors, which are attached to the networking device beyond the OPM.

Note: Serial over LAN is not supported by this module.

OPM supports speeds of up to 2 Gb Fibre Channel. If you plan to use external 4 Gb FC switches and internal 4 Gb FC expansion cards together with OPM, the FC fabric will operate at 2 Gb.

For more information, refer to *IBM BladeCenter Optical Pass-thru Module Installation Guide*:

`http://www.ibm.com/support/docview.wss?uid=psg1MIGR-52864`

2.6.4 Multi-switch Interconnect Module

The Multi-switch Interconnect Module (MSIM), part number 39Y9314, feature code 1465, is a switch module container that fits in the high speed switch bays (bays 7 and 8 and bays 9 and 10) of the BladeCenter H chassis. The MSIM accepts two supported standard switch modules and enables a blade server to have up to a total of eight network connections.

To be able to connect to the switch modules in the MSIM, a blade server must have a CFFh-type adapter installed. See 3.1, “Expansion cards” on page 234, for a description and list of CFFh expansion cards.

The MSIM is shown in Figure 2-36.

Note: The MSIM comes standard without any I/O modules installed. They need to be ordered separately. In addition, the use of MSIM modules requires that all four power modules be installed in the chassis.

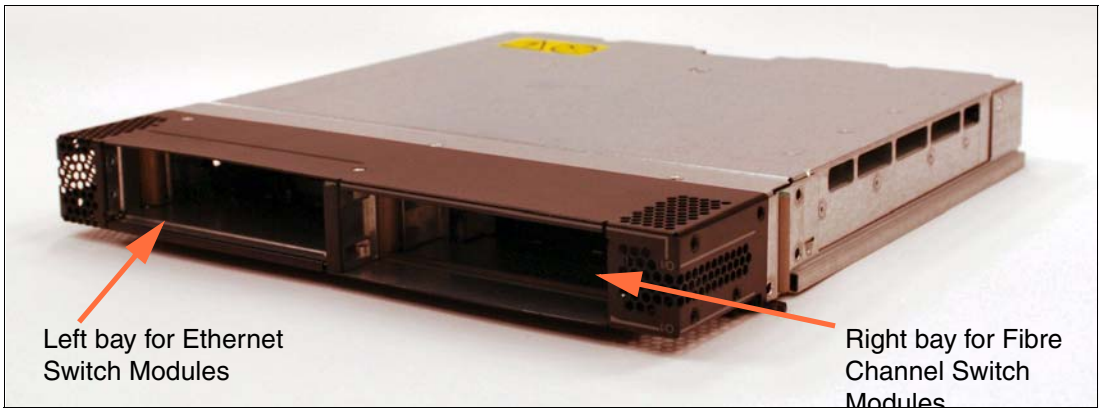

Figure 2-36 Multi-switch Interconnect Module

Each MSIM occupies high-speed slots 7 and 8 or 9 and 10, and up to two MSIMs can be installed into IBM BladeCenter H chassis. To use the MSIM, each blade must have certain CFFh expansion cards installed; refer to Table 2-3 on page 65 for more information.

The ports on the CFFh expansion cards in each server are hardwired to specific bays in the switch modules in each MSIM. Table 2-40 lists the mappings of expansion card ports to the I/O bays of MSIMs.

Table 2-40 Mapping of expansion card ports to the I/O bays of MSIMs

Port number of the CFFh expansion card	Corresponding switch module bay in the MSIM
1	7 (upper left interconnect module bay)
2	8 (upper right interconnect module bay)
3	9 (lower left interconnect module bay)
4	10 (lower right interconnect module bay)

In general, the MSIM supports almost all standard switch modules (high-speed switch modules are not supported with the MSIM). However, some restrictions might apply. Table 2-41 on page 176 lists supported I/O module options for the MSIM.

Currently, only three expansion cards support connectivity to the MSIM. As shown in Table 2-41 on page 176, the placement of modules in the MSIM varies between these expansion cards as shown in the rightmost two columns.

Table 2-41 Supported I/O module options for the MSIM

Part number	Feature code	Description	MSIM I/O bay[a]	
		Supported expansion card	QLogic Eth. & 4 Gb FC Exp. Crd (CFFh), 39Y9306, f/c 2968 QLogic Eth. & 8 Gb FC Exp. Crd (CFFh), 44X1940, f/c 5485	2/4 Port Ethernet Exp. Card (CFFh), 44W4479, f/c 5476
Ethernet				
32R1892	1498	Cisco Systems Intelligent Gb Ethernet Switch	Left	Both bays
32R1888	1497	Cisco Systems Intelligent Gb Fiber Ethernet Switch	Left	Both bays
41Y8523	2989	Cisco Catalyst Switch Module 3110G	Left	Both bays
41Y8522	2988	Cisco Catalyst Switch Module 3110X	Left	Both bays
43W4395	5450	Cisco Catalyst Switch Module 3012	Left	Both bays
39Y9324	1484	Server Connectivity Module	Left	Both bays
32R1859	1494	Nortel Layer 2-7 Gb Ethernet Switch	Not supported	Not supported
32R1860	1495	Nortel Layer 2/3 Copper Gb Ethernet Switch	Left	Both bays
32R1861	1496	Nortel Layer 2/3 Fiber Gb Ethernet Switch	Left	Both bays
32R1783	1493	Nortel 10 Gb Uplink Ethernet Switch	Left	Both bays
44W4404	1590	Nortel 1/10 Gb Uplink ESM	Left	Both bays
39Y9267	2952	Nortel 10 Gb Ethernet Switch Module	Not supported	Not supported
Fibre Channel				
32R1812	1569	Brocade 20 port 4 Gb SAN Switch	Right	Not supported
32R1813	1571	Brocade 10 port 4 Gb SAN Switch	Right	Not supported
39Y9280	2983	Cisco Systems 4 Gb 20 port Fibre Channel Switch	Right	Not supported
39Y9284	2984	Cisco Systems 4 Gb 10 port Fibre Channel Switch	Right	Not supported
26R0881	1560	QLogic 20 port 4 Gb Fibre Channel Switch	Right	Not supported
43W6725	2987	QLogic 20-port 4 Gb SAN Switch Module	Right	Not supported
43W6724	2986	QLogic 10-port 4 Gb SAN Switch Module	Right	Not supported
43W6723	2985	QLogic 4 Gb Intelligent Pass-thru Module	Right	Not supported
44X1905	5478	QLogic 20-Port 8 Gb SAN Switch Module	Right	Not supported
44X1907	5482	QLogic 8 Gb Intelligent Pass-thru Module	Right	Not supported
SAS				
39Y9195	2980	SAS Connectivity Module	Not supported	Not supported
43W3584	3734	SAS RAID Controller Module	Not supported	Not supported

Part number	Feature code	Description	MSIM I/O bay[a]	
		Supported expansion card	QLogic Eth. & 4 Gb FC Exp. Crd (CFFh), 39Y9306, f/c 2968 QLogic Eth. & 8 Gb FC Exp. Crd (CFFh), 44X1940, f/c 5485	2/4 Port Ethernet Exp. Card (CFFh), 44W4479, f/c 5476
InfiniBand				
32R1756	1574	Cisco Systems 4X InfiniBand Switch Module	Not supported	Not supported
39Y9211	2942	QLogic InfiniBand Fibre Channel Bridge Module	Not supported	Not supported
39Y9207	2941	QLogic InfiniBand Ethernet Bridge Module	Not supported	Not supported
Pass-thru Modules				
39Y9316	1556	Optical Pass-thru Module	Both bays	Both bays
39Y9320	2900	Copper Pass-thru Module	Left	Both bays
44W4483	5452	Intelligent Copper Pass-thru Module	Left	Both bays
43W4419	2990	4X InfiniBand DDR Pass-thru Module	Not supported	Not supported
43W6723	2985	QLogic 4 Gb Intelligent Pass-thru Module	Right	Not supported
44X1907	5482	QLogic 8 Gb Intelligent Pass-thru Module	Right	Not supported

a. *Right* means that it is only supported in the rightmost I/O slot of the MSIM. *Left* means that it is only supported in the leftmost I/O slot of the MSIM. *Both* means that it is supported in both the rightmost and leftmost I/O slots.

Figure 2-37 illustrates how it is possible to have eight I/O paths on one blade server. The blade must support the CFFh expansion card, it must be installed in a BladeCenter H or BladeCenter HT, and MSIM (or MSIM-HT, in the case of BladeCenter HT) must be installed in chassis. This is just one sample configuration. The connections are:

- The blade server has two onboard Gigabit Ethernet adapters that connect to Ethernet Switch Modules (ESMs) in bay 1 and bay 2.
- The blade server has a CFFv Gigabit Ethernet expansion card installed. See “Ethernet Expansion Card (CFFv)” on page 242 for a description of this card. The CFFv Gigabit Ethernet Expansion card connects to the ESMs in bay 3 and bay 4.
- The blade server also has a CFFh expansion card installed that provides two 1 Gb Ethernet ports and two 4 Gb Fibre Channel ports. See “QLogic Ethernet & 4Gb Fibre Channel Expansion Card (CFFh)” on page 246 for a description of this card.

  The ports on the CFFh card connect to switch modules that are installed in a MSIM. One MSIM is installed in bays 7 and 8 and a second MSIM is installed in bays 9 and 10. Each MSIM must have an ESM installed in the left I/O slot and a Fibre Channel Switch Module (FCSM) installed in the right I/O slot.

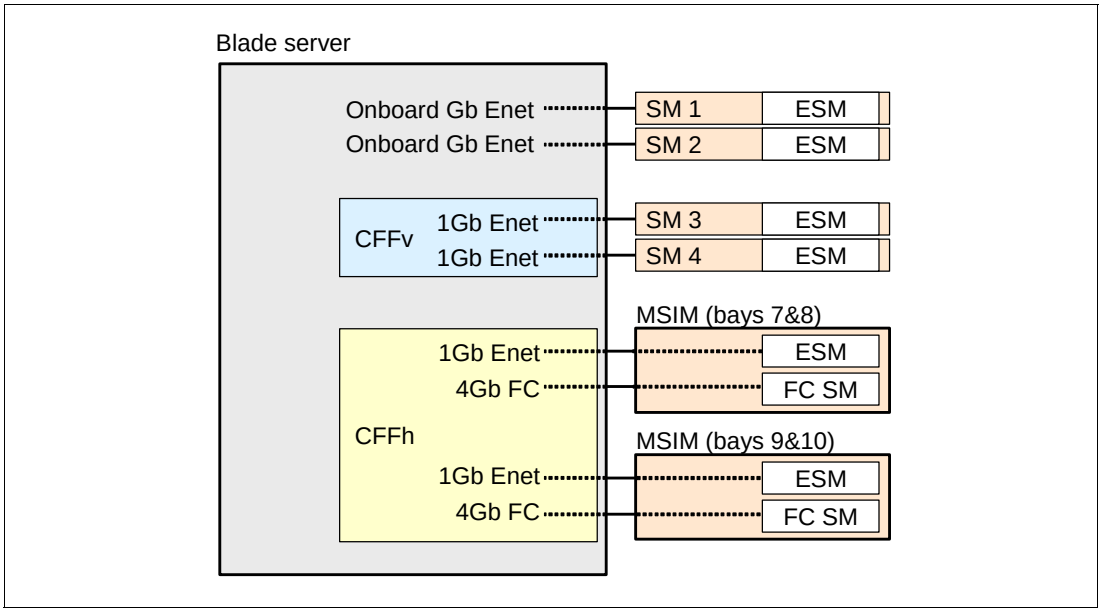

Figure 2-37 Example of eight I/O paths - six 1 Gb Ethernet and two 4 Gb Fibre Channels

Note: A BladeCenter chassis cannot have both an MSIM and high-speed switch module (HSSM). This is because the MSIM requires a CFFh-type expansion card, whereas the HSSM requires a high-speed form factor (HSFF)-type expansion card. All blades in one chassis must have the same high-speed card type (CFFh or HSSF).

For more information, refer to *Installation and User's Guide Multi-Switch Interconnect Module - IBM BladeCenter:*

`http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5070476`

2.6.5 Multi-Switch Interconnect Module for BladeCenter HT

The Multi-switch Interconnect Module for BladeCenter HT (MSIM-HT), part number 44R5913, feature code 5491, is a switch module container that fits in the high-speed switch bays (bays

7 and 8 and bays 9 and 10) of the BladeCenter HT chassis. The MSIM-HT accepts two supported standard switch modules as shown in Figure 2-36 on page 175, and enables a blade server to have a total of eight network connections.

Note: The MSIM-HT comes standard without any I/O modules installed; they need to be ordered separately. In addition, the use of MSIM-HT modules requires that all four power modules be installed in the chassis.

Figure 2-38 Multi-switch Interconnect Module for BladeCenter HT (without interposers)

Each MSIM-HT occupies high-speed slots 7 and 8 or 9 and 10. Up to two MSIM-HTs can be installed into the IBM BladeCenter HT chassis. To use the MSIM-HT, each blade must have compatible CFFh expansion cards installed; see Table 2-6 on page 72.

The ports on the CFFh expansion cards in each server are hardwired to specific bays in the switch modules in each MSIM-HT. Table 2-42 lists mappings of expansion card ports to the I/O bays of MSIMs.

Table 2-42 Mapping of expansion card ports to the I/O bays of MSIM-HTs

Port number of the CFFh expansion card	Corresponding switch module bay in the MSIM
1	7 (upper left interconnect module bay)
2	8 (upper right interconnect module bay)
3	9 (lower left interconnect module bay)
4	10 (lower right interconnect module bay)

MSIM-HT comes standard with two special interposers that allow you to fit it into high-speed bays of the BladeCenter HT chassis. One interposer is used when the MSIM is to be installed in the upper bays (7 and 8). The other interposer is used when the MSIM is to be installed in the lower bays (9 and 10). Figure 2-39 on page 180 shows the interposer for use with the lower bays.

Figure 2-39 Interposer used when the MSIM-HT is installed in the lower high-speed bays

MSIM-HT supports a limited set of standard switch modules (high-speed switch modules are not supported with the MSIM-HT). Table 2-43 lists supported I/O module options for the MSIM-HT.

Note: At the time of writing, HS21 (8853) and HS21 XM (7995) are the only supported blade servers for use with MSIM-HT.

Currently, only two expansion cards support connectivity to the MSIM-HT. As illustrated in Table 2-43, the placement of modules in the MSIM-HT varies between these expansion cards as shown in the rightmost two columns.

Table 2-43 Supported I/O module options for the MSIM-HT

Part number	Feature code	Description	MSIM-HT I/O bay[a]	
		Supported expansion card	**QLogic Eth. & 4 Gb FC Exp. Crd (CFFh), 39Y9306, f/c 2968**	**2/4 Port Ethernet Exp. Card (CFFh), 44W4479, f/c 5476**
Ethernet				
32R1892	1498	Cisco Systems Intelligent Gb Ethernet Switch	Left	Both bays
32R1888	1497	Cisco Systems Intelligent Gb Fiber Ethernet Switch	Not supported	Not supported
41Y8523	2989	Cisco Catalyst Switch Module 3110G	Not supported	Not supported
41Y8522	2988	Cisco Catalyst Switch Module 3110X	Not supported	Not supported
43W4395	5450	Cisco Catalyst Switch Module 3012	Not supported	Not supported
39Y9324	1484	Server Connectivity Module	Not supported	Not supported
32R1859	1494	Nortel Layer 2-7 Gb Ethernet Switch	Not supported	Not supported
32R1860	1495	Nortel Layer 2/3 Copper Gb Ethernet Switch	Left	Both bays

Part number	Feature code	Description	MSIM-HT I/O bay[a]	
		Supported expansion card	QLogic Eth. & 4 Gb FC Exp. Crd (CFFh), 39Y9306, f/c 2968	2/4 Port Ethernet Exp. Card (CFFh), 44W4479, f/c 5476
32R1861	1496	Nortel Layer 2/3 Fiber Gb Ethernet Switch	Not supported	Not supported
32R1783	1493	Nortel 10 Gb Uplink Ethernet Switch	Not supported	Not supported
44W4404	1590	Nortel 1/10 Gb Uplink ESM	Left	Both bays
39Y9267	2952	Nortel 10 Gb Ethernet Switch Module	Not supported	Not supported
Fibre Channel				
32R1812	1569	Brocade 20 port 4 Gb SAN Switch	Not supported	Not supported
32R1813	1571	Brocade 10 port 4 Gb SAN Switch	Not supported	Not supported
39Y9280	2983	Cisco Systems 4 Gb 20 port Fibre Channel Switch	Not supported	Not supported
39Y9284	2984	Cisco Systems 4 Gb 10 port Fibre Channel Switch	Not supported	Not supported
26R0881	1560	QLogic 20 port 4 Gb Fibre Channel Switch	Not supported	Not supported
43W6725	2987	QLogic 20-port 4 Gb SAN Switch Module	Right	Not supported
43W6724	2986	QLogic 10-port 4 Gb SAN Switch Module	Right	Not supported
43W6723	2985	QLogic 4 Gb Intelligent Pass-thru Module	Not supported	Not supported
44X1905	5478	QLogic 20-Port 8 Gb SAN Switch Module	Not supported	Not supported
44X1907	5482	QLogic 8 Gb Intelligent Pass-thru Module	Not supported	Not supported
SAS				
39Y9195	2980	SAS Connectivity Module	Not supported	Not supported
43W3584	3734	SAS RAID Controller Module	Not supported	Not supported
InfiniBand				
32R1756	1574	Cisco Systems 4X InfiniBand Switch Module	Not supported	Not supported
39Y9211	2942	QLogic InfiniBand Fibre Channel Bridge Module	Not supported	Not supported
39Y9207	2941	QLogic InfiniBand Ethernet Bridge Module	Not supported	Not supported
Pass-thru Modules				
39Y9316	1556	Optical Pass-thru Module	Not supported	Not supported
39Y9320	2900	Copper Pass-Thru Module	Not supported	Not supported
44W4483	5452	Intelligent Copper Pass-thru Module	Not supported	Not supported
43W4419	2990	4X InfiniBand DDR Pass-thru Module	Not supported	Not supported

Part number	Feature code	Description	MSIM-HT I/O bay[a]	
		Supported expansion card	QLogic Eth. & 4 Gb FC Exp. Crd (CFFh), 39Y9306, f/c 2968	2/4 Port Ethernet Exp. Card (CFFh), 44W4479, f/c 5476
43W6723	2985	QLogic 4 Gb Intelligent Pass-thru Module	Not supported	Not supported
44X1907	5482	QLogic 8 Gb Intelligent Pass-thru Module	Not supported	Not supported

a. *Right* means that it is only supported in the rightmost I/O slot of the MSIM-HT. *Left* means that it is only supported in the leftmost I/O slot of the MSIM-HT. *Both* means that it is supported in both the rightmost and leftmost I/O slots.

For more information, refer to *Multi-Switch Interconnect Module Installation Instructions - IBM BladeCenter HT*:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5073195

2.6.6 4X InfiniBand Pass-thru Module

With 14 InfiniBand 4X DDR ports internal to the servers and 14 InfiniBand 4X DDR ports external toward the upstream network, the 4X InfiniBand DDR Pass-thru Module, part number 43W4419, feature code 2990, offers full non-blocking 4X DDR InfiniBand support to all 14 blade servers in a BladeCenter H chassis. It supports both 10 Gbps single-data rate (SDR) and 20 Gbps double-data rate (DDR) InfiniBand connectivity.

The InfiniBand Pass-thru Module is a double-height module, and up to two can be installed in an IBM BladeCenter H, utilizing either switch bays 7 and 8, or switch bays 9 and 10 in the rear of the chassis.

Figure 2-40 on page 183 shows the 4X InfiniBand Pass-thru Module.

Note: Although the Pass-thru Module is a double-high module (to physically fit all the external connectors), it is only connected to the upper module bay. In other words, the module connects to bay 7 and bay 9. It is not connected to bay 8 and bay 10.

Figure 2-40 4X InfiniBand Pass-thru Module

Here are some important pass-thru module features:

- Provides a pass-through connection for 14 InfiniBand 4X DDR ports
- External ports: 14 copper InfiniBand 4X connectors
- Internal ports: 14 InfiniBand 4X ports connected to the midplane and on to the server slots
- Uplink ports that support hot-pluggable media converter for optical cable
- Two status LEDs: OK (green) and Fault (yellow)

The pass-thru module is managed by the IBM Advanced Management Module for IBM BladeCenter through a connection to the chassis midplane in the BladeCenter unit. Thus, the pass-thru module does not require and does not contain an InfiniBand Subnet Management (SM) function.

To enable communication between blade servers and the pass-thru module, you must install a compatible high-speed expansion card in each blade server that you want to communicate with the pass-thru module. See 3.1.15, “InfiniBand DDR Host Channel Adapters” on page 251 for information about these expansion cards.

Table 2-44 lists 4X InfiniBand Pass-thru module-related options

Table 2-44 4X InfiniBand Pass-thru module related options

Part number	Description
43W4419	4X InfiniBand Pass-thru Module for BladeCenter
43W6742	InfiniBand 3m DDR Cable
43W7243	InfiniBand 8m DDR Cable

Refer to *4X InfiniBand Pass-thru Module Installation and User Guide* for more information:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5072709

2.7 Management modules

The management module is a hot-swap device that you use to configure and manage all installed BladeCenter components. It provides system management functions and keyboard/video/mouse (KVM) multiplexing for all the blade servers in the BladeCenter unit. It controls an Ethernet and serial port connections for remote management access.

There are two Advanced Management Modules currently available as shown in Table 2-45:

- The Advanced Management Module for BladeCenter S, BladeCenter E, BladeCenter H, and BladeCenter HT, 25R5778
- The Advanced Management Module for BladeCenter T, 32R0835

Note: The original management modules (48P7055 for BladeCenter E and 90P3741 for BladeCenter T) have been withdrawn, but we describe them here in case you plan to add hardware to existing chassis installations. We refer them as MM in this section.

Table 2-45 Management modules supported in each BladeCenter chassis

Management module	Part Number	Feature Code	BC S	BC E	BC H	BC T	BC HT
Advanced management module	25R5778	1604	Yes	Yes	Yes	No	Yes
Advanced management module for BladeCenter T	32R0835	4239	No	No	No	Yes	No
Management module[a]	48P7055[a]	None	No	Yes	No	No	No
Management module for BladeCenter T[a]	90P3741[a]	None	No	No	No	Yes	No

a. This management module is now withdrawn and is not included with any currently shipping BladeCenter chassis.

To obtain the latest firmware for your management module, go to the Advanced Management Module firmware release matrix:

http://www.ibm.com/support/docview.wss?uid=psg1SERV-AMM

Table 2-46 lists all possible supported combinations of BladeCenter chassis and management modules, and summarizes a list of features that are supported by each configuration.

The term *standard switch module* refers to the traditional switch module such as Nortel L2/3 Gigabit Ethernet Switch Module that has a form-factor compatible with I/O bays 1, 2, 3 and 4 of each BladeCenter chassis (refer to the drawings of the rear of each of chassis in 1.1.4, “Supported operating systems” on page 6).

The term *high-speed switch module* refers to a high-speed switch module such as Cisco Systems 4X InfiniBand Switch Module that has a form-factor compatible with I/O bays 7, 8, 9, and 10 of IBM BladeCenter H and HT chassis (see Figure 1-8 on page 18).

Table 2-46 Features supported by chassis and management modules installed

	BC S	BC E		BC H	BC T		BC HT
Feature	AMM	MM	AMM	AMM	MM	AMM	AMM
Blade servers	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Standard switch modules	Yes	Yes[a]	Yes	Yes	Yes	Yes	Yes

	BC S	BC E		BC H	BC T		BC HT
Feature	AMM	MM	AMM	AMM	MM	AMM	AMM
High-speed switch modules	No	No	No	Yes	No	No	Yes
Concurrent KVM (cKVM)[b]	Yes	No	Yes	Yes	No	Yes	Yes
External USB ports	Yes	No	Yes	Yes	No	No	Yes
External PS/2 ports	No	Yes	No	No	Yes	Yes	No
External serial port	Yes	No	Yes	Yes	No	Yes	Yes
Active Energy Manager advancements	Yes	No	Yes	Yes	No	Yes	Yes

a. Some switches are supported with the Advanced Management Module only. Refer to Table 1-2 on page 3 for more information.

b. The Concurrent KVM feature card must be installed in the blade to support this functionality. Refer to Table 1-3 on page 5 for a list of supported blades.

Note: Mixing different types of management modules in the same chassis is not supported. Both the primary management module and the backup management module must be of the same type. BladeCenter S supports only one management module, and it comes standard with the BC S chassis.

The Advanced Management Module has different options for local and remote management:

- The 10/100 Mbps Ethernet port can be used for remote management using Web-based, CLI-based, or IBM Director management interfaces.
- The serial port in the front of the Advanced Management Module can be used for local CLI-based management. CLI-based management is useful for unattended remote configurations and batch processing. The Management Module supports only remote management via Ethernet port.
- The Advanced Management Module has two USB ports for a local keyboard and mouse. The Management Module uses PS/2 ports for the same function. You select which blade to control using the appropriate button on the blade.
- The Web interface of both the Advanced Management Module and the Management Module allows remote control of the mouse and keyboard (remote KVM).

The Advanced Management Module also supports concurrent remote KVM access if each blade server you want to be accessible concurrently has a cKVM feature card installed. See 3.17, “Concurrent KVM Feature Card (cKVM)” on page 356, for details.

Note: The Advanced Management Module is required for almost all newly announced blades and I/O modules.

In addition to power draw monitoring, which is also supported by the management module, the Advanced Management Module supports Active Energy Manager advancements such as the capability to cap power (see 2.14.3, “Power considerations” on page 218, for additional information).

2.7.1 Advanced Management Module

The Advanced Management Module is a hot-swap module that you use to configure and manage all installed BladeCenter components. The Advanced Management Module provides system management functions and keyboard/video/mouse (KVM) multiplexing for all blade servers in the BladeCenter unit that support KVM. It controls a serial port for remote connection; the external keyboard, mouse, and video connections for use by a local console; and a 10/100 Mbps Ethernet remote management connection. Figure 2-41 shows an Advanced Management Module.

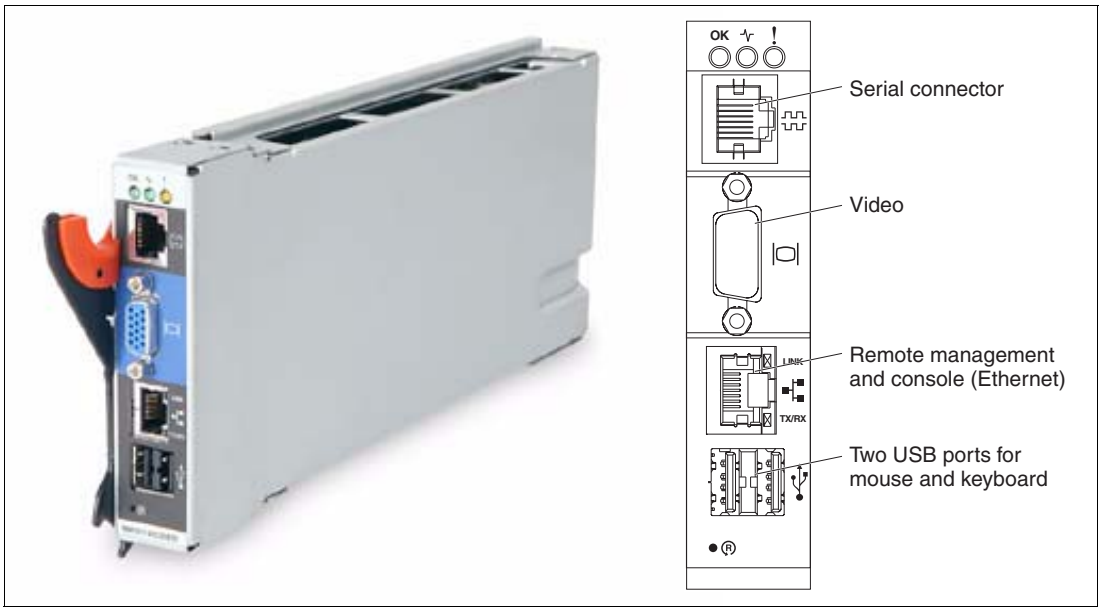

Figure 2-41 Advanced Management Module

All BladeCenter chassis come standard with at least one Advanced Management Module. Each chassis also supports a second management module for redundancy, part number 25R5778, feature code 1604. One of the management modules is active, and the second one, if installed, remains inactive until the management functions are manually switched over to it, or if the primary management module fails.

The service processor in the management module communicates with the service processor in each blade server to support features such as blade-server power-on requests, error and event reporting, KVM requests, and requests to use the BladeCenter shared media tray (removable-media drives and USB connector).

You configure BladeCenter components by using the management module, setting information such as IP addresses. The management module communicates with all components in the BladeCenter unit, detecting their presence or absence, reporting their status, and sending alerts for error conditions when required.

Management module network interface

Unlike the older management modules that required two IP addresses for their external and internal network interfaces, the Advanced Management Module requires only one IP address used for both external and internal communications. You should therefore plan your management network accordingly. For instance, use the same subnet for management module and I/O module management interfaces. For more information see Chapter 4, “Network integration” on page 371.

Note: We recommend that you use static IP addresses (or at least reservations for DHCP) for the Advanced Management Module to avoid management issues.

For more information, refer to the following product publications:

- *IBM BladeCenter Advanced Management Module User's Guide*
- *Advanced Management Module Command Line Interface Reference Guide*

Management module redundant configuration

You can configure one optional Advanced Management Module for higher availability by ordering part number 25R5778, feature code 1604.

Table 2-47 Advanced Management Module

Part number	Feature code	Description
25R5778	1604	IBM BladeCenter Advanced Management Module

If you plan to use management module redundancy, we recommend that you use another set of KVM cables for cabling redundant management modules. If a management module fails in this case, you can simply select another console port without having to manually reconfigure the KVM infrastructure.

The same rule applies to Ethernet cabling. We recommend that you connect Ethernet ports on both management modules to the network to be able to manage remotely in case of failover without any reconfiguration.

For more information about Advanced Management Module, refer to the following product documentation:

- *IBM BladeCenter Advanced Management Module for BladeCenter and BladeCenter H Installation Guide*
- *IBM BladeCenter Advanced Management Module User's Guide*
- *Advanced Management Module Command Line Interface Reference Guide*

For more information about hardware management features, refer to *IBM eServer xSeries and BladeCenter Server Management*, SG24-6495.

2.7.2 Advanced Management Module for BladeCenter T

The Advanced Management Module for BladeCenter T is a hot-swap unit that is installed into the front of the BladeCenter T unit; see Figure 2-42 on page 188. The BladeCenter T unit supports two management modules: one ships standard and the second is optional. If you

install two management modules in the BladeCenter T unit, they must be identical; that is, it is not supported to have one management module and one Advanced Management Module installed.

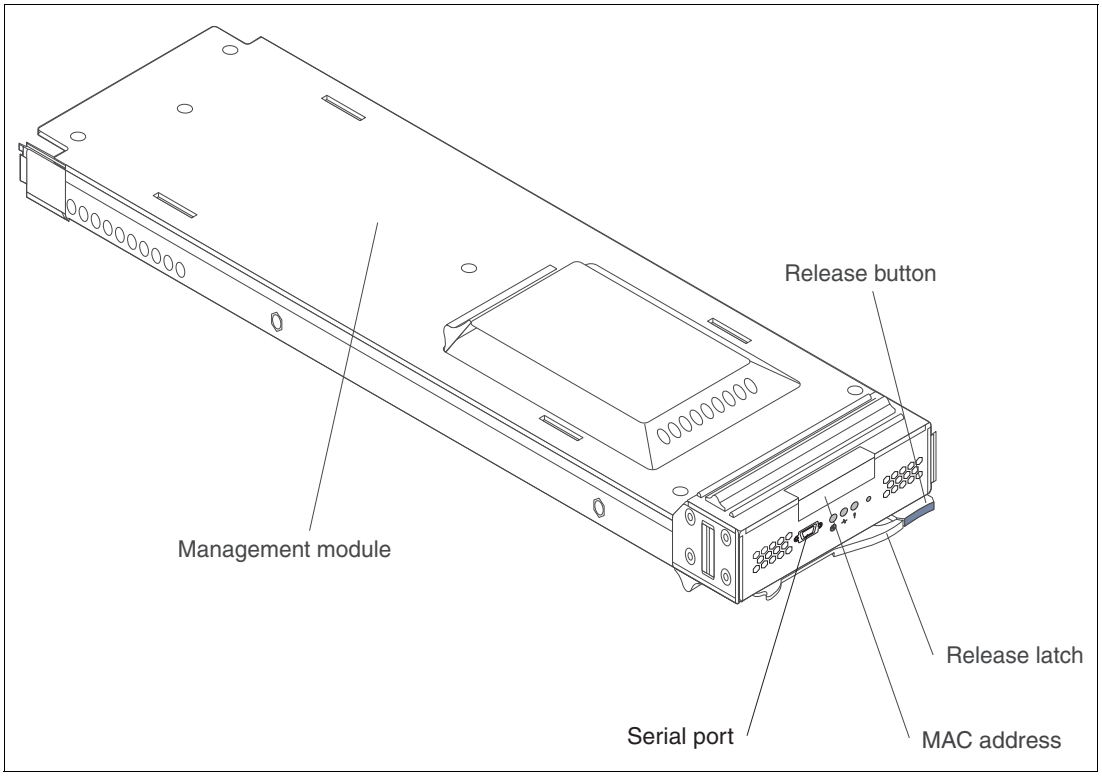

Figure 2-42 IBM BladeCenter T Advanced Management Module

Only one management module is active. The active management module is also known as the primary management module. If a secondary management module is installed in the BladeCenter T unit, this management module is in passive or standby mode by default. If the active or primary management module fails, the secondary management module is automatically enabled with all the configuration settings of the primary management module.

To configure and manage the installed BladeCenter components, the Advanced Management Module communicates through the external I/O connectors on the associated KVM and LAN modules at the rear of the BladeCenter T unit. The front of the management module also contains a serial connector (port) for use with the command-line interface (CLI).

The interface connections (such as PS/2 ports or Ethernet ports) of the Advanced Management Module for BladeCenter T are actually present, not on the management modules itself, but on the separate hot-swap LAN and KVM modules located at the rear of BladeCenter T chassis (see Figure 1-12 on page 21).

The LAN module, shown in Figure 2-43, features two RJ-45 connectors to provide 10/100 Mbps Ethernet connections to either of the two Management Modules at the front of the chassis. Ethernet port 1 is dedicated to management module 1 and Ethernet port 2 is dedicated to management module 2. If you have two management modules installed, you should have connections to both of these Ethernet ports to maintain connectivity. A DSUB 15P telco alarm connector is also provided.

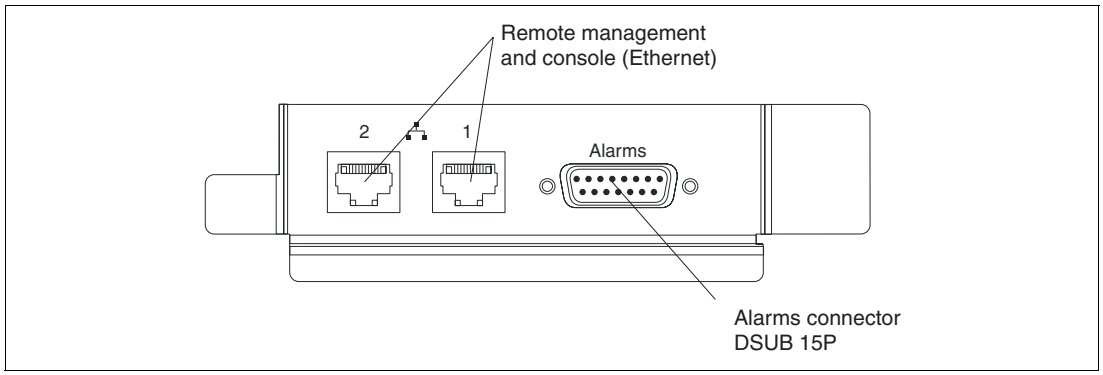

Figure 2-43 IBM BladeCenter T LAN module

The KVM module, shown in Figure 2-44, contains two PS/2 connectors for the keyboard and mouse, a system-status panel, and an HD-15 analog video connector.

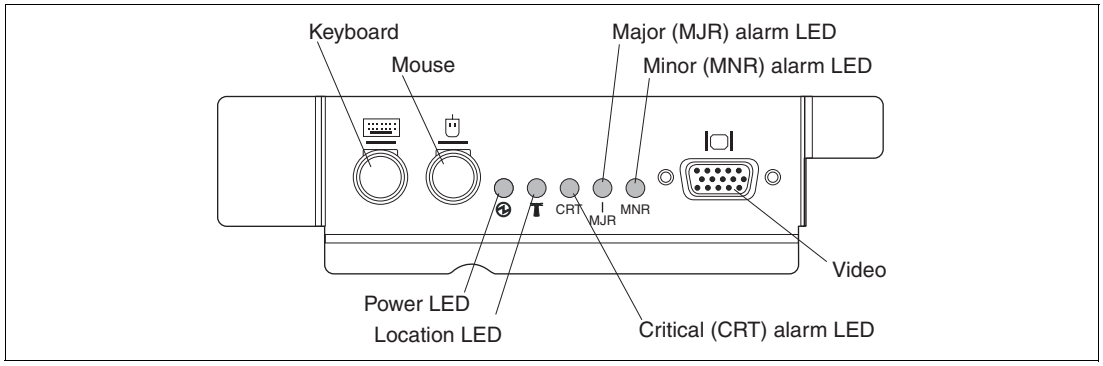

Figure 2-44 IBM BladeCenter T KVM Module

The Advanced Management Module performs system-management functions for the BladeCenter T unit. Through the management module, you can configure the BladeCenter T unit and modules as well as information such as the Internet protocol (IP) addresses of the management module and I/O modules. The management module also performs PS/2-to-USB conversions for the system keyboard and mouse and can send the VGA data stream to a remote console for viewing.

The management module communicates with the service processor in each blade server for functions such as blade server power-on requests, blade server error and event reporting, blade server requests for keyboard, mouse, and video, blade server requests for optical drive and USB ports.

In addition to blade servers, the management module also communicates with the I/O modules, power modules, blower modules, and the media tray in the BladeCenter T unit to detect their presence or absence and any error conditions, sending alerts when required.

Management module network interface

Unlike previous management modules that require two IP addresses for their external and internal network interfaces, the Advanced Management Module requires only one IP address, which is used for both external and internal communications. So your management network should be planned accordingly. For instance, use the same subnet for management module and I/O modules management interfaces. For more information see Chapter 4, “Network integration” on page 371.

Note: We recommend that you use static IP addresses (or at least reservations for DHCP) for the management module to avoid some management issues.

Management module redundant configuration

You can configure one optional redundant management module for higher availability by ordering part number 32R0835, feature code 4239.

Table 2-48 Redundant MM options for IBM BladeCenter T

Part Number	Feature Code	Description
32R0835	4239	IBM BladeCenter T Advanced Management Module

Only one KVM cable is required both for single and redundant management modules because only one KVM module is in the chassis.

LAN module has two Ethernet ports, as shown in Figure 2-43 on page 189. The right Ethernet port is dedicated to management module 1 and the left Ethernet port is dedicated to management module 2. If you want to maintain remote management capabilities in case of failover without reconnecting cables, connect both interfaces to the network.

For more information about the Advanced Management Module for BladeCenter T, refer to the following product documentation:

- *IBM BladeCenter Advanced Management Module for BladeCenter T Installation Guide*
- *IBM BladeCenter Advanced Management Module User's Guide*
- *Advanced Management Module Command Line Interface Reference Guide*

For more information about hardware management features, refer to *IBM eServer xSeries and BladeCenter Server Management*, SG24-6495.

2.8 Blower modules

Each of the different BladeCenter chassis have unique blower modules.

2.8.1 Blower module for BladeCenter E

The BladeCenter E chassis comes with two hot-swap blowers for 1+1 cooling redundancy; see Figure 1-4 on page 14.

Figure 2-45 IBM BladeCenter E blower

The blower speeds vary depending on the ambient air temperature at the front of the BladeCenter E:

- If the ambient temperature is 72°F or below, the BladeCenter E blowers will run at 30% of their maximum rotational speed, increasing their speed as required to control internal BladeCenter temperature.
- If the ambient temperature is above 72°F, the blowers will run at 80% of their maximum rotational speed, increasing their speed as required to control internal BladeCenter temperature.

If a blower fails, the remaining blower continues to cool the BladeCenter E unit and blade servers. Replace a failed blower as soon as possible, to restore cooling redundancy.

Refer to the *IBM BladeCenter Installation and User Guide* for more information.

In noise-sensitive environments, you can use the acoustic attenuation module; see 2.8.2, “Acoustic attenuation module” on page 192. The other way of limiting noise level is to use the “acoustic mode” setting in the Advanced Management Module. With this mode, the Advanced Management Module will throttle the processor speeds of the blades in order to stay within noise limits.

For more information refer to *IBM BladeCenter Advanced Management Module User Guide*.

2.8.2 Acoustic attenuation module

The Acoustic Attenuation Module (part number 39M4674), colloquially referred to as “the muffler” is an option for BladeCenter E that you can install over the blower modules in the rear of the chassis to reduce decibels in sound-sensitive environments. BladeCenter E generates 74 decibels (7.4 bels) at maximum performance levels. The Acoustic Attenuation Module reduces the decibel level to 69 decibels using a T-shaped baffle; see Figure 2-46.

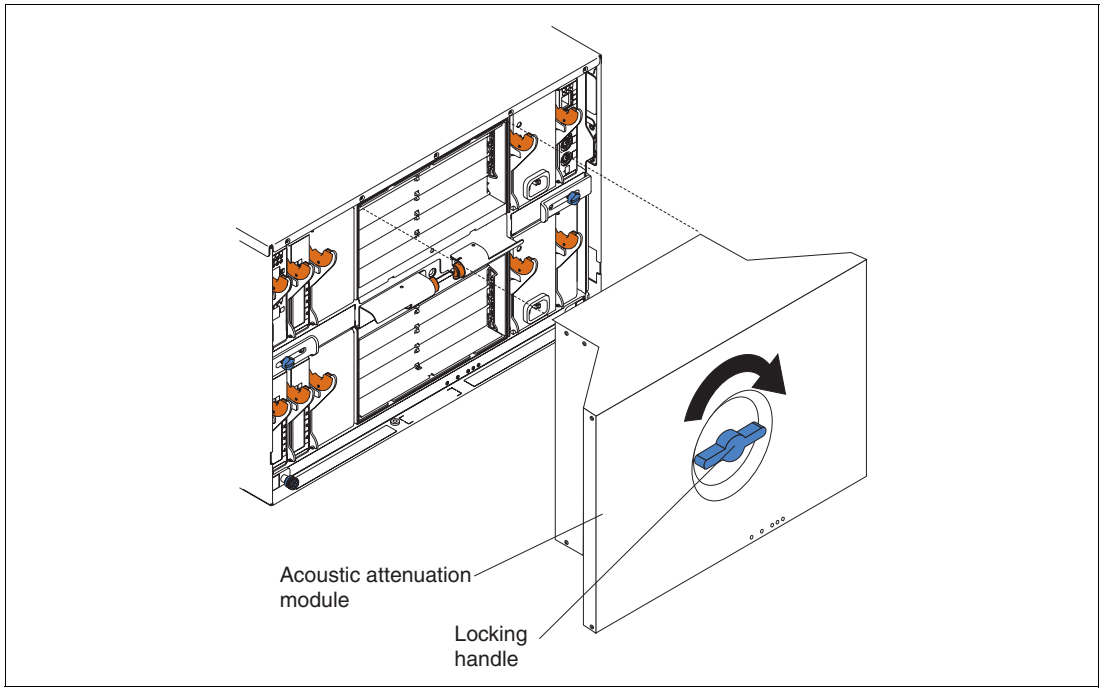

Figure 2-46 IBM BladeCenter E Acoustic Attenuation Module

For more information, refer to *Acoustic Attenuation Module User's Guide*.

Note: There is no Acoustic Attenuation Module for the BladeCenter H, HT or T chassis.

2.8.3 Blower modules for BladeCenter H

The BladeCenter H chassis comes with two hot-swap blowers for 1+1 cooling redundancy; see Figure 1-8 on page 18.

Figure 2-47 IBM BladeCenter H blower

The blower speeds vary depending on the ambient air temperature at the front of the BladeCenter H unit and the temperature of internal BladeCenter H components.

- If the ambient temperature is 25°C (77°F) or below, the BladeCenter H unit blowers will run at their minimum rotational speed, increasing their speed as required to control internal BladeCenter temperature.
- If the ambient temperature is above 25°C (77°F), the blowers will run faster, increasing their speed as required to control internal BladeCenter H unit temperature.

If a blower fails, the remaining blower will run full speed and continues to cool the BladeCenter H unit and blade servers. Replace a failed blower as soon as possible, to restore cooling redundancy.

Refer to the *IBM BladeCenter H Installation and User Guide* for more information.

2.8.4 Blower modules for BladeCenter T

The blower modules are hot-swap units that are installed into the rear of the system; see Figure 1-11 on page 21. The BladeCenter T unit comes with four blowers that are in a 3+1 redundancy configuration. A blower module is shown in Figure 2-48.

Figure 2-48 IBM BladeCenter T Blower Module

All the cooling requirements are still met if one blower fails. All blowers contain a backflow device that prevents the system from drawing air into the exhaust port of a failed blower. The management module in the BladeCenter T unit controls the blower speed and detects blower failures.

Refer to *IBM BladeCenter T Installation and User Guide* for more information.

2.8.5 Fan modules for BladeCenter HT

The BladeCenter HT unit comes standard with four hot swap fan modules for N+1 cooling redundancy. The fan module for BladeCenter HT is shown in Figure 2-49.

Figure 2-49 Fan module for BladeCenter HT

Each fan module contains two fans operating as a pair in a series. If one fan fails, the remaining fan runs at full speed and continues to cool the BladeCenter HT unit. Replace a failed fan module as soon as possible to restore cooling redundancy.

Fan module speeds vary depending on the ambient air temperature within the BladeCenter HT unit, which is reported by the media tray.

- If the ambient temperature is 25°C (77°F) or below, the BladeCenter HT unit's fan modules will run at their minimum rotational speed, increasing their speed as required to control the internal BladeCenter HT temperature.
- If the ambient temperature is above 25°C (77°F), the fan modules will run faster, increasing their speed as required to control the internal BladeCenter HT unit temperature.

Refer to *IBM BladeCenter HT Installation and User's Guide* for more information.

2.8.6 Fan modules for BladeCenter S

The BladeCenter S chassis comes with four installed hot-swap fan modules. The fan modules (sometimes called fan packs) are designed to provide N+1 redundant cooling airflow to the blade servers and I/O modules. Storage modules are cooled by fans located in the power supply modules.

Each fan module contains two fans. The fan module for BladeCenter S is shown in Figure 2-50.

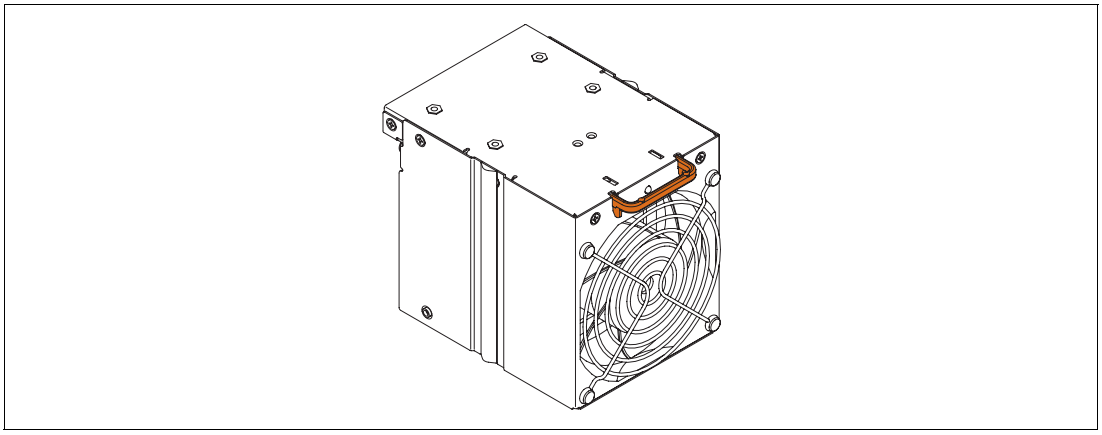

Figure 2-50 Fan module for BladeCenter S

If one fan module fails, the remaining fan modules run at full speed and continue to cool the BladeCenter S unit. Replace a failed fan module as soon as possible to restore cooling redundancy.

Refer to *IBM BladeCenter S Installation and User's Guide* for more information.

2.9 Power modules

Different IBM BladeCenter chassis use different power supplies supporting the unique power requirements of different data center computing environments. They include:

- 2000W AC power supplies for IBM BladeCenter E chassis for traditional enterprise environments
- 1300W AC or DC power supplies for IBM BladeCenter T for Telco environments
- 2900W AC power supplies for IBM BladeCenter H
- 3160W AC and 2535W DC power supplies for IBM BladeCenter HT supporting current and future computing power requirements for high-speed environments

- 950W/1450W (110/220 V) power supplies for BladeCenter S that can be plugged into traditional wall jack

2.9.1 Power modules for BladeCenter E

There are two power modules standard with the IBM BladeCenter E, with support for a maximum of four. A power module is shown in Figure 2-51.

Figure 2-51 Power Supply for IBM BladeCenter E

The IBM BladeCenter E unit comes with one pair of 2000W hot-swap power modules in power bays 1 and 2 and with two IEC 320-C19 to C20 power cables. Each power supply has its own IEC 320-C20 power connector. The power supplies in bays 1 and 2 provide power to all the I/O and management modules and to blade bays 1 through 6.

The BladeCenter E unit supports a second pair of power modules in power bays 3 and 4 that provide power to blade bays 7 through 14 (see Figure 1-2 on page 13 and Figure 1-4 on page 14). Table 2-49 summarizes the application for each power module.

Table 2-49 Power module functions

Bays	Power module function
1 and 2	Provides power to all modules and to blade bays 1 through 6
3 and 4	Provides power to blade bays 7 through 14

Power modules are not needed in bays 3 and 4 until you begin installing blade servers and options in blade bays 7 through 14. Table 2-50 lists the power supply option available for the IBM BladeCenter E chassis.

Table 2-50 Power supply option for IBM BladeCenter E

Part Number	Feature Code	Description
39M4675	1966	IBM BladeCenter E 2000W Power Supply Modules[a]

a. This option contains two power supplies and two IEC 320-C19 to C20 power cables.

Note: If you install a blade server that has a storage expansion unit option attached in blade bay 6, the option will use blade bay 7. Power modules will be required in power bays 1, 2, 3 and 4.

Each pair of power modules is redundant. If either power module fails, the remaining power module continues to supply power, but there is no redundancy; the failed power module must be replaced as soon as possible.

Important: Because the existing chassis might have lower capacity power supplies (1200, 1400 or 1800 W), it is important to ensure that the pair of power modules in the same power domain has the same capacity.

To provide true redundant power, power modules 1 and 3 must connect to a different AC power source than power modules 2 and 4. Connect power modules 1 and 3 to a different PDU than power modules 2 and 4. Then, connect each PDU to an AC power source (building power source or service entrance) that is controlled by a separate circuit breaker. See 2.14.3, “Power considerations” on page 218, for more information about power requirements for BladeCenter E.

2.9.2 Power modules for BladeCenter H

The BladeCenter unit comes with two 2900W hot-swap power modules. The power module for BladeCenter H is shown in Figure 2-52.

Figure 2-52 IBM BladeCenter H Power Supply

The standard power modules in bays 1 and 2 are used to power blade servers in blade bays 1 through 7 and I/O modules in I/O module bays 1 through 4 and 7 through 10.

Optional power modules are needed in power module bays 3 and 4 if you install blade servers in blade bays 8 through 14, or if you install I/O modules in any of I/O module bays 5 through 10 (see Table 2-51 on page 198).

Table 2-51 Devices powered by each power module

Devices[a]	Power module 1 (Standard)	Power module 2 (Standard)	Power module 3 (Optional)	Power module 4 (Optional)
Blades 1-7	Yes	Yes		
Blades 8-14			Yes	Yes
I/O modules 1 & 2 (SM 1-2)[b]	Yes	Yes	Yes	Yes
I/O modules 3 & 4 (SM 1&2/BM 3&4)	Yes	Yes		
I/O modules 5 & 6 (BM 1 & 2)			Yes	Yes
I/O modules 7 & 8 (HSSM 1 & 3)	Yes		Yes	
I/O modules 9 & 10 (HSSM 2 & 4)		Yes		Yes
Management modules[b]	Yes	Yes	Yes	Yes

a. SM = switch module; BM = bridge module; HSSM = high speed switch module.
b. The two management modules and I/O modules 1 and 2 are powered by any power module.

Note: Management module bays 1 and 2 and I/O modules 1 and 2 are powered from any or all power; that is, any of four power supplies can supply power to them.

See Table 2-52 for the power supply option available for IBM BladeCenter H chassis.

Table 2-52 Power supply option for IBM BladeCenter H

Part number	Feature code	Description
31R3335	1962	IBM BladeCenter H 2900W AC Power Module Pair with Fan Packs

Each pair of power modules is redundant. If either power module fails, the remaining power module continues to supply power, but there is no redundancy; the failed power module must be replaced as soon as possible.

Each power supply has its own 3-fan pack used for power supply cooling. The power supply itself has no external power connectors. Two special power connectors are located at the rear of the chassis (see Figure 1-8 on page 18).

No power cables are shipped with either the BladeCenter H chassis or the power module option. You should order them separately. Table 2-53 on page 199 lists the cable options available for BladeCenter H. You need two cables for each BladeCenter H chassis. Refer to 2.14.3, “Power considerations” on page 218, for more information about power considerations for IBM BladeCenter H.

Note: The photos in Table 2-53 show BladeCenter chassis connectors with silver rings. These are pre-production level connectors. The production-level connectors have brown plastic connectors; the color matches the sockets at the back of the chassis.

You cannot connect a pre-production power cable with a silver ring to a production-level chassis with the brown power socket.

Table 2-53 Power cable options for IBM BladeCenter H

Part Number	Feature code	Description	Photo
25R5783	6270	4.3m 208V Double 30A NEMA L6-30P	
25R5784	6271	4.3m 230V Dual 32A IEC 309 P+N+G/16A IEC 320-C20	
25R5785	6226	2.8m 200-240V Triple 16A IEC 320-C20	
25R5811	6273	4.3m 220V Double 30A KSC 8305 (for South Korea)	
25R5812	6272	4.3m 230V Dual 32A AS/NZS 3112/16A IEC 320-C20 (for Australia/NZ)	

See 2.14.3, “Power considerations” on page 218, for more information about power requirements for BladeCenter H.

2.9.3 Power modules for BladeCenter T

The IBM BladeCenter T chassis is available in two types with either DC power (8720) or AC power (8730) capabilities.

Note: After a decision has been made for one type of power supply, a conversion to the other type of power supply is impossible. If an initial configuration comes with AC power supplies, then it cannot be changed to the DC power. It is not simply a matter of different power supply modules; rather, the internal cabling within the chassis is different and cannot be changed after it ships from the factory.

Figure 2-53 IBM BladeCenter T DC power supply

The BladeCenter T 8720 unit comes with one pair of 1300W DC hot-swap power modules in the power module bays 1 and 2. The chassis supports a second pair of power modules in the power module bays 3 and 4. Power modules are not needed in bays 3 and 4 unless you install blade servers and options in blade bays 5 through 8, or I/O modules in I/O-module bays 3 or 4 (see Table 2-54 on page 202).

The BladeCenter T unit does not have a power switch. The BladeCenter T 8720 (DC power) unit has two DC power terminal connectors on the rear, each powering two power modules (there is no external power connector on the power supply itself).

The left connection (looking from the rear) supplies power to the power module bays 2 and 4. The connection on the right supplies power to the power module bays 1 and 3. Each DC terminal has four #M6 (0.25-inch) studs, one for -48 V DC, one for RETURN, and two for connecting the safety ground wire.

To provide true redundant power, BladeCenter T power modules 1 and 3 must be connected to a different power source than power modules 2 and 4.

Figure 2-54 IBM BladeCenter T AC power supply

The BladeCenter T 8730 unit comes with one pair of 1300W hot-swap AC power modules in the power module bays 1 and 2. The BladeCenter T unit supports a second pair of power modules in the power module bays 3 and 4.

There are four IEC320 power connectors on the rear of the Type 8730 (AC power) chassis, marked 1- 4 on the rear panel (there is no external power connector on the power supply itself). Power is applied to the corresponding power modules according to the numbering convention on the rear panel. For example, power connector 1 supplies power to the power module 1, and so on.

The BladeCenter T unit does not have a power switch. To start the BladeCenter T Type 8730 (AC power) unit, connect one end of a power cord into the input power connector 1 and 2 on the rear of the BladeCenter T unit. Connect the other end of each power cord into a 220-volt power distribution unit (PDU) that is connected to an appropriate electrical outlet.

Important: In a redundant pair of power modules, a power module that is not connected to a 220-volt power source creates a nonredundant condition.

If you have a second pair of power modules to install in the power module bays 3 and 4, then connect power cords to the input power connectors 3 and 4 on the rear of the BladeCenter T unit. Connect the other end of each power cord into a 220-volt power distribution unit (PDU) that is connected to an appropriate electrical outlet.

To provide true redundant power, the BladeCenter T power modules 1 and 3 must be connected to a different power source than power modules 2 and 4. Connect the BladeCenter T power modules 1 and 3 to a different PDU than power modules 2 and 4. Then connect each PDU to an AC power source (a building power source or service entrance) that is controlled by a separate circuit breaker.

Table 2-54 on page 202 lists the power module functions for BladeCenter T.

Table 2-54 Power module functions for BladeCenter T

Power module bays	Power module function
1 and 2	Provides power to ► I/O bays 1 and 2 ► Blade bays 1 through 4 ► Both management module bays ► Media tray
3 and 4	Provides power to ► I/O bays 3 and 4 ► Blade bays 4 through 8

Table 2-55 lists the power supply options available for the IBM BladeCenter T chassis.

Table 2-55 Power supply options for IBM BladeCenter T

Part Number	Feature code	Description
32R0833	1958	IBM BladeCenter T 1300W DC Power Supply Option[a]
32R0834	1960	IBM BladeCenter T 1300W AC Power Supply Option[b]

a. This option contains two power supplies.
b. This option contains two power supplies and two IEC 320-C13 to C14 power cables.

See 2.14.3, “Power considerations” on page 218, for more information about power requirements for BladeCenter T.

2.9.4 Power Modules for BladeCenter HT

The IBM BladeCenter HT chassis is available in two types with either DC power (8740) or AC power (8750) capabilities.

Note: After a decision has been made for one type of power supply, a conversion to the other type of power supply is impossible. If an initial configuration comes with AC power supplies, then it cannot be changed to DC power.

Both AC and DC power modules for BladeCenter HT look the same (see Figure 2-52 on page 197).

Figure 2-55 IBM BladeCenter HT Power Supply

The BladeCenter HT unit comes with one pair of DC (8740) or AC (8750) hot-swap power modules in the power module bays 1 and 2. The chassis supports a second pair of power modules in the power module bays 3 and 4. Power modules are not needed in bays 3 and 4 unless you install blade servers and options in blade bays 7 through 12, or I/O modules in I/O-module bays 7 through 10 (see Table 2-56).

The BladeCenter HT unit does not have a power switch. The BladeCenter HT 8740 (DC power) unit has four DC power terminal connectors on the rear, each powering one power module. Each DC terminal has four #M6 (0.25-inch) studs, one for -48 V DC, one for RETURN, and two for connecting the safety ground wire. There are no power cables supplied with BladeCenter HT 8740.

The BladeCenter HT 8750 (AC power) has four standard IEC309-C20 power connectors on the rear, each powering one power module. BladeCenter HT8750 comes standard with two C19/C20 power cables.

Each power supply has its own 3-fan pack used for power supply cooling. The power supply itself has no external power connectors.

As viewed from the rear of the BladeCenter HT unit, power connector 1 (top right connector) supplies power to the power module bay 1, power connector 2 (bottom left connector) supplies power to the power module 2, power connector 3 (top left connector) supplies power to the power module 3, and power connector 4 (bottom right connector) supplies power to the power module 4 (see Figure 1-16 on page 29).

To provide true redundant power, BladeCenter HT power modules 1 and 3 must be connected to a different power source than power modules 2 and 4.

The standard power modules in bays 1 and 2 are used to power blade servers in blade bays 1 through 6 and I/O modules in I/O module bays 1 through 4. Optional power modules are needed in power module bays 3 and 4 if you install blade servers in blade bays 7 through 12 or if you install I/O modules in high-speed I/O module bays 7 through 10 (see Table 2-56).

Table 2-56 Devices powered by power modules

Power module bays	Power module function
1 and 2	Provides power to management modules 1 and 2, I/O modules 1 through 4, and to blade bays 1 through 6
3 and 4	Provides power to high-speed I/O modules 7 through 10, and blade bays 7 through 12

Table 2-57 lists the power supply options available for the IBM BladeCenter HT chassis.

Table 2-57 Power supply options for IBM BladeCenter HT

Part number	Feature code	Description
42C5279	1984	IBM BladeCenter HT DC Power Supply Option[a]
42C5280	1983	IBM BladeCenter HT AC Power Supply Option[b]

a. This option contains two DC power supplies with fan packs.
b. This option contains two AC power supplies with fan packs and two 2.8m C19/C20 power cables.

Each pair of power modules is redundant. If either power module fails, the remaining power module continues to supply power, but there is no redundancy. The failed power module must be replaced as soon as possible.

Refer to 2.14.3, “Power considerations” on page 218, for more information about power considerations for IBM BladeCenter HT.

2.9.5 Power modules for BladeCenter S

The BladeCenter S supports up to four auto-sensing power modules, which are capable of supporting either 110 V or 220 V AC power. There are two power supplies for BladeCenter S based on the input connector, as listed in Table 2-58.

Table 2-58 Power supply connectors by chassis model

BladeCenter S model	Input connector	Availability
8886-1MY	IEC 320 C20	Worldwide
8886-1NG	IEC 320 C14	Denmark, Switzerland and Sweden

Two power modules are standard, and a maximum of four power modules are supported. Figure 2-56 shows one of the power modules.

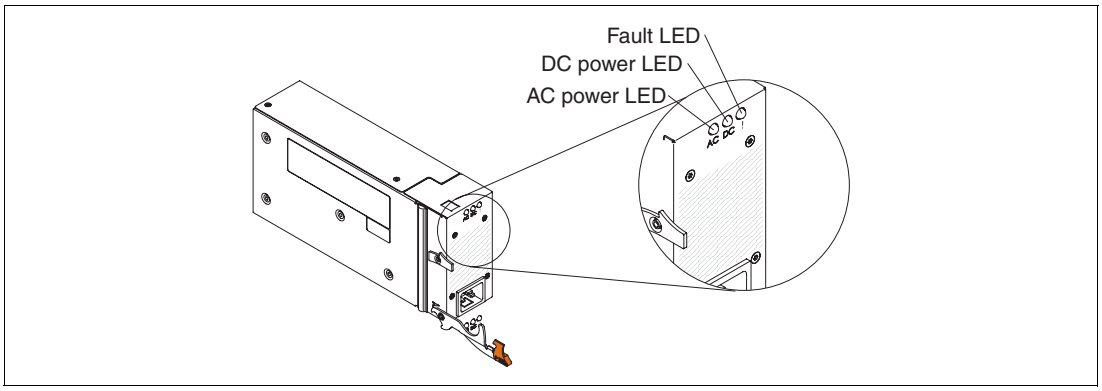

Figure 2-56 BladeCenter S power module

Within the BladeCenter S chassis, all power supplies are combined into a single power domain that distributes power to each of the blade servers and modules through the system midplane.

The second pair of power modules is required if *any* of these situations occur:

- The power requirements of the installed components (servers, I/O modules, disks, and so forth) exceed the capacity of the standard two power modules.
- You install the second storage module, because power modules 3 and 4 also provide the necessary fans to cool this second storage module.
- The power profile selected requires more power supplies for redundancy. See “IBM BladeCenter S specific considerations” on page 221 for more information about this topic.

Use the BladeCenter Power Configurator to determine whether your configuration requires the second pair of power supplies:

http://www.ibm.com/systems/bladecenter/powerconfig

The power modules are hot-swappable components and can easily be replaced during normal BladeCenter operation, assuming a redundant power policy has been selected in the AMM. If a power supply fails, the cooling fans inside the power supply continue to operate normally, because the power supply fans are powered from the “common” voltage from the

midplane. This is important to note, because the power supply fans provide the airflow to cool the storage modules.

The power modules are auto-sensing and can support either 110 V or 220 V AC power. However, do not mix voltage power sources within the same BladeCenter S chassis. If you have a mix of 110 V and 220 V input power sources, the chassis detects this mix and does not allow some power supplies to function. In this situation, the DC LEDs of the power supplies blocked will not illuminate. The AMM will also post messages to the error log.

The BladeCenter S 8886-1MY ships standard with two 2.5 m IEC 320 C19 - C20 power jumper cords and four country-specific line cords (except model 8886-1MG for the EMEA geography, which does not include country-specific power cords).

BladeCenter S 8886-1NG ships with two 2 m IEC 320 C13 - C14 power jumper cords.

Table 2-59 lists the power supply and power cord options available for BladeCenter S models.

Table 2-59 Power supply and power cable options for IBM BladeCenter S

Part number	Feature code	Description
43W3582	1992	950W/1450W Auto-Sensing Power Supplies 3 and 4[a]
46C7438	2102	C14 950W/1450W Auto-Sensing Power Supplies 3 and 4[b]
40K9766	None	C19 4.3m Line Cord - Europe
40K9767	None	C19 4.3m Line Cord - UK
40K9768	None	C19 4.3m Line Cord - Italy
40K9769	None	C19 4.3m Line Cord - Switzerland/Denmark
40K9770	None	C19 4.3m Line Cord - South Africa
40K9771	None	C19 4.3m Line Cord - Israel
39Y7917	None	2.8m, 10A/230V, C13 to CEE7-VII (Europe) Line Cord
39Y7922	None	2.8m, 10A/230V, C13 to SABS 164 (South Africa) Line Cord
39Y7923	None	2.8m, 10A/230V, C13 to BS 1363/A (UK) Line Cord
39Y7921	None	2.8m, 10A/230V, C13 to CEI 23-16 (Italy) Line Cord
39Y7920	None	2.8m, 10A/230V, C13 to SI 32 (Israel) Line Cord
39Y7918	None	2.8m, 10A/230V, C13 to DK2-5a (Denmark) Line Cord
39Y7919	None	2.8m, 10A/230V, C13 to SEV 1011 (Sws) Line Cord

a. This option ships with two 2.5 m IEC 320 C19 - C20 power jumper cords.
b. This option ships with two 2.0 m IEC 320 C13 - C14 power jumper cords.

2.10 Direct serial connections

Direct serial connections are used to provide hardwired, dedicated terminal console access to blade servers for management purposes. Signals from a blade's serial port are routed through the chassis midplane connections to the outside. To support this a blade server must have the required wire, and be installed in a chassis that supports this wire.

Direct serial connections are supported by:

- BladeCenter H and HT chassis with Serial Port Breakout Cable
- BladeCenter S with the Serial Module installed

Table 2-60 lists information about the blades that support this feature.

Table 2-60 Blade servers that support direct serial connections

Blade	Machine type	Direct serial connection support
HS12	8014	Supported
HS12	8028	Supported
HS20	8678	No
HS20	8832	No
HS20	8843	Supported
HS20	7981	Supported
HS21	8853	Supported
HS21 XM	7995	Supported
HS40	8839	No
LS20	8850	Supported
LS21	7971	Supported
LS22	7901	Supported
LS41	7972	Supported
LS42	7902	Supported
JS12	7998-60X	Supported
JS20	8842	No
JS21	8844	Supported
JS22	7998	Supported
QS21	0792	Supported
QS22	0793	Supported
PN41	3020	No
HC10	7996	Supported

2.10.1 Serial Port Breakout Cable for IBM BladeCenter H

The Serial Port Breakout Cable for BladeCenter H, part number 40K9605, feature code 4811, connects to the rear of the BladeCenter H (lower-left corner of the rear of the chassis, see Figure 1-8 on page 18). This cable provides 14 serial connections for terminal access, one to each supported blade server. The cable is shown in Figure 2-57 on page 207.

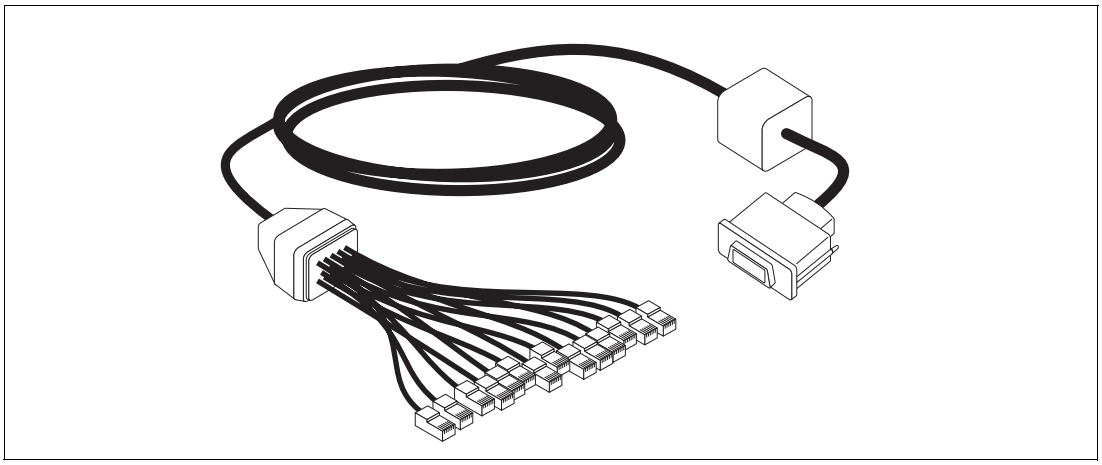

Figure 2-57 BladeCenter H Serial Port Breakout Cable

Not all blade servers support this option. Table 2-60 on page 206 lists the blades that support this option when the blade is installed into a BladeCenter H or HT chassis.

Each serial connector is an RJ-45. The pin assignments of the RJ-45s are listed in Table 2-61. The maximum baud rate for serial connection is 19.2 Kbps.

Table 2-61 Pin assignments for the RJ-45 connectors

Contact (pin number)	Signal name	Signal direction
1	RTS (request to send)	Output from blade
2	Not used	N/A
3	RXD (receive data)	Input to blade
4	GND	N/A
5	Not used	N/A
6	TXD (transfer data)	Output from blade
7	Not used	N/A
8	CTS (clear to send)	Input to blade

Note: There are four signals plus a ground routed to each of the RJ45 connectors as listed in Table 2-61 on page 207. However, not all standard RS232 signals are implemented. For example, DSR/RI, DCD, and DTR are not implemented. Devices such as serial modems which require all of the standard serial port signals will not work and are not supported when attached to the direct serial cable connection. The RJ45 port connection was designed for use with a serial terminal TTY concentrator.

2.10.2 BladeCenter S Serial Pass-thru Module

The Serial Pass-thru Module, part number 43W3583, feature code 1584, provides six serial port connectors that can be used to directly attach to each blade server in the BladeCenter S chassis via a four-wire serial RJ-45 connector. Port connector links bypass the AMM and provide a dedicated link directly to each blade.

The port connections function at speeds of up to 19.2k baud. They are intended for serial console access only. If used, the module must be installed in the Serial Pass-thru Module Bay (see Figure 1-18 on page 32). The port connectors are numbered from 1 to 6, from top to bottom, and correspond to blade servers in blade server bays 1 through 6.

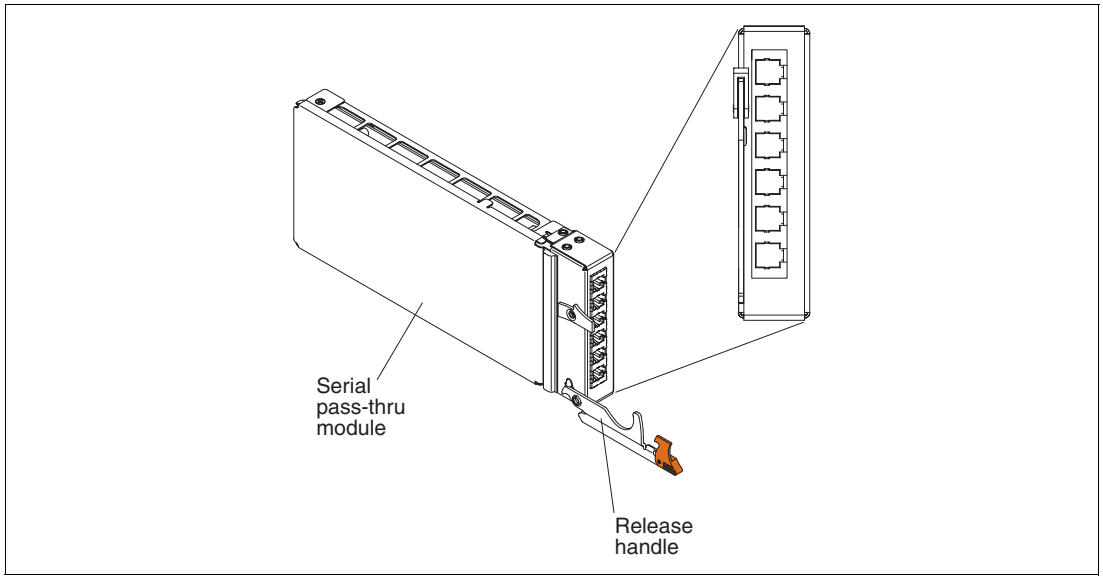

Figure 2-58 Serial Pass-thru Module

The serial cable required to access the ports on the Serial Pass-thru Module uses standard CTS/RTS and TXD/RXD signaling to provide console access. The cable's RJ-45 pin configuration is described in Figure 2-59.

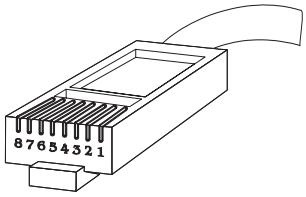

Contact (Pin)	Signal name	Signal direction
1	RTS - Request to send	Output from blade server
2	Not used	N/A
3	RXD - Receive Data	Input from blade server
4	GND	N/A
5	Not used	N/A
6	TXD - Transfer Data	Output from blade server
7	Not used	N/A
8	CTS - Clear to send	Input to blade server

Figure 2-59 RJ-45 pin location reference diagram

Note: There are four signals plus a ground routed to each of the RJ45 connectors as described in Figure 2-59. However, not all standard RS232 signals are implemented. For example, DSR/RI, DCD, and DTR are not implemented. Devices such as serial modems which require all of the standard serial port signals will not work and are not supported when attached to the direct serial cable connection. The RJ45 port connection was designed for use with a serial terminal TTY concentrator.

2.11 BladeCenter HT interposers

Because of the design features of the IBM BladeCenter HT chassis, interposers must be used when installing any I/O module or redundant management module. These interposers basically extend the connections from the midplane to the rear of the module. Possible options are listed in Table 2-62.

Table 2-62 Interposers for IBM BladeCenter HT

Part number	Description	Photo
42C5300	IBM BladeCenter HT Interposer for Gb Switch and Bridge Bays	
42C5301	IBM BladeCenter HT Interposer for Gb Switch and Bridge Bays with Interswitch Links (ISL)	Same as above
42C5302	IBM BladeCenter HT Interposer for HS Switch Bay	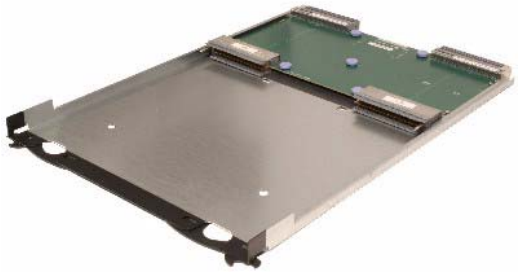
42C5315	IBM BladeCenter HT Advanced Management Module Interposer	

2.12 BladeCenter S Office Enablement Kit

The BladeCenter S Office Enablement Kit, part number 201886X, is an enclosure specifically designed for the BladeCenter S chassis for use in offices without a dedicated server room, or where the dust level is high. The enclosure, with the BladeCenter S chassis and the Flat Panel Monitor kit installed, is shown in Figure 2-60 on page 210.

Figure 2-60 BladeCenter S Office Enablement Kit with BladeCenter S and Flat Panel Monitor kit

Based on the NetBAY11, the Office Enablement Kit is an 11U enclosure with security doors and special acoustics and air filtration to suit office environments. With the BladeCenter S chassis installed, this leaves an extra 4U of space to hold other rack devices.

The Office Enablement Kit features:

- An acoustical module

  The Office Enablement Kit comes with an acoustical module that helps make BladeCenter S quiet for the office environment, while allowing easy access to BladeCenter S components.

- A locking door

  Security is important in any office environment. The Office Enablement Kit comes with a front locking door that helps keep your data safe and secure.

- 4U of extra space for other devices

  Different businesses use different tools to enable their office IT. The Office Enablement Kit includes 4U of extra space for other types of IT that an office might need. This space can take any IT that fits into a 4U or smaller standard rack space.

- Mobility

  The Office Enablement Kit comes with lockable wheels to make your BladeCenter S easily transportable.

The Office Enablement Kit also supports an optional Air Contaminant Filter to assist BladeCenter S functionality when deployed in dusty environments. IBM BladeCenter Airborne Contaminant Filter, part number 43X0340, feature code 4024, is an optional hardware kit that enables the Office Enablement Kit to use air filters, and one air filter is included. In addition, replacement air filters can be ordered in quantities of four (IBM BladeCenter Airborne Contaminant Replacement Filter (4-Pack), part number 43X0437, feature code 4025).

The enclosure has the following approximate dimensions: height 24 inches, width 24 inches, and depth 42 inches.

2.13 BladeCenter S storage modules

The BladeCenter S supports up to two storage modules, as shown in Figure 1-17 on page 31. These modules provide integrated SAS storage functionality to the BladeCenter S chassis. No storage modules are standard with the BladeCenter S chassis.

The storage module is fundamentally a collection of disk drives, which are made accessible to blade servers through a SAS I/O module (see 2.5, “SAS I/O modules” on page 164) and SAS expansion cards installed in the blades (see 3.1.17, “SAS Expansion Card” on page 252).

Each of the two storage module contains up to six 3.5 inch hot-swap hard drives, for a total of 12 internal drives. The storage module supports SAS, SATA, and Near Line SAS (NL SAS) drives. Intermixing SAS and SATA or SAS and NL SAS drives within the same storage module is supported.

Note: SAS Connectivity Module (39Y9195, feature code 2980) only supports SAS and SATA drives. SAS RAID Controller Module (43W3584, feature code 3734) only supports SAS and NL SAS drives.

The IBM BladeCenter S 6-Disk Storage Module, part number 43W3581, feature code 4545, is shown in Figure 2-61.

Figure 2-61 BladeCenter S 6-disk Storage Module

The BladeCenter S chassis can accommodate up to six 3.5 inch disk drives in each storage module. You can easily and quickly assign the drives directly to blades using built-in predefined configurations or through user-defined custom configurations.

Table 2-63 lists hard disk drives supported by BladeCenter S disk storage modules.

Table 2-63 Disk drives supported by BladeCenter S disk storage module

Description	Part Number	Feature Code
73 GB 15K 3.5” SAS Hot-swap HDD	43W7523	3748
146 GB 15K 3.5” SAS Hot-swap HDD	43W7524	3749
300 GB 15K 3.5” SAS Hot-swap HDD	43X0802	3747
IBM 450 GB 15K SAS 3.5” HS HDD	42D0519	5586

Description	Part Number	Feature Code
500 GB SATA2 drive	39M4558	3743
750 GB Dual Port SATA Hot Swap Option	43W7580	3746
IBM 1 TB 7200 Dual Port SATA 3.5” HS HDD	43W7630	5561
IBM 750GB 7200 NL SAS 3.5” HS HDD	42D0546	1084
IBM 1 TB 7200 NL SAS 3.5” HS HDD[a]	42D0547	5581

a. Support for this HDD is planned for February 2009.

Refer to 5.10, “IBM BladeCenter S integrated storage” on page 423 for more information about BladeCenter S integrated storage planning.

2.14 Installation and physical site plans

The IBM BladeCenter E chassis, the BladeCenter T chassis, the BladeCenter H chassis, the BladeCenter HT chassis, and all blade servers are designated as customer setups. Customer setup instructions are shipped with the systems or can be downloaded from IBM Support Web site. Refer to product publications for installation instructions. These product publications can be downloaded in PDF format from the BladeCenter Support site:

http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/brandmain?brandind=5000020

Topics in this section include:

2.14.1 Physical planning

This section includes the physical planning specifications and weight considerations.

BladeCenter E physical specifications

An IBM BladeCenter E type 8677 unit has the following specifications:

- Height: 304 mm (12 in. or 7 U)
- Depth: 711 mm (28 in.)
- Width: 444 mm (17.5 in.)
- Weight:
  - Fully configured with modules and blades: Approximately 108.86 kg (240 lb.)
  - Fully configured without blades: Approximately 44.91 kg (99 lb.)
- A BladeCenter unit uses 7U of vertical rack space.
- Acoustical noise emissions:
  - Without acoustics module option:
    - Sound power, idle: 7.4 bels maximum
    - Sound power, operating: 7.4 bels maximum
  - With acoustics module option:
    - Sound power, idle: 6.9 bels maximum
    - Sound power, operating: 6.9 bels maximum
- Electrical input:
  - Sine-wave input (50-60 Hz single-phase) required
  - Input voltage:
    - Minimum: 200 V ac
    - Maximum: 240 V ac

BladeCenter T physical specifications

An IBM BladeCenter T type 8720 unit has the following specifications:

- Height: 349 mm (13.75 in. or 8 U)
- Depth: 508 mm (20 in.) from front of chassis to rear I/O connector plane. Maximum depth: 600 mm (23.62 in.) including bezel, handles, and cable bend radius

- Width: 442 mm (17.4 in.)
- Weight:
  - Fully configured with modules and blade servers: Approximately 86.64 kg (191 lb.)
  - Fully configured without blade servers: Approximately 44.45 kg (98 lb.)
- A BladeCenter T unit uses 8U of vertical rack space
- Declared acoustical noise emission levels for normal operations:
  - Sound-power levels (upper-limit): 7.8 bels
  - Sound-pressure levels (average), for four one-meter bystander positions: 63 dBA
- Electrical input:
  - dc power
  - Input voltage: -48 V dc (-40 V dc to -60 V dc)

An IBM BladeCenter T type 8730 unit has the following specifications:

- Height: 349 mm (13.75 in. or 8 U)
- Depth: 508 mm (20 in.) from front of chassis to rear I/O connector plane. Maximum depth: 600 mm (23.62 in.) including bezel, handles, and cable bend radius
- Width: 442 mm (17.4 in.)
- Weight:
  - Fully configured with modules and blade servers: Approximately 86.64 kg (191 lb.)
  - Fully configured without blade servers: Approximately 44.45 kg (98 lb.)
- A BladeCenter T unit uses 8U of vertical rack space
- Declared acoustical noise emission levels for normal operations:
  - Sound-power levels (upper-limit): 7.8 bels
  - Sound-pressure levels (average), for four one-meter bystander positions: 63 dBA
- Electrical input:
  - Sine-wave input (50 or 60 Hz single-phase) required
  - Input voltage:
    - Minimum: 200 V ac
    - Maximum: 240 V ac

BladeCenter H physical specifications

An IBM BladeCenter H unit has the following dimensions:

- Height: 400 mm (15.75 in. or 9 U)
- Depth: 711 mm (28 in.)
- Width: 482 mm (19 in.)
- Weight:
  - Full configured weight with blade servers: Approximately 158.8 kg (350 lb.)
  - Empty chassis without modules or blade servers: Approximately 40.82 kg (90 lb.)
- Acoustics:
  - Declared sound power level: 7.5 bels

- Electrical input:
  - Sine-wave input (50-60 Hz single-phase) required
  - Input voltage:
    - Minimum: 200 V ac
    - Maximum: 240 V ac

BladeCenter HT physical specifications

An IBM BladeCenter HT type 8740 unit has the following specifications:

- Height: 533.4 mm (21 in. or 12 U).
- Depth:
  - Without optional bezel: 617 mm (24.29 in.)
  - With optional bezel: 706 mm (27.80 in.) bezel, handles, and cable bend radius
- Width: 442 mm (17.4 in.).
- Weight:
  - Fully configured with modules and blade servers: approximately 158.8 kg (350 lb.)
  - Fully configured without blade servers: approximately 65.32 kg (144 lb.)
- A BladeCenter HT unit uses 12 U of vertical rack space.
- Declared acoustical noise emission levels for normal operations: Sound-power levels (upper-limit): 7.8 bels.
- Electrical input:
  - DC isolated
  - Four inputs at 60 A rating each
  - Input voltage: -40 V dc to -72 V dc

An IBM BladeCenter HT type 8750 unit has the following specifications:

- Height: 533.4 mm (21 in. or 12 U)
- Depth:
  - Without optional bezel: 617 mm (24.29 in.)
  - With optional bezel: 706 mm (27.80 in.) bezel, handles, and cable bend radius
- Width: 442 mm (17.4 in.)
- Weight:
  - Fully configured with modules and blade servers: approximately 158.8 kg (350 lb.)
  - Fully configured without blade servers: approximately 65.32 kg (144 lb.)
- A BladeCenter HT unit uses 12 U of vertical rack space.
- Declared acoustical noise emission levels for normal operations: sound-power levels (upper-limit): 7.8 bels.
- Electrical input:
  - Four inputs at 16 A rating each
  - Input voltage: 180 V ac to 265 V ac
  - Sine-wave input (50/60 Hz single-phase)

BladeCenter S physical specifications

An IBM BladeCenter S unit has the following specifications:

- Height: 306 mm (12 in. or 7 U)
- Depth: 733 mm (28.9 in.)
- Width: 444 mm (17.5 in.)
- Weight:
  - Minimum configuration: 40.8 kg (90 lb)
  - Maximum configuration: 108.8 kg (240 lb)
- A BladeCenter unit uses 7U of vertical rack space.
- Electrical input:
  - Sine-wave input (50-60 Hz single-phase) required
  - Input voltage:
    - 200 to 240 (nominal) V ac; 50 Hz or 60 Hz
    - 110 to 127 (nominal) V ac; 50 Hz or 60 Hz
- Power consumption
  - Minimum configuration: 0.40 kVA (two power supplies)
  - Maximum configuration: 3.50 kVA (four power supplies)
  - Leakage current: 3 mA
- BTU heat output:
  - Ship configuration — 1365 Btu/hr (400 watts)
  - Full configuration — 11942 Btu/hr (3500 watts)
- Acoustical noise emissions for BladeCenter with six blade servers:
  - 63-69 dB blade dependent (operating)
  - 63-69 dB blade dependent (idling)

Weight considerations

The following factors are important when planning for the physical installation of BladeCenter units:

- Racks have a maximum weight load limit.
- Floors, particularly raised floors, have maximum weight limits.

Specific requirements are available in the following documentation:

- BladeCenter type 8677 - *Planning and Installation Guide*
- BladeCenter T types 8720 and 8730 - *Planning and Installation Guide*
- BladeCenter H type 8852 - *Installation and User Guide*
- BladeCenter HT types 8740 and 8750 - *Installation and User Guide*
- IBM BladeCenter S Type 8886 - *Installation and User's Guide*

You can find this documentation under publications in the relevant product section in the BladeCenter Support site:

http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/brandmain?brandind=5000020

2.14.2 Rack installation

IBM intends for you to install and secure the BladeCenter chassis in a rack. The IBM System x Rack Configurator is no longer available as a separate tool. Its functionality is incorporated into IBM Standalone Solution Configuration Tool (SSCT). You can download the IBM SSCT from:

http://www.ibm.com/systems/x/configtools.html

When planning for racks, you must consider the following factors:

- Floor space that provides adequate access to the front and rear of the rack, and space on both sides of the rack or a row of racks for ventilation and walk through space.
- Rack weight load limits and floor weight load limits.
- Rack height – the chassis require the following rack space:
  - BladeCenter E: 7U
  - BladeCenter T: 8U
  - BladeCenter H: 9U
  - BladeCenter HT: 12U
  - BladeCenter S: 7U

  Also consider the space needed for other equipment that might be installed in the same rack with the BladeCenter units.
- The chassis unit can only be installed horizontally.
- Three or more people are required to install the device in a rack cabinet, after all the blade servers are removed.
- Do not leave any unused U space within a rack cabinet open. Blank filler panels must be used to prevent recirculation of warm air.
- Install your BladeCenter chassis only in a rack cabinet with perforated front and rear doors.
- Install your BladeCenter E or BladeCenter H in a rack that meets the following requirements:
  - Minimum depth of 70 mm (2.76 in.) between the front mounting flange and inside of the front door.
  - Minimum depth of 157 mm (6.18 in.) between the rear mounting flange and inside of the rear door.
  - Minimum depth of 559 mm (22 in.) and maximum depth of 762 mm (30 in.) between the front and rear mounting flanges.
- IBM BladeCenter T can be installed into standard rack, 2-post rack for Telco, 4-post rack for Telco, or UTF (Universal Telco Frame) rack.
- IBM BladeCenter HT can be installed into a 2-post or a 4-post rack.
- The BladeCenter HT unit is designed to be rack-compatible with NEBS-compliant 2-post racks that have size 12–24 tapped holes.

Specific requirements are available in the following documentation:

- BladeCenter type 8677 - *Rack Installation Instructions*
- BladeCenter T types 8720 and 8730 - *Rack Installation Instructions*
- BladeCenter H type 8852 - *Rack Installation Instructions*
- BladeCenter HT types 8740 and 8750 - *2-post and 4-post Installation instructions*
- IBM BladeCenter S Type 8886 *Planning Guide*

You can find the documentation under publications in the relevant product section at:

http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/brandmain?brandind=5000020

Note: IBM has tested and offers rack kits for NEBS and ETSI Seismic racks. Do not place racks directly against a wall because, if service is needed, the service representative will need access to the back of the BladeCenter T unit.

The following components are not included and must be supplied separately:

- Mouse
- Keyboard
- Display
- Network cables

2.14.3 Power considerations

Different IBM BladeCenter chassis have different power topologies inside, as discussed in 2.9, “Power modules” on page 195. You will find additional information in this section.

Note: The BladeCenter chassis type 8677, BladeCenter T chassis type 8730 (AC), BladeCenter H chassis type 8852, and BladeCenter HT chassis type 8750 (AC) can only be powered with 220VAC 50/60 Hz worldwide. Your data center must be able to supply this voltage.

The 48 VDC-powered BladeCenter T chassis type 8720 and BladeCenter HT chassis type 8740 must be installed by personnel trained in DC wiring installations. Setup instructions are shipped with the systems.

IBM BladeCenter power management

IBM BladeCenter hardware architecture provides energy efficiency through its design. In addition to this, enterprise-level power management features allow you to control electrical power consumption. IBM Active Energy Manager software is the key component in the BladeCenter power management solution.

The BladeCenter E, BladeCenter T, BladeCenter H, and BladeCenter HT chassis each have two power domains, with a specific set of components powered. Each chassis is also powered by two redundant power supplies. The BladeCenter S chassis has only one power domain, containing up to four power supplies. Specific considerations about power domains can be found in 2.9, “Power modules” on page 195.

As a part of the Active Energy Manager solution, the BladeCenter power management policies have been implemented for better distribution and control of power available to blades on a chassis level. All BladeCenter chassis have the same set of power management policies, and BladeCenter S has two additional policies.

The power management policies are:

- Redundant with potential performance impact (“AC power source redundancy with blade throttling allowed” for BladeCenter S)
- Redundant without performance impact (“AC power redundancy” for BladeCenter S)
- Non-redundant
- Power module redundancy (BladeCenter S only)
- Power module redundancy with blade throttling allowed (BladeCenter S only)

These policies control how much electrical power is available to power domain components, as well as what happens in case of a power source or power supply failure in this domain.

One possible response to loss of power redundancy could be throttling. *Throttling* means that blades reduce their power consumption by decreasing the processor frequency, and as soon as power is restored, the blades will return to their full frequency operations. For fast response time, this process is controlled by the service processor on the blade, and blades must support throttling. This technique is part of the IBM Active Energy Management initiative. Power policies with a performance impact achieve power management through throttling when needed.

In general, “non-redundant” means that maximum power is available for the domain, and power supply failure does not necessarily mean that blades will be powered off. If throttling can accommodate the capacity of the remaining power supplies in the domain, then blades continue their operations with reduced processor frequency.

“Redundant without performance impact (AC power redundancy)” means N+N redundancy for the domain. It also means that the power capacity available is equal to the capacity of one or two power supplies, depending on the number of power supplies in one domain (two or four). For example, if a power domain has two power supplies (applicable to all chassis), then maximum power capacity is equal to the capacity of one power supply. If a power domain has four power supplies (applicable to the BladeCenter S only), then maximum capacity is the sum of capacities for two power supplies. No throttling occurs in this scenario. If the blade exceeds the power capacity of the domain, it will not be powered up.

“Redundant with potential performance impact (AC power source redundancy with blade throttling allowed)” is almost the same as the preceding policy, except that the power capacity available is higher and throttling is possible.

BladeCenter S has two additional policies that provide N+1 redundancy for four power supplies in the same manner as described. When “Power module redundancy” is selected, then maximum power capacity is equal to the sum of capacities for three power supplies, and no throttling occurs. “Power module redundancy with blade throttling allowed” means that maximum power capacity is higher than in the previous case, and throttling is possible.

IBM System x and the BladeCenter Power Configurator tool, together with the power policy selected, can be used to estimate the possibility of throttling based on a specific configuration.

The IBM BladeCenter Power Configurator tool can be used for the estimation of power consumption for a specific configuration. It shows power consumption for each power domain and overall power consumption, as well as other characteristics. This tool can be downloaded from:

http://www.ibm.com/systems/bladecenter/powerconfig/

Note: If throttling is absolutely unacceptable, then you must fit the configuration into the available power envelope, either by limiting the number of blades installed into the chassis or by selecting blade models that fall within the power limits necessary for operation in the “Redundant without performance impact” power policy.

Additional power considerations for specific chassis are covered later in the document.

For a detailed discussion of power management in the IBM BladeCenter, refer to *BladeCenter packaging, power, and cooling* in the IBM Journal of Research and Development, available at:

http://www.research.ibm.com/journal/rd49-6.html

IBM BladeCenter E specific considerations

IBM BladeCenter E chassis shipped earlier might have different types of power supplies: 1200W, 1400W, 1800W, or 2000W. Only 2000W power supplies are supported with blade servers shipping currently, so if you plan to install new blades into existing BladeCenter E chassis, you must plan to upgrade your power supplies to 2000W in the power domains where you plan to install these new blades. You must install the power modules in pairs in a domain, and they must match each other in capacity (wattage, amperage, and so on).

Only 2000W power supplies are available now. All others were withdrawn from market. See 2.9.1, “Power modules for BladeCenter E” on page 196, for more information about this topic.

To provide true redundant power, BladeCenter E power modules 1 and 3 must be connected to a different 200-240 Volt AC power source than power modules 2 and 4. An installed power module must be connected to an AC power source and must not be used as a filler.

In any configuration, the total power consumption for the single IBM BladeCenter E chassis must not exceed 5400W.

IBM BladeCenter E should be attached to high-voltage PDUs, as described in 2.14.5, “Cabling considerations” on page 225.

Refer to 2.9.1, “Power modules for BladeCenter E” on page 196, for additional information.

IBM BladeCenter H- specific considerations

BladeCenter H chassis comes with two 2900W power supplies. Additional 2900W power supplies for power domain B are available as an option as described in 2.9.2, “Power modules for BladeCenter H” on page 197.

Remember that power cables for BladeCenter H are not shipped with the chassis or with optional power supplies, and must be purchased separately. IBM BladeCenter H has only two specific power connectors, which must be connected to a different 200-240 Volt AC power sources to provide true redundancy.

In any configuration, the total power consumption for the single IBM BladeCenter H chassis must not exceed 8000W.

IBM BladeCenter H should be attached to high-voltage PDUs, as described in 2.14.5, “Cabling considerations” on page 225.

See 2.9.2, “Power modules for BladeCenter H” on page 197, for additional information.

IBM BladeCenter T specific considerations

IBM BladeCenter T ships with two 1300W AC (8730) or DC (8720) power supplies. Additional 1300W power supplies for power domain B are available as an option as described in 2.9.3, “Power modules for BladeCenter T” on page 200.

To provide true redundant power, BladeCenter T power modules 1 and 3 must be connected to a different power source than power modules 2 and 4.

In any configuration, the total power consumption for the single IBM BladeCenter T chassis must not exceed 3000W.

IBM BladeCenter T 8730 should be attached to high-voltage PDUs as described in 2.14.5, “Cabling considerations” on page 225.

IBM BladeCenter T 8720 should be attached to the appropriate DC wiring as described in 2.14.5, “Cabling considerations” on page 225.

See 2.9.3, “Power modules for BladeCenter T” on page 200, for additional information.

IBM BladeCenter HT specific considerations

IBM BladeCenter HT ships with two AC (8750) or DC (8740) power supplies. Additional AC or DC power supplies for power domain B are available as an option, as described in 2.9.4, “Power Modules for BladeCenter HT” on page 202.

To provide true redundant power, BladeCenter HT power modules 1 and 3 must be connected to a different power source than power modules 2 and 4.

IBM BladeCenter HT 8750 should be attached to high-voltage PDUs, as described in 2.14.5, “Cabling considerations” on page 225.

IBM BladeCenter HT 8740 should be attached to the appropriate DC wiring, as described in 2.14.5, “Cabling considerations” on page 225.

See 2.9.4, “Power Modules for BladeCenter HT” on page 202, for additional information.

IBM BladeCenter S specific considerations

Unlike the other BladeCenter chassis, the BladeCenter S can operate from a 110V AC power source as well as 220V. Also unlike the other BladeCenter chassis, within the BladeCenter S chassis all power supplies are combined into a single power domain, which distributes power to each of the blade servers and modules through the system midplane.

IBM BladeCenter S ships with two 950 W (110 V)/1450 W (220 V) AC power supplies. An additional power supplies option is available, as described in 2.9.5, “Power modules for BladeCenter S” on page 204.

BladeCenter S can be attached to a supported PDU or to office wall jacks, as described in 2.14.5, “Cabling considerations” on page 225.

See 2.9.5, “Power modules for BladeCenter S” on page 204, for additional information.

2.14.4 Cooling considerations

When planning a data center, take the following points into consideration regarding cooling.

Air temperature

The permissible temperatures and humidity limits for different BladeCenter chassis are:

- IBM BladeCenter E
  - On: 10.0 to 35.0 degrees C (50 to 95 degrees F) at 0 to 914 m (0 to 3000 ft)
  - On: 10.0 to 32.0 degrees C (50 to 90 degrees F) at 914 to 2133 m (3000 to 7000 ft)
  - Off: -40 to 60 degrees C (-40 to 140 degrees F)

- Relative humidity:
  - On: 8% to 80%
  - Off: 8% to 80%
- Maximum altitude: 2133 m (7000 ft)

- IBM BladeCenter H
  - On: 10.0 to 35.0 degrees C (50 to 95 degrees F) at 0 to 914 m (0 to 3000 ft)
  - On: 10.0 to 32.0 degrees C (50 to 90 degrees F) at 914 to 2,133 m (3000 to 7000 ft)
  - Off: -40 to 60 degrees C (-40 to 140 degrees F)
  - Relative humidity:
    - On: 8% to 80%
    - Off: 8% to 80%
  - Maximum altitude: 2,133 m (7000 ft)
- IBM BladeCenter T 8720 (NEBS/Telecom)
  - On: 5 to 40 degrees C (41 to 104 degrees F) Altitude: -60 m (-197 ft) to 1800 m (6000 ft)
  - On (short term): -5 to 55 degrees C (23 to 131 degrees F) Altitude: -60 m (-197 ft) to 1800 m (6000 ft)
  - On: 5 to 30 degrees C (41 to 86 degrees F) Altitude: 1800 m (600 ft) to 4000 m (13,000 ft)
  - On (short term): -5 to 45 degrees C (23 to 113 degrees F) Altitude: 1800 m (6000 ft) to 4000 m (13,000 ft)
  - Off: -40 to 70 degrees C (-40 to 158 degrees F)
  - Rate of temperature change 30 degrees C/hr (54 degrees F/hr)
  - Humidity:
    - On: 5% to 85%
    - On (short term): 5% to 90% but not to exceed 0.024 kg water/kg of dry air
    - Off: uncontrolled

    Short term is a period of not more than 96 consecutive hours and a total of not more than 15 days in one year. (This refers to a total of 360 hours in any given year but no more than 15 occurrences during that one-year period.)
  - Maximum altitude: 4000 m (13 000 ft)
- IBM BladeCenter T 8730 (non-NEBS)
  - On: 5.0 to 35.0 degrees C (41 to 95 degrees F) at -60 to 1800 m (-197 to 6,000 ft)
  - On: 5.0 to 32.0 degrees C (41 to 90 degrees F) at 1800 to 4000 m (6000 to 13 000 ft)
  - Off: uncontrolled
  - Relative humidity:
    - On: 5% to 85%
    - Off: uncontrolled
  - Maximum altitude: 4,000 m (13,000 ft)
- IBM BladeCenter HT 8740 and 8750
  - On: 5° to 40°C (41° to 104°F) at altitude: -60 to 1800 m (-197 ft to 6000 ft)
  - On: 5° to 30°C (41° to 86°F) at altitude: 1800 m to 4000 m (6000 ft to 13 000 ft)
  - Off: -40° to 70°C (-40° to 158°F)

- Relative humidity: 5% to 85%

- IBM BladeCenter S
  - Temperature:
    - 10° to 35°C (50° to 95°F) at 0 to 914 m (0 to 3000 ft)
    - 10° to 32°C (50° to 90°F) at 914 to 2,133 m (3000 to 7000 ft)
  - Relative humidity: 8% to 80%
  - Maximum altitude: 2133 m (7000 ft)

Air flow

Air flow is critical for ensuring that the operating air temperature stays within permissible limits. The blowers used in each chassis type are described in 2.8, “Blower modules” on page 191.

- Air flow direction is from front to back.
- All BladeCenter chassis bays must be populated with a module, a blade server, or a filler in place of the component. If a location is left empty for more than one minute while the BladeCenter unit is operating, performance degradation or thermal failures might occur.
- All equipment installed in a rack with a BladeCenter unit must use front-to-back air flow to prevent warm air recirculation problems. Devices that use back-to-front air flow cause warm air to enter the front of the BladeCenter unit. This can result in reduced reliability, component failure, data loss, or server shutdown.
- In racks with multiple BladeCenter units, populate the BladeCenter chassis starting with the bottom chassis in the rack and working up towards the top of the rack.
- Any unused rack space must be covered with a blank rack filler panel to ensure proper air circulation.
- The BladeCenter T unit comes with a bezel assembly containing a removable and replaceable air filter. There are software features in the management module that detect a clogged filter and generate system alerts based on the severity of the airflow reduction. The typical service interval for the filter is approximately three to six months, depending on the environment.
- The IBM BladeCenter HT has an option of using air filtration (an optional bezel is required). There are software features in the management module that detect a clogged filter and generate system alerts based on the severity of the airflow reduction. The typical service interval for the filter is approximately three to six months, depending on the environment.

Heat output

The amount of heat output of BladeCenter chassis in BTU per hour is as listed in Table 2-64.

Table 2-64 Heat output for IBM BladeCenter chassis configurations

Chassis	Heat output for maximum configuration
IBM BladeCenter E	18425 BTU/hour (5400 W)
IBM BladeCenter H	27280 BTU/hour (8000 W)
IBM BladeCenter T 8720	11229 BTU/hour (3291 W)
IBM BladeCenter T 8730	10440 BTU/hour (3060 W)
IBM BladeCenter HT 8740	19680 BTU/hour (5766 W)

Chassis	Heat output for maximum configuration
IBM BladeCenter HT 8750	21850 BTU/hour (6400 W)
IBM BladeCenter S	11942 BTU/hr (3500 W)

Prevention of air recirculation

Consider these factors when planning for single or multiple rack installations:

- When racks are positioned adjacent to each other, ensure that the racks fit tightly together from side to side to prevent inter-rack air recirculation from the back to the front.
- Air recirculation occurs over the top or around the side of a rack in a room that does not have a cooling system with sufficient airflow volume capacity. Ensure that the cooling system has adequate capacity for the room cooling load.

Room cooling

To prevent possible BladeCenter thermal failures proper room cooling is vital, keep these points in mind when planning for the installation:

- Ensure that the site cooling system has adequate capacity for the room cooling load.
- Ensure that cool air is provided to the front of the BladeCenter unit and rack.
- Ensure that the room cooling system is positioned so warm exhaust air is directed away from all BladeCenter units towards the room cooling system without passing in front of a BladeCenter unit.
- A significant air temperature gradient can occur from the bottom to the top of a rack in a room which has a cooling system that does not have sufficient airflow volume and cooling capacity. This might cause equipment at the top of the rack to run hot resulting in reduced reliability, component failure, data loss, or server shutdown.

Rear Door Heat eXchanger

For data centers that have limited cooling capacity, using the Rear Door Heat eXchanger (see Figure 2-62 on page 225) is a more cost-effective solution than adding another air conditioning unit.

Note: The Rear Door Heat eXchanger is not a requirement for BladeCenter cooling. It is a solution for clients who cannot upgrade a data center room's air conditioning units due to space, budget, or other constraints.

The Rear Door Heat eXchanger (part number 32R0712) has the following features:

- A water-cooled heat exchanger door is designed to dissipate heat generated from the back of the computer systems before it enters the room.
- An easy-to-mount rear door design attaches to client-supplied water, using industry standard fittings and couplings.
- Up to 50,000 BTUs (or approximately 15 kW) of heat can be removed from air exiting the back of a rack full of servers.
- It fits to an IBM Enterprise 42U rack (930842S, 930842E, 141042X).
- It offers a 1-year limited warranty.

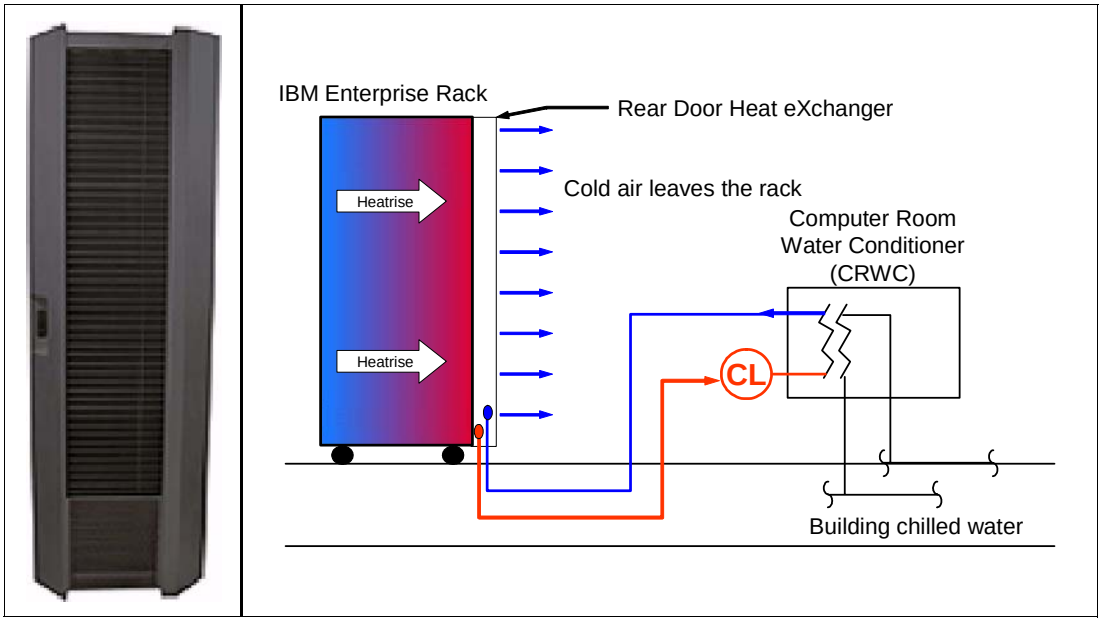

Figure 2-62 Rear Door Heat eXchanger (left) and functional diagram

The IBM Rear Door Heat eXchanger is an effective way to assist your air conditioning system in keeping your data center cool. It removes heat from the rack before the heat enters the room, allowing your air conditioning unit to handle the increasingly dense system deployment your organization requires to meet its growing computing needs. The IBM Rear Door Heat eXchanger also offers an convenient way to handle dangerous “hot spots” in your data center.

2.14.5 Cabling considerations

This section describes specific considerations related to console and power cabling.

Power cabling

The IBM BladeCenter chassis all have different power supplies, and therefore different cabling requirements. However, they must be connected to supported power distribution units (PDU).

Because power requirements are country-specific, refer to the available documentation. The available PDUs allow for different types of power connections worldwide, making it possible to support six fully configured BladeCenter chassis in one 42U Enterprise Rack with a total of four power connections to the data center, either as one-phase or three-phase connections with different ampere loads.

Note: Line cords are not included for countries in the EMEA (Europe, Middle East, Africa) region.

Some general rules can be followed for choosing correct PDUs and power cables:

- IBM BladeCenter E (8677)
  - Each IBM BladeCenter unit can have two or four C20 power connectors (on power supplies) depending on the number of 2000W power supplies installed.
  - The BladeCenter unit comes standard with two IEC 320 C19 - C20 power cables, and the power supplies option contains two additional IEC 320 C19 - C20 power cables.
  - Connect power supplies in the same domain to different power sources to provide true redundancy.
  - BladeCenter E must be attached to the PDUs that provide a sufficient number of C19 connectors, power capacity, voltage and amperage.
- IBM BladeCenter H (8852)
  - Each IBM BladeCenter H unit has two specific power connectors on the chassis.
  - The BladeCenter unit does not have any power cables included into the ship group; they must be ordered separately (see Table 2-53 on page 199 for part numbers and photos of the required cables).
  - Two power cables are required per single BladeCenter H unit. Connect these cables to different power sources to provide true redundancy.
  - IBM BladeCenter H triple IEC 320 C20 power cables can be used in most cases, so each BladeCenter H unit requires six C19 power connectors (three per PDUs connected to different power sources).
  - The PDUs selected should provide sufficient power capacity, voltage and amperage.
- IBM BladeCenter T (8720)

  The BladeCenter T 8720 unit is connected to the DC power circuit. This connection should performed by trained personnel.
- IBM BladeCenter T (8730)
  - Each BladeCenter T 8730 unit has four C14 power connectors on the chassis.
  - The BladeCenter T unit comes standard with two IEC 320 C13 - C14 power cables, and the power supplies option contains two additional IEC 320 C13 - C14 power cables.
  - Connect power supplies in the same domain to different power sources to provide true redundancy.
  - BladeCenter T must be attached to the PDUs that provide a sufficient number of C13 connectors, power capacity, voltage and amperage.
- IBM BladeCenter HT (8740)

  The BladeCenter HT 8740 unit is connected to the DC power circuit. This connection should be performed by trained personnel.
- IBM BladeCenter HT (8750)
  - Each BladeCenter HT 8750 unit has four C20 power connectors on the chassis.
  - The BladeCenter HT unit comes standard with two IEC 320 C19 - C20 power cables, and the power supplies option contains two additional IEC 320 C19 - C20 power cables.
  - Connect power supplies in the same domain to different power sources to provide true redundancy.
  - The BladeCenter HT must be attached to the PDUs that provide a sufficient number of C19 connectors, power capacity, voltage, and amperage.

- IBM BladeCenter S
  - Each IBM BladeCenter S chassis requires either two or four power cables, each with IEC 320-C20 power connectors, when connected to each installed power module.
  - The BladeCenter S unit comes standard with two IEC 320 C19 - C20 power cables, and four country-specific line cords. The additional power supplies option contains two additional IEC 320 C19 - C20 power cables.
  - The power modules are auto-sensing and can support either 110 V or 220 V AC power. However, do *not* mix voltage power sources within the same BladeCenter S chassis.

The latest information about available PDU products and supported cabling is in the *Configuration and Options Guide* which is available at:

`http://www.ibm.com/systems/x/configtools.html`

We also strongly recommend that you read the information available in:

- *BladeCenter type 8677 - Planning and Installation Guide*
- *BladeCenter T types 8720 and 8730 - Planning and Installation Guide*
- *BladeCenter H type 8852 - Installation and User Guide*
- *BladeCenter HT types 8740 and 8750 - Installation and User Guide*
- *BladeCenter S type 8886 Planning and Installation Guide*

These documents cover the following details:

- Physical planning
- Rack consideration
- Power consideration
- Cooling consideration
- Power connections (PDUs and cables)
- Physical installation time

You can find this documentation under publications in the relevant product section (for example, BladeCenter 8677 or BladeCenter T 8720 or 8730, or BladeCenter H 8852, or BladeCenter HT 8740 or 8750) at:

`http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/brandmain?brandind=5000020`

Refer to 2.9, “Power modules” on page 195, for additional information.

Console cabling

The IBM BladeCenter management module acts as a KVM switch for blade servers installed into the chassis. The console (keyboard, video, mouse) can be connected to the management module directly, or through the KVM switch.

The standard management module for BladeCenter has PS/2 connectors for keyboard and mouse. The advanced management module for IBM BladeCenter E, BladeCenter H, BladeCenter HT, and BladeCenter S has USB ports for keyboard and mouse. BladeCenter T uses PS/2 ports for keyboard and mouse for both standard and advanced management modules.

Ensure you select the right KVM cables to connect BladeCenter to the console switches. Depending on the KVM switch and management module used, select the corresponding KVM cables. Supported combinations are listed in Table 2-65.

Table 2-65 Supported console switches and cables

	IBM 1x8 Console Switch (17351LX) **IBM 2x16 Console Switch (17352LX)**	**IBM Local 2x8 Console Manager (17351GX)** **IBM Global 2x16 Console Manager (17352GX)** **IBM Global 4x16 Console Manager (17354GX)**
IBM BladeCenter with MM	31R3130 3M Console Switch Cable (PS/2)	39M2896 250 mm KVM Conversion Option 39M2897 1.5 M KVM Conversion Option
IBM BladeCenter with AMM	31R3132 3M Console Switch Cable (USB)	39M2895 IBM USB Conversion Option
IBM BladeCenter H with AMM	31R3132 3M Console Switch Cable (USB)	39M2895 IBM USB Conversion Option
IBM BladeCenter T with MM	31R3130 3M Console Switch Cable (PS/2)	39M2896 250mm KVM Conversion Option 39M2897 1.5 M KVM Conversion Option
IBM BladeCenter T with AMM	31R3130 3M Console Switch Cable (PS/2)	39M2896 250 mm KVM Conversion Option 39M2897 1.5 M KVM Conversion Option
IBM BladeCenter HT with AMM	31R3132 3M Console Switch Cable (USB)	39M2895 IBM USB Conversion Option

Providing a SOL console for the JS20/JS21 blades

The BladeCenter JS20 differs from other blade servers that are available for the BladeCenter in that it does not provide an interface to a keyboard, video monitor, and mouse (KVM) console. The JS21 provides a keyboard and video starting with the GA3 firmware update that was made available in June 2007.

The JS20 and JS21 support Serial over LAN (SOL) remote text console functions. Use of SOL is optional for other types of blade servers that support KVM consoles.

The SOL remote text console function involves several components in the BladeCenter infrastructure as illustrated in Figure 2-63.

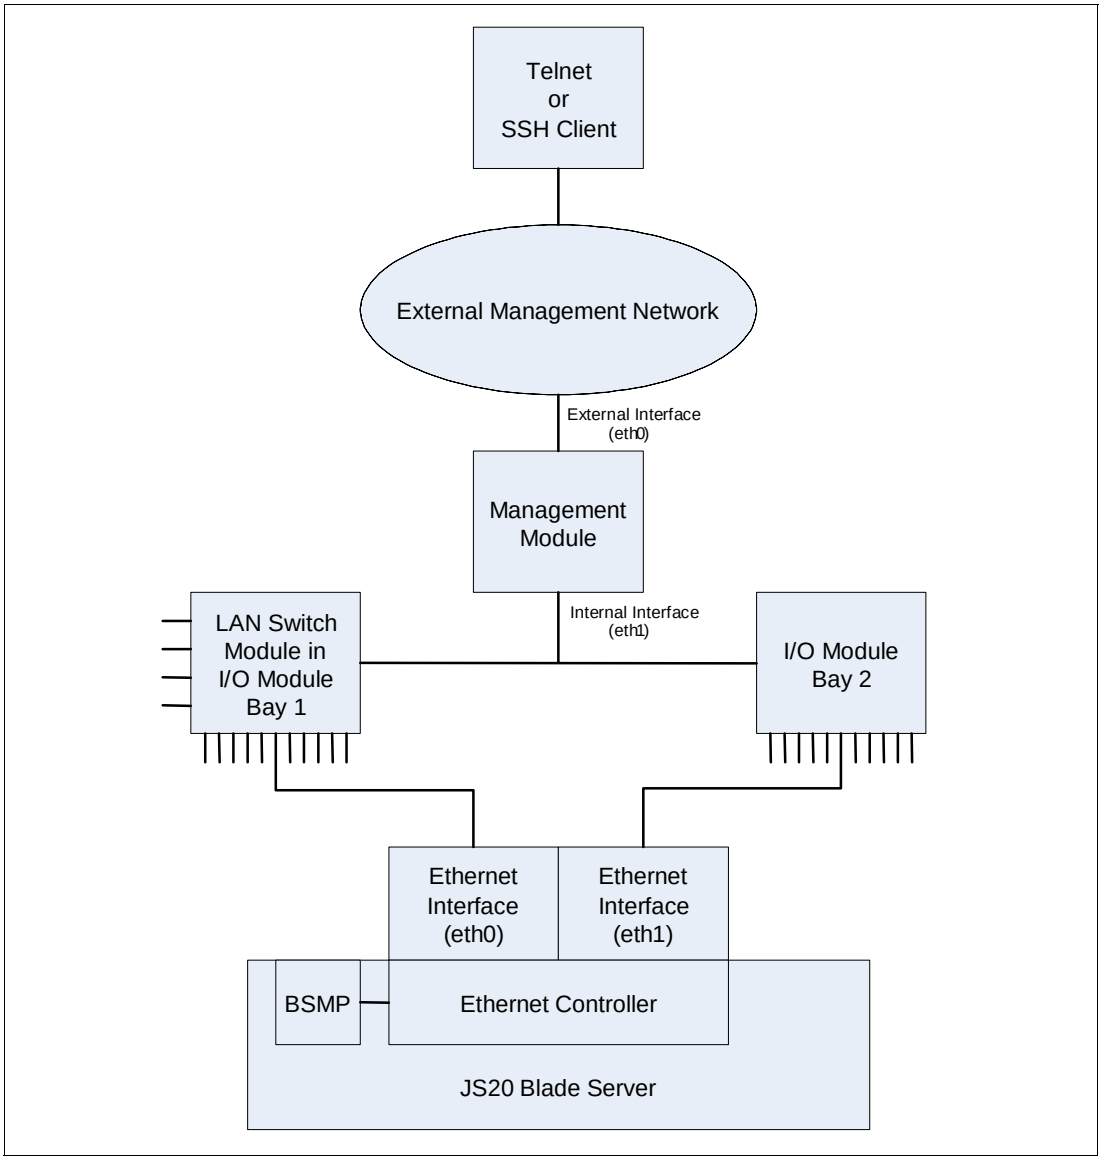

Figure 2-63 Serial over LAN components

In the BladeCenter environment, the integrated service processor (BMC or ISMP) and network interface controller (NIC) on each blade server route the serial data from the blade server serial communications port to the network infrastructure of the BladeCenter unit, including an Ethernet-compatible I/O module that supports SOL communication. BladeCenter components are configured for SOL operation through the BladeCenter Management Module.

The Management Module also acts as a proxy in the network infrastructure to couple a client running a Telnet session with the Management Module to an SOL session running on a blade server. This enables the Telnet client to interact with the serial port of the blade server over the network. Because all SOL traffic is controlled by and routed through the Management Module, administrators can segregate the management traffic for the BladeCenter unit from the data traffic of the blade servers.

To start an SOL connection with a blade server, you must first start a Telnet command line interface session with the Management Module. When this Telnet command line interface session is running, you can start a remote-console SOL session with any blade server in the BladeCenter unit that is set up and enabled for SOL operation. You can establish up to 20 separate Telnet sessions with a BladeCenter Management Module. For a BladeCenter unit, this enables you to have 14 simultaneous SOL sessions active (one for each of up to 14 blade servers), with six additional CLI sessions available for BladeCenter unit management.

For a BladeCenter T unit, this enables you to have eight simultaneous SOL sessions active (one for each of up to eight blade servers), with 12 additional command line interface sessions available for BladeCenter unit management. If security is a concern, you can use Secure Shell (SSH) sessions to establish secure Telnet command line interface sessions with the BladeCenter Management Module before starting an SOL console redirect session with a blade server.

The steps involved are as follows:

1. Using a Telnet or SSH client, connect to the BladeCenter Management Module CLI. This is usually via an external management network that is connected to the Management Module's 10/100BaseT Ethernet interface.
2. From the Management Module's CLI, initiate an SOL remote console session to the desired blade server.
3. The Management Module uses a private VLAN provided by the LAN switch module in I/O module bay 1 to transport the SOL data stream to the Ethernet interface of the target blade server.
4. The Ethernet controller of the target blade server passes the SOL data stream received from the private VLAN to the blade system management processor (BSMP), which manages the text console for the blade server.

Restriction: You can only have one active SOL remote console connection to each server.

Take the following points into consideration when using Serial over LAN (SOL):

- A SOL-capable Ethernet network switch must be present in module bay 1 of the BladeCenter chassis.
- You cannot use the OPM (Optical Pass-Thru Module) or the CPM (Copper Pass-Thru Module) for an SOL connection.
- SOL uses the first network interface, Ethernet 1 (eth1 or Planar Ethernet 1), of the blade server to communicate. When this network interface attempts to boot through PXE or DHCP, the network interface is reset, causing the current SOL session to be dropped and have a new status of Not Ready.

  If you require booting through PXE or DHCP, use the second network interface, Ethernet 2 (eth2 or Planar Ethernet 2), of the blade server and install an SOL-capable Ethernet I/O module in I/O-module bay 1.
- Access to the SOL remote text console function is via the Management Module CLI (command line interface). The CLI is documented in *IBM BladeCenter Management Module Command Line Interface Reference Guide*.
- The Management Module can support up to 20 concurrent CLI connections. This is sufficient to support the concurrent use of a SOL remote text console to each blade server in a full BladeCenter chassis. At the same time, the Management Module supports six additional CLI connections for other administrative activities.

For detailed documentation about how to set up and manage SOL connections for your specific environment, refer to *Serial over LAN (SOL) Setup Guide - IBM BladeCenter T*:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-54666

3

Blade server hardware configuration

This chapter describes, in detail, the hardware configuration of different blade servers and expansion options. It includes the following topics:

Section 1.1, “Support matrixes” on page 2 provides the following compatibility information:

- The servers supported in each chassis, Table 1-1 on page 2
- The I/O modules supported in each chassis, Table 1-2 on page 3
- The expansion cards supported in each server, Table 1-3 on page 5
- The operating systems supported on each server, Table 1-4 on page 7

3.1 Expansion cards

Each IBM BladeCenter server blade has the optional capability of accommodating one or more I/O expansion cards to provide redundant connections to the BladeCenter switch modules. Most expansion cards feature two I/O ports.

The routing of the expansion card's two ports is done through the BladeCenter's midplane to the I/O modules. One port from the expansion card is wired via the midplane to one I/O module bay, while the other port is wired to a second I/O module bay to provide redundancy against a single point of failure.

This section describes the five different expansion card form factors. This section also explains the different BladeCenter expansion cards that are available and what types are supported with each blade server.

Table 1-3 on page 5 is a matrix that shows which expansion cards are supported in each blade server.

Ethernet expansion cards are discussed in the following sections:

- 3.1.2, "Ethernet Expansion Card (CFFv)" on page 242
- 3.1.3, "2/4 Port Ethernet Expansion Card (CFFh)" on page 242
- 3.1.4, "BladeCenter Gigabit SFF Ethernet Expansion Card" on page 243
- 3.1.5, "NetXen 10 Gb Ethernet Expansion Card (CFFh)" on page 243
- 3.1.6, "Broadcom 2-port 10 Gb Ethernet Expansion Card (CFFh)" on page 244
- 3.1.7, "Broadcom 4-port 10 Gb Ethernet Expansion Card (CFFh)" on page 245

Combo Ethernet and Fibre Channel expansion cards are discussed in the following section:

- 3.1.8, "QLogic Ethernet & 4Gb Fibre Channel Expansion Card (CFFh)" on page 246
- 3.1.9, "QLogic Ethernet and 8Gb Fibre Channel Expansion Card (CFFh)" on page 247

Fibre Channel expansion cards are discussed in the following sections:

- 3.1.11, "BladeCenter 2 Gb Fibre Channel Expansion Card" on page 248
- 3.1.12, "QLogic 4 Gb Fibre Channel Expansion Card (SFF)" on page 249
- 3.1.13, "Emulex 4 GB SFF Fibre Channel Expansion Card" on page 249
- 3.1.14, "Emulex 4 GB Fibre Channel Expansion Card (CFFv)" on page 250

InfiniBand expansion cards are discussed in the following sections:

- 3.1.15, "InfiniBand DDR Host Channel Adapters" on page 251

Storage and cluster expansion cards are discussed in the following sections:

- 3.1.16, "QLogic iSCSI Expansion Card" on page 252
- 3.1.17, "SAS Expansion Card" on page 252
- 3.1.18, "SAS Connectivity Card" on page 254
- 3.1.19, "ServeRAID MR10ie (CIOv) Controller" on page 255

3.1.1 Form factors

IBM blade servers have a PCI-X connector (for older blade servers) or both a PCI-X and a PCI Express connector (newer servers). Different form factors of expansion cards are available for the blade servers—some use the PCI-X connector and some use the PCI Express connector.

There are currently five different form factors for expansion cards, as shown in Table 3-1.

Table 3-1 BladeCenter expansion card form factors

Form factor	Description	Protocol/connector used
StFF	Standard Form Factor	PCI-X
SFF	Small Form Factor	PCI-X
CFFv	Combination Form Factor (vertical)[a]	PCI-X
CFFh	Combination Form Factor (horizontal)[a]	PCI Express
HSFF	High Speed Form Factor	PCI Express

a. CFF is sometimes referred to as Combo Form Factor or Compact Form Factor.

Support matrix

The different form factors are supported in the blade servers, as per Table 3-2.

For specific supported expansion cards and blade servers, see Table 1-3 on page 5 or the next sections in this chapter for each blade.

Table 3-2 Blade support matrix for expansion cards

Blade	Type	StFF	SFF	CFFv	CFFh	HSFF
HC10	7996	No	No	No	No	No
HS12	8014	No	Yes	Yes	Yes	Yes
HS12	8028	No	Yes	Yes	Yes	Yes
HS20	8678	Yes	Yes	No	No	No
HS20	8832	Yes	Yes	No	No	No
HS20	8843	Yes	Yes	No	No	No
HS20	7981	Yes	Yes	No	No	No
HS21	8853	Yes	Yes	Yes	Yes	Yes
HS21 XM	7995	Yes	Yes	Yes	Yes	Yes
HS40	8839	Yes	Yes	No	No	No
LS20	8850	Yes	Yes	No	No	No
LS21	7971	Yes	Yes	Yes	Yes	Yes
LS22	7901	Yes	Yes	Yes	Yes	Yes
LS41	7972	Yes	Yes	Yes	Yes	Yes
LS42	7902	Yes	Yes	Yes	Yes	Yes
JS12	7998-60x	Yes	Yes	Yes	Yes	Yes
JS20	8842	Yes	No	No	No	No
JS21	8844	Yes	Yes	Yes	Yes	Yes
JS22	7998-61x	Yes	Yes	Yes	Yes	Yes
QS21	0792	No	Yes	Yes	Yes	Yes

Blade	Type	StFF	SFF	CFFv	CFFh	HSFF
QS22	0793	No	Yes	Yes	Yes	Yes
PN41	3020	No	No	No	No	No

Standard Form Factor (StFF)

The standard form factor is the original form factor for expansion cards in blade servers. An StFF expansion card uses the PCI-X connector on the blade server and expansion units. When a StFF expansion card is installed using the PCI-X connector it is not possible to install a CFFh or HSFF expansion card on the adjacent PCI Express connector (if one exists in the server). Also, when an StFF expansion card is installed, the second hard disk drive on some blade servers cannot be installed.

An StFF expansion card requires compatible switch modules to be installed in bay 3 and bay 4 of the BladeCenter chassis.

Figure 3-1 shows the placement of an StFF expansion card in a blade server.

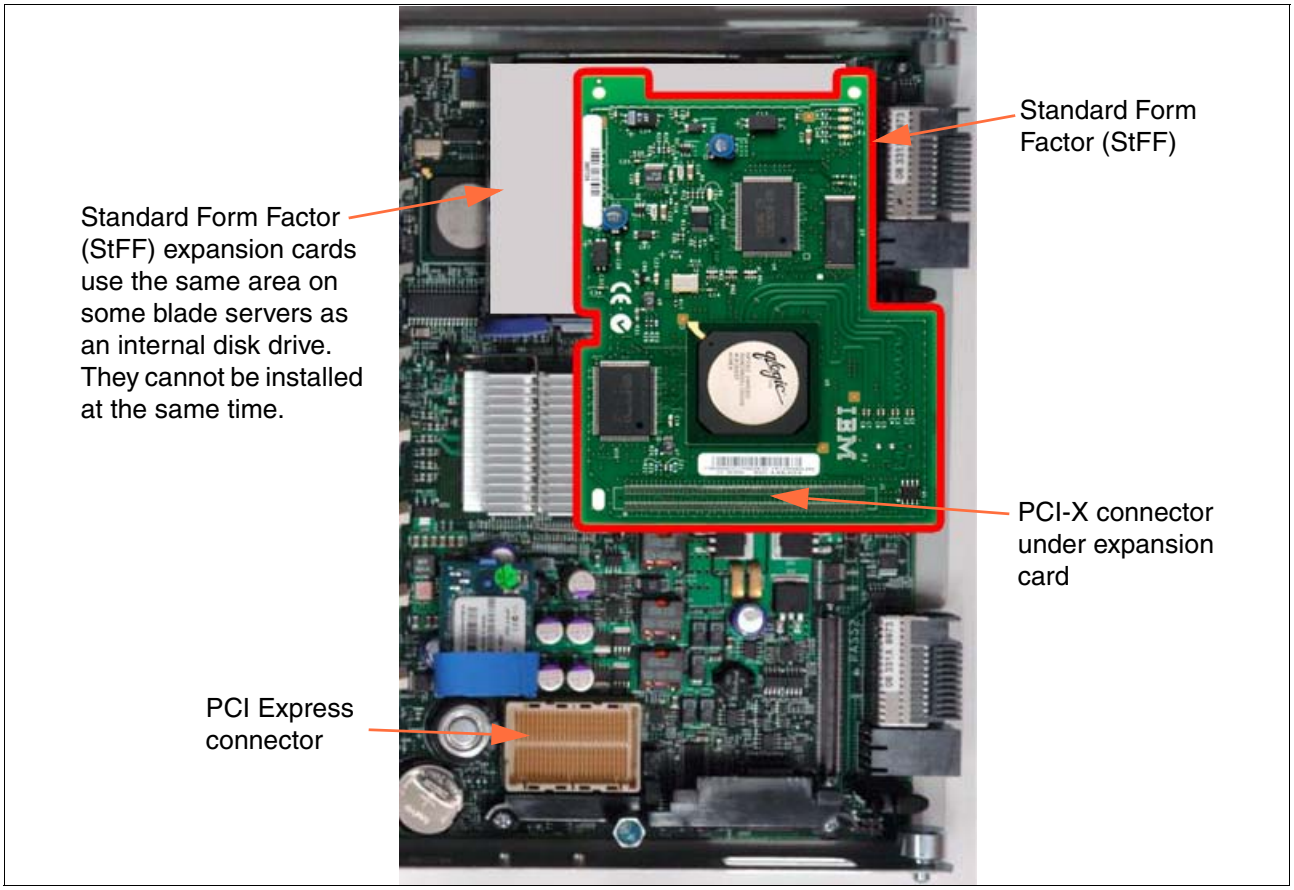

Figure 3-1 Rear section of a blade server showing StFF expansion card placement

Table 3-3 shows the expansion cards that are available in the Standard Form Factor size.

Table 3-3 Current Standard Form Factor (StFF) BladeCenter expansion cards

Part number	Feature Code	Description
32R1923	1458	QLogic iSCSI Expansion Card

Small Form Factor (SFF)

A small form factor adapter uses the PCI-X connector on the blade server or expansion unit. When an SFF expansion card is installed on the PCI-X connector it is not possible to install a CFFh or HSFF expansion card on the adjacent PCI Express connector. However, unlike the StFF card, the SFF expansion card does not interfere with the installation of a hard disk drive. An SFF expansion card can be installed in the blade server and can also be installed in an expansion unit.

A SFF adapter requires compatible switch modules to be installed in bay 3 and bay 4 of the BladeCenter chassis.

Figure 3-2 shows the placement of an SFF BladeCenter expansion card on a blade server.

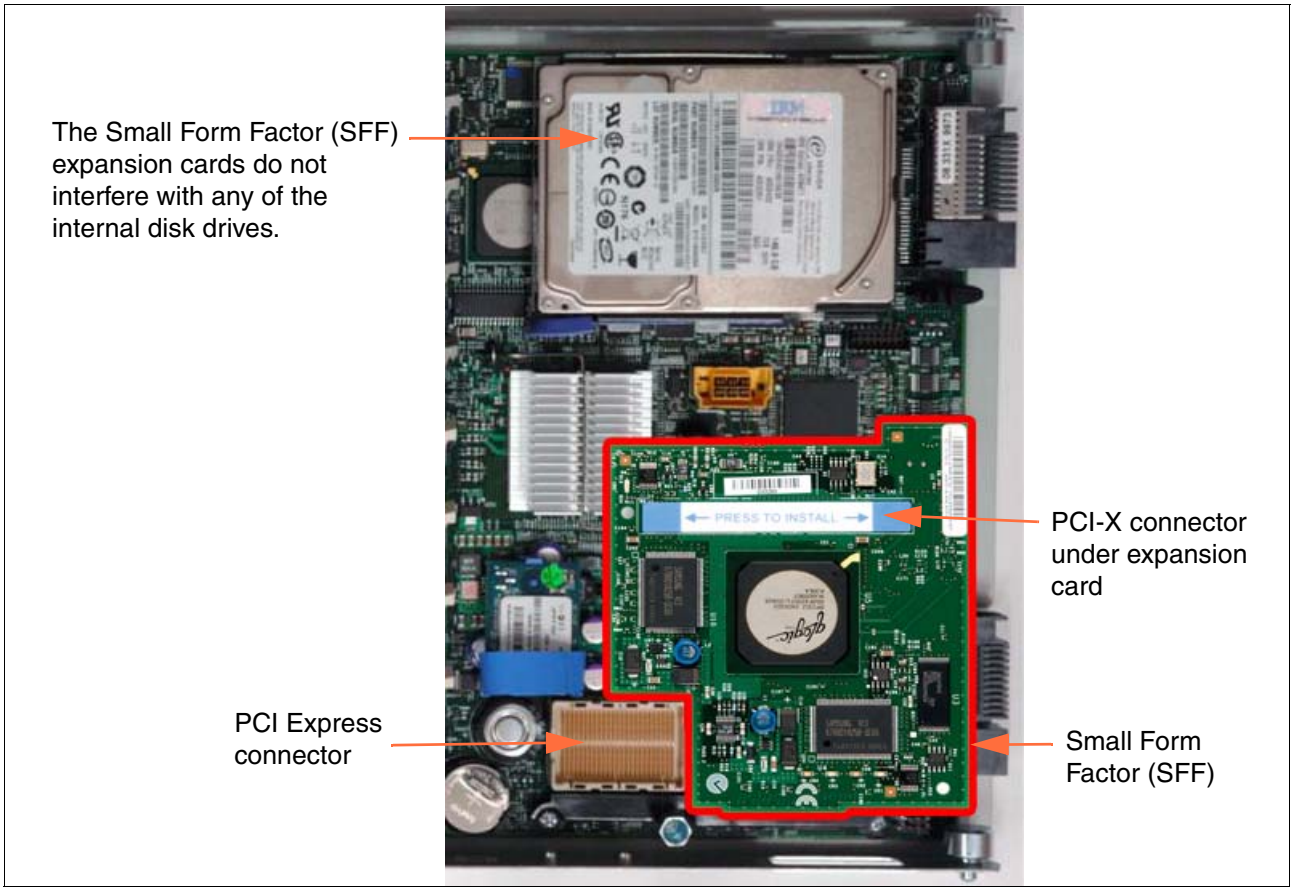

Figure 3-2 Rear section of a blade server showing SFF expansion card placement

Table 3-4 shows the BladeCenter expansion cards that are available in the Small Form Factor (SFF).

Table 3-4 Current Small Form Factor (SFF) BladeCenter expansion cards

Part number	Feature Code	Description
39R8624	1548	IBM BladeCenter SFF Gigabit Ethernet Expansion Card
26R0890	1577	QLogic 4 Gb SFF Fibre Channel Expansion Card
39Y9186	2925	Emulex 4 Gb SFF Fibre Channel Expansion Card

Combination Form Factor vertical (CFFv)

A CFFv expansion card uses the PCI-X connector on the blade server or expansion unit. The CFFv card can be used in conjunction with a CFFh card in servers that have both a PCI-X and PCI-Express connector (see Figure 3-6 on page 241). The function of a CFFv card is similar to the function of the equivalent SFF adapter. A CFFv adapter requires compatible switch modules to be installed in bay 3 and bay 4 of the BladeCenter chassis.

Figure 3-3 shows the placement of a CFFv expansion card on a blade server.

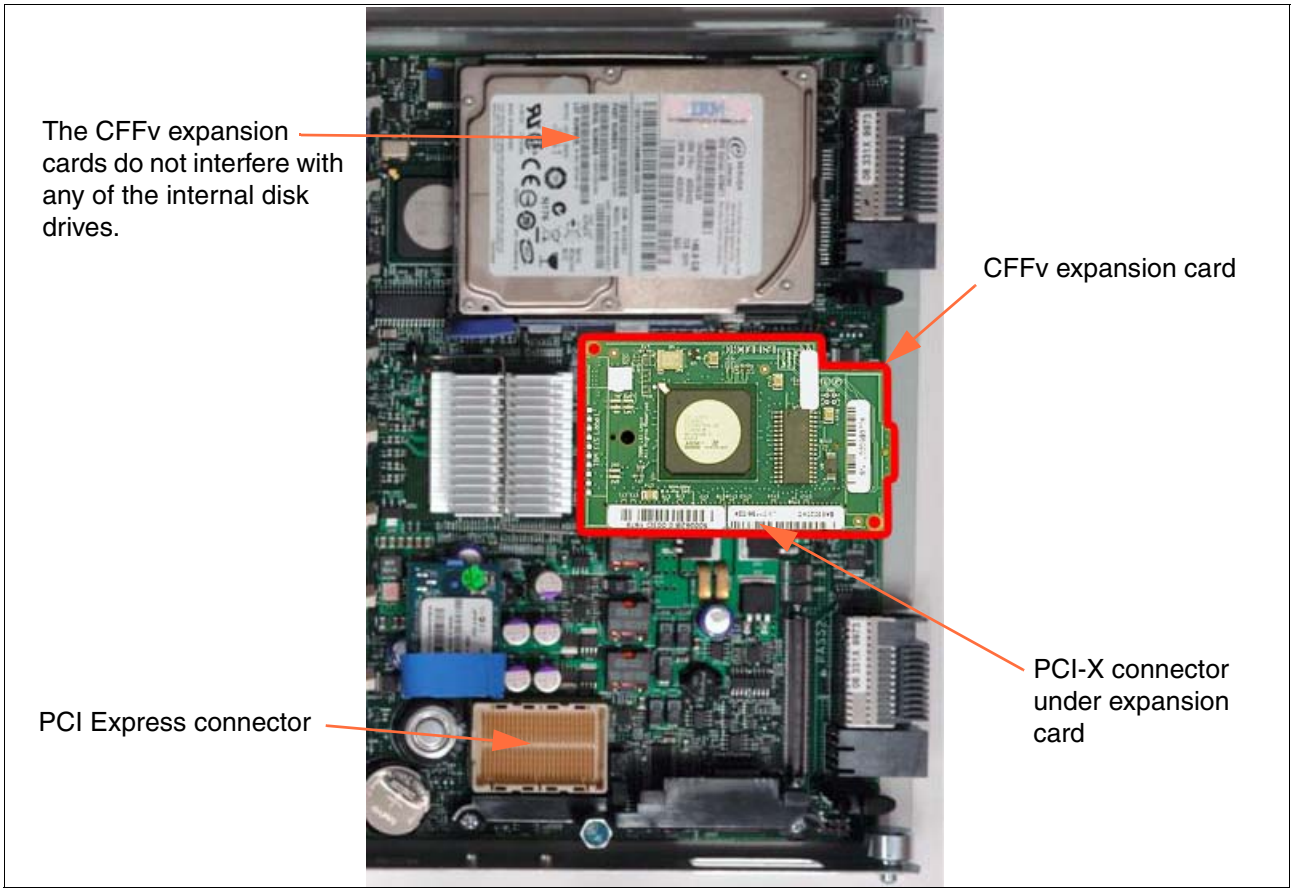

Figure 3-3 Rear section of a blade server showing CFFv expansion card placement

Table 3-5 shows the available CFFv expansion cards.

Table 3-5 Current CFFv expansion cards

Part number	Feature Code	Description
39Y9310	2969	Ethernet Expansion Card
41Y8527	2970	QLogic 4 Gb Fibre Channel Expansion Card
43W6859	2994	Emulex 4 GB Fibre Channel Expansion Card
39Y9190	2979	SAS Expansion Card
43W3974	1591	SAS Connectivity Card

Combination Form Factor horizontal (CFFh)

The CFFh expansion card attaches to the PCI Express connector of the blade server and therefore cannot be installed in a blade server that does not have a PCI Express connector. In addition, the CFFh adapter can only be used in servers that are installed in the BladeCenter H, BladeCenter HT, or BladeCenter S chassis.

The CFFh card can be used in conjunction with a CFFv adapter (Figure 3-5 on page 240). Only one CFFh card can be installed in a blade server. If an expansion unit (such as the SIO or MIO) is installed on the blade server, then the CFFh adapter must be installed in the top-most PCI Express slot. This is because the PCI Express connector on the blade server is also used to connect to the expansion unit.

A CFFh adapter requires that either:

- A Multi-Switch Interconnect Module (MSIM) is installed in bays 7 and 8, bays 9 and 10, or both.
- A high speed switch module be installed in bay 7 and bay 9.
- In the BladeCenter S, a compatible switch module is installed in bay 2.

The requirement of either the MSIM or high-speed switch modules depends on the type of CFFh expansion card installed. The MSIM must contain compatible switch modules. See 2.6.4, “Multi-switch Interconnect Module” on page 174, for more information about the MSIM.

Figure 3-4 shows the placement of a CFFh expansion card on a blade server.

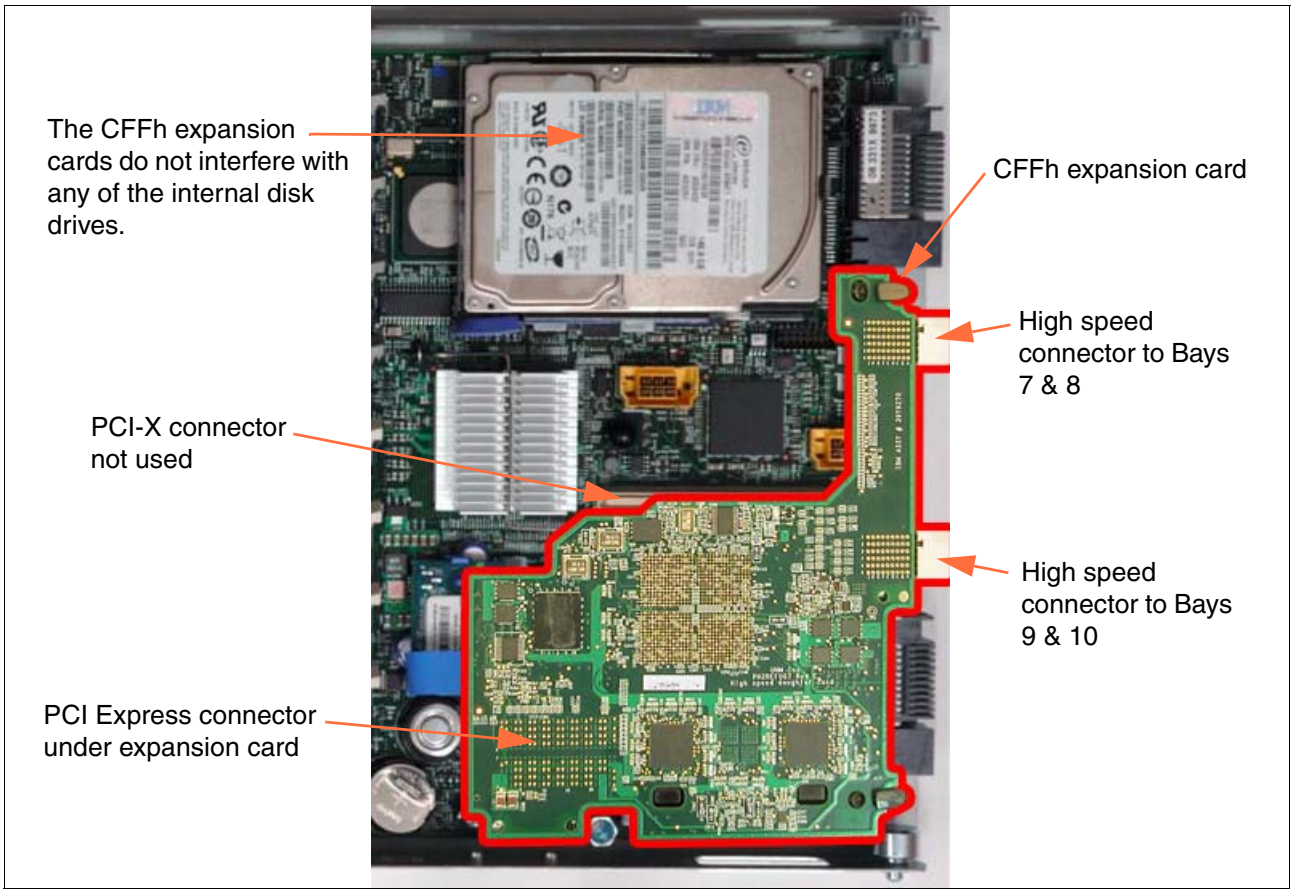

Figure 3-4 Rear section of a blade server showing CFFh expansion card placement

Table 3-6 shows the expansion cards that are available in the CFFh form factor.

Table 3-6 Current CFFh expansion cards

Part number	Feature code	Description	Required in bays 7 and 8, and bays 9 and 10
39Y9271	2967	NetXen 10 Gb Ethernet Expansion Card	Nortel 10 Gb Ethernet Switch Module
39Y9306	2968	QLogic Ethernet and 4 Gb Fibre Channel Expansion Card	MSIM with compatible Ethernet and FC Switch Modules
44W4466	5489	Broadcom 2-port 10 Gb Ethernet Expansion Card	Nortel 10 Gb Ethernet Switch Module in bay 7 and bay 9 only
44W4465	5479	Broadcom 4-port 10 Gb Ethernet Expansion Card	Nortel 10 Gb Ethernet Switch Module
44X1940	5485	QLogic Ethernet and 8 Gb Fibre Channel Expansion Card	MSIM with compatible Ethernet and FC Switch Modules
44W4479	5476	2/4 Port Ethernet Expansion Card	MSIM with compatible Ethernet Switch Modules
43W4423 43W4420	None	InfiniBand DDR Host Channel Adapters	InfiniBand Pass-thru Module

Figure 3-5 shows the CFFh expansion card installed together with a CFFv expansion card. These are the only form factors that can be installed at the same time.

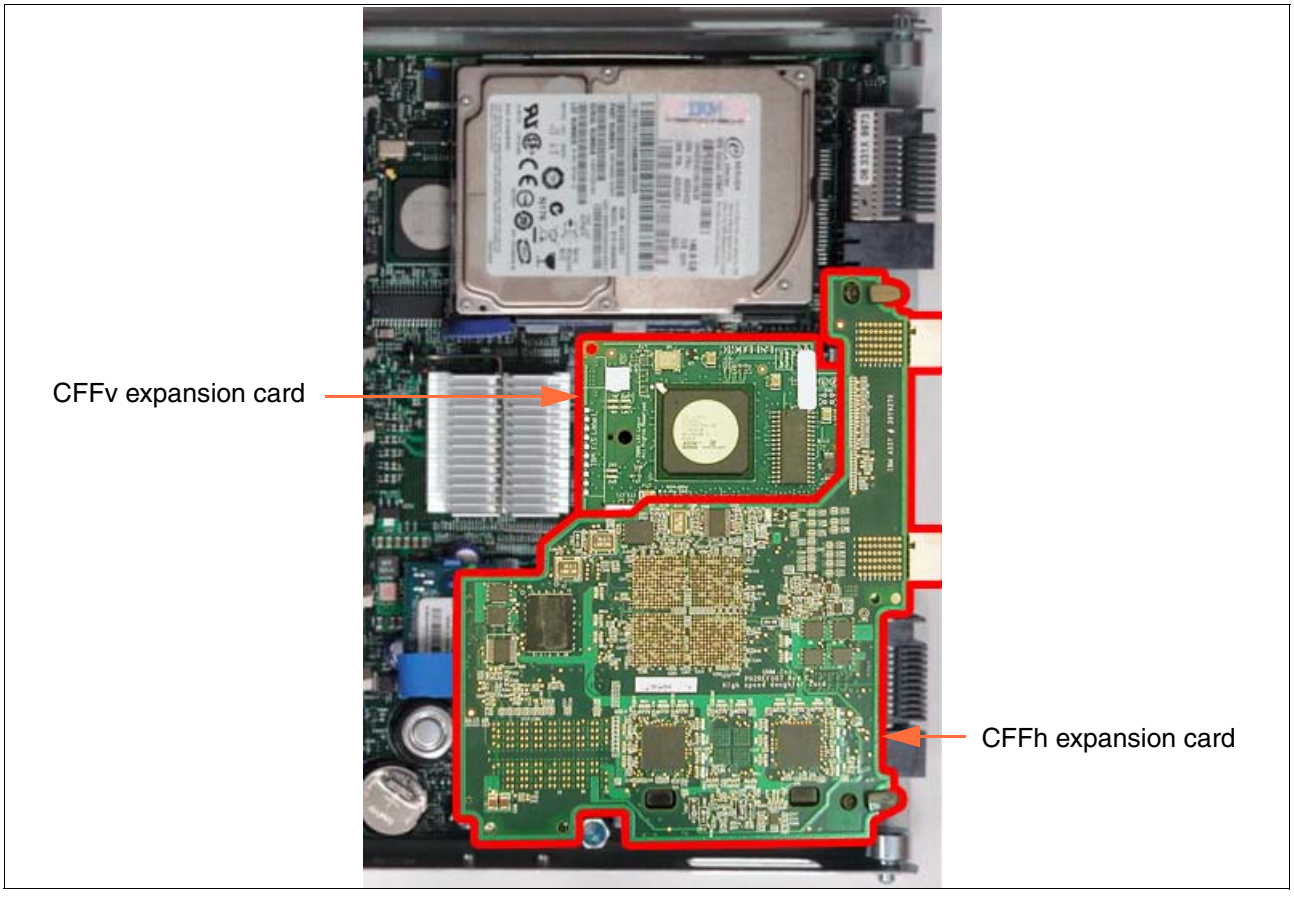

Figure 3-5 Rear section of a blade server showing CFFh and CFFv expansion cards installed together

High Speed Form Factor (HSFF)

With an HSFF adapter installed on the PCI Express connector, no other adapter can be installed on the adjacent PCI-X connector. Only one HSFF adapter can be installed in a blade server. If an expansion unit (like the SIO or MIO) is installed on the blade server, then the HSFF adapter must be installed in the top-most PCI Express slot. This is because the PCI Express connector on the blade server is also used to connect to the expansion unit.

The HSFF adapter cannot be installed in all blade servers. The HSFF adapter attaches to the PCI Express connector of the blade server and therefore cannot be installed in a blade server that does not have a PCI Express connector.

The HSFF adapter requires a compatible High Speed Switch Module to be installed in the high speed module bays 7 and 9.

The HSFF adapter is only supported in servers in the BC-H and BC-HT chassis.

Figure 3-6 show the placement of an HSFF BladeCenter expansion card on a blade server. Table 3-7 shows the BladeCenter expansion cards that are available in the High Speed Form Factor (HSFF).

Table 3-7 Current High Speed Form Factor BladeCenter expansion cards

Part Number	FC	Description
32R1760	1466	Cisco Systems 4X InfiniBand Expansion Card

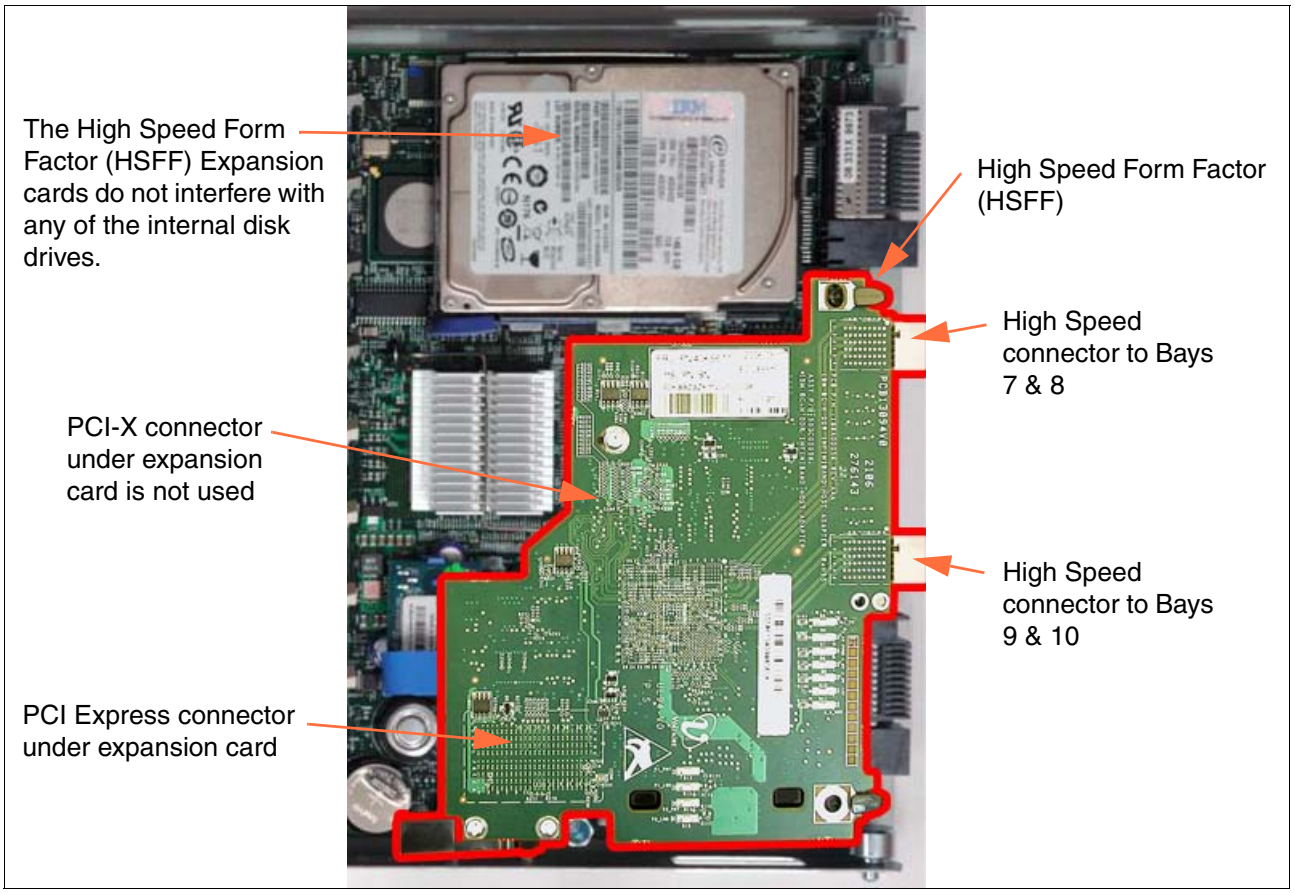

Figure 3-6 Rear section of a blade server showing HSFF expansion card placement

3.1.2 Ethernet Expansion Card (CFFv)

Part number: 39Y9310

Form factor: CFFv

This is a blade server expansion card with two 1 Gb Ethernet ports. It provides a PCI-X 1.0 interface via the Broadcom 5704 chipset. This expansion card can be combined with a CFFh expansion card on the same blade server.

An ESM, CPM, or OPM needs to be installed in switch bays 3 and 4.

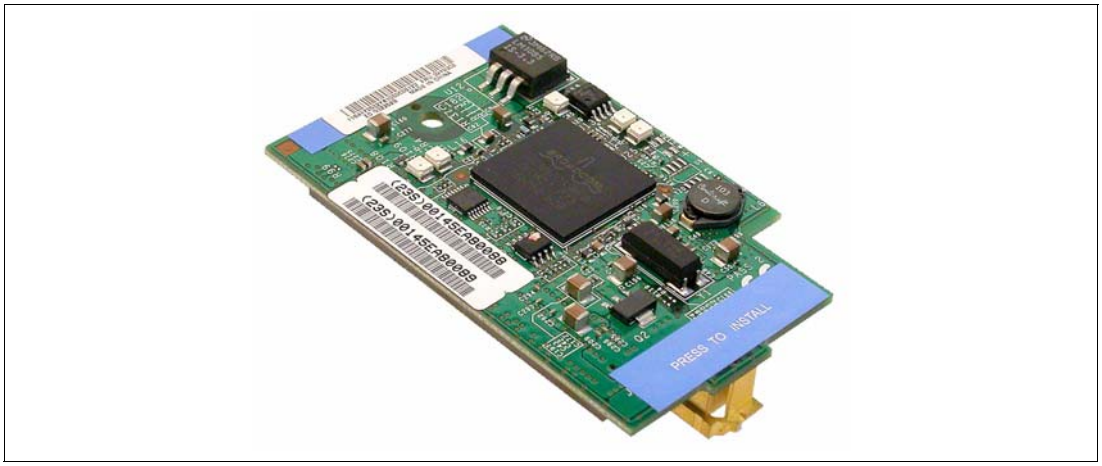

Figure 3-7 Ethernet Expansion Card (CFFv)

3.1.3 2/4 Port Ethernet Expansion Card (CFFh)

Part number: 44W4479

Form factor: CFFh

The 2/4 Port Ethernet Expansion Card is based on the Broadcom 5709 Gigabit Ethernet chipset. It enables 2 ports in the BladeCenter S Chassis. It enables 4 ports in the BladeCenter H and HT chassis.

The 2/4 Port Ethernet Expansion Card features:

- TCP offload engine (TOE)
- Full fast-path TCP offload
- TCP/IP checksum offload
- TCP segmentation offload
- PXE 2.0 remote boot support

This expansion card can be combined with a CFFv expansion card on the same blade server.

When attached to a blade server that is installed in a BladeCenter H or BladeCenter HT chassis, the 2/4 Port Ethernet Expansion Card is used in conjunction with the Multi-Switch Interconnect Module (MSIM), and requires that Gigabit Ethernet Switch Modules be installed

in the MSIM. In order for all four Ethernet ports to be used, two MSIMs and four Gigabit Ethernet Switch Modules are required.

When attached to a blade server that is installed in a BladeCenter S chassis, only two ports are used and both of these two ports are routed to bay 2. A supported Gigabit Ethernet Switch Module will need to be installed in bay 2 of the BladeCenter S chassis.

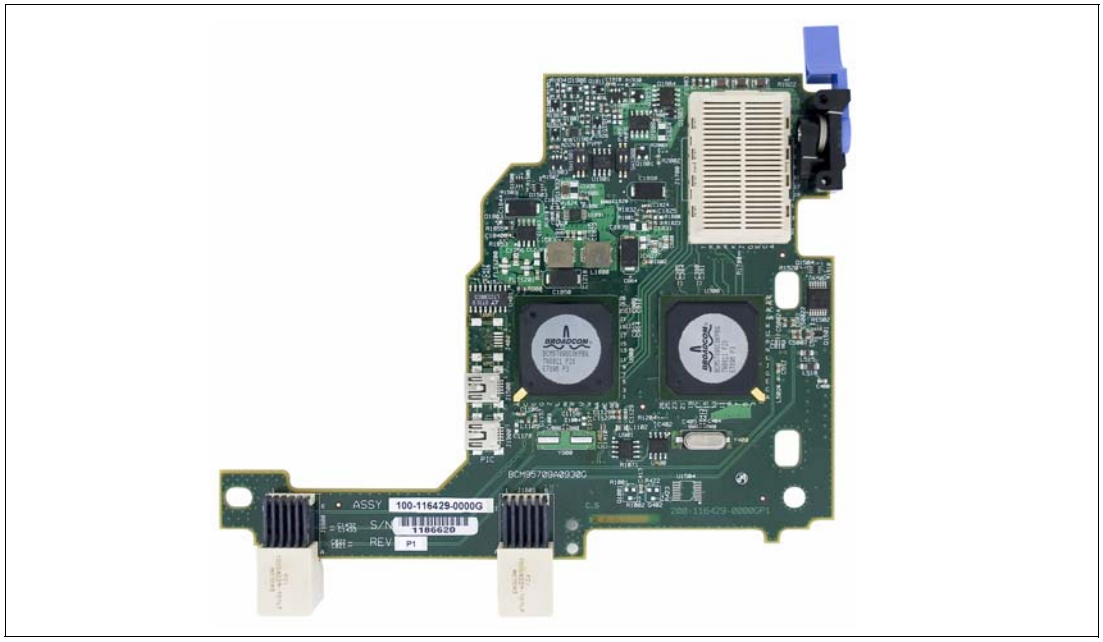

Figure 3-8 2/4 Port Ethernet Expansion Card (CFFh)

3.1.4 BladeCenter Gigabit SFF Ethernet Expansion Card

Part number: 39R8624

Form factor: Small Form Factor

The BladeCenter SFF Gigabit Ethernet Expansion Card, part number 39R8624, provides each blade server with two additional Gigabit ports. This option increases capacity, allows for expansion of the Ethernet I/O subsystem, enables additional bandwidth, and allows for connection to multiple LAN segments. It requires an Ethernet switch module to be installed in bay 3 and bay 4 of the BladeCenter chassis.

3.1.5 NetXen 10 Gb Ethernet Expansion Card (CFFh)

Part number: 39Y9271

Form factor: CFFh

This is a blade server expansion card with two 10 Gigabit Ethernet interfaces. It is a high-speed expansion card especially designed to interface with the Nortel 10 Gb Ethernet Switch Module, and it uses the NetXen NX2031 network controller chip.

This expansion card can be combined with a CFFv expansion card on the same blade server.

This expansion card requires that a Nortel 10 Gb Ethernet Switch Module be installed in switch bay 7 and bay 9.

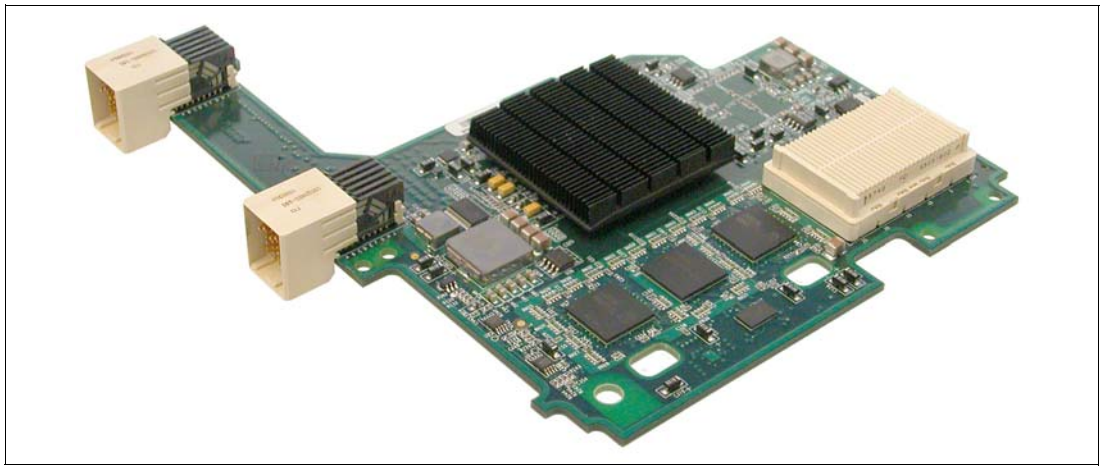

Figure 3-9 NetXen 10 Gb Ethernet Expansion Card (CFFh)

3.1.6 Broadcom 2-port 10 Gb Ethernet Expansion Card (CFFh)

Part number: 44W4466

Form factor: CFFh

The Broadcom 2-port 10 Gb Ethernet Expansion Card is based on the Broadcom 57710 module. It has two 10 Gigabit Ethernet interfaces. It is a high-speed expansion card especially designed to interface with a 10 Gb Ethernet Switch Module.

The Broadcom 2-port 10 Gb Ethernet Expansion Card features:

- TCP offload engine (TOE)
- Full fast-path TCP offload
- TCP/IP checksum offload
- TCP segmentation offload

This expansion card can be combined with a CFFv expansion card on the same blade server.

This expansion card connects to two 10 Gb Ethernet Switch Modules in high speed switch bays 7 and 9.

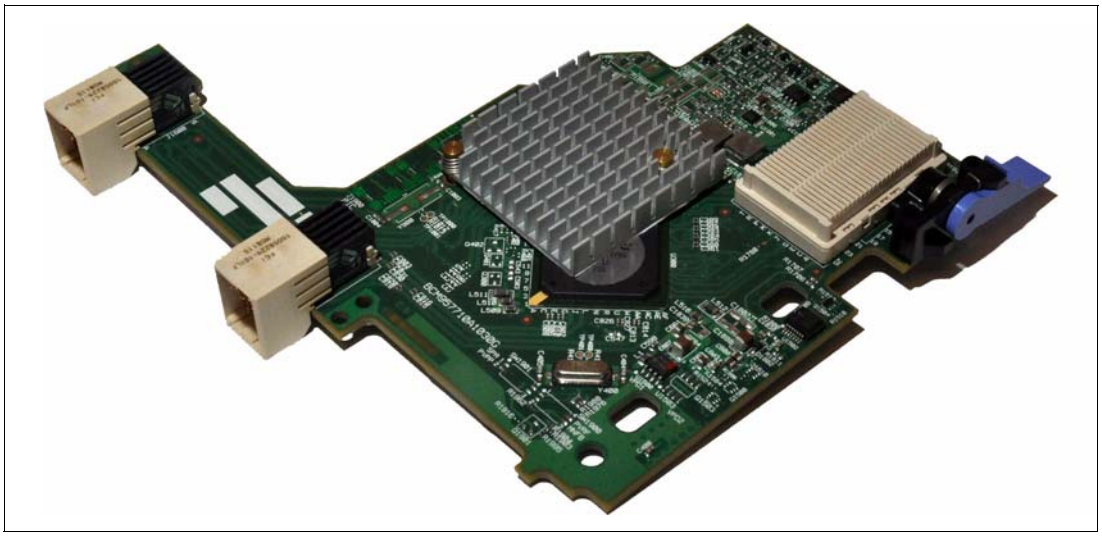

Figure 3-10 Broadcom 2-port 10 Gb Ethernet Expansion Card (CFFh)

3.1.7 Broadcom 4-port 10 Gb Ethernet Expansion Card (CFFh)

Part number: 44W4465

Form factor: CFFh

This is a blade server expansion card with four 10 Gigabit Ethernet interfaces. It is a high-speed expansion card especially designed to interface with a 10 Gb Ethernet Switch Module. The Broadcom 4-port 10 Gb Ethernet Expansion Card includes two dual port Broadcom 57710 modules.

The Broadcom 4-port 10 Gb Ethernet Expansion Card features:

- TCP offload engine (TOE)
- Full fast-path TCP offload
- TCP/IP checksum offload
- TCP segmentation offload

This expansion card can be combined with a CFFv expansion card on the same blade server.

This expansion card connects to four 10 Gb Ethernet Switch Modules in high speed switch bays 7, 8, 9, and 10.

Figure 3-11 Broadcom 4-port 10 Gb Ethernet Expansion Card (CFFh)

The 2-port and 4-port Broadcom 10 Gb cards are very similar in appearance. The visual differences between them are shown in Figure 3-12.

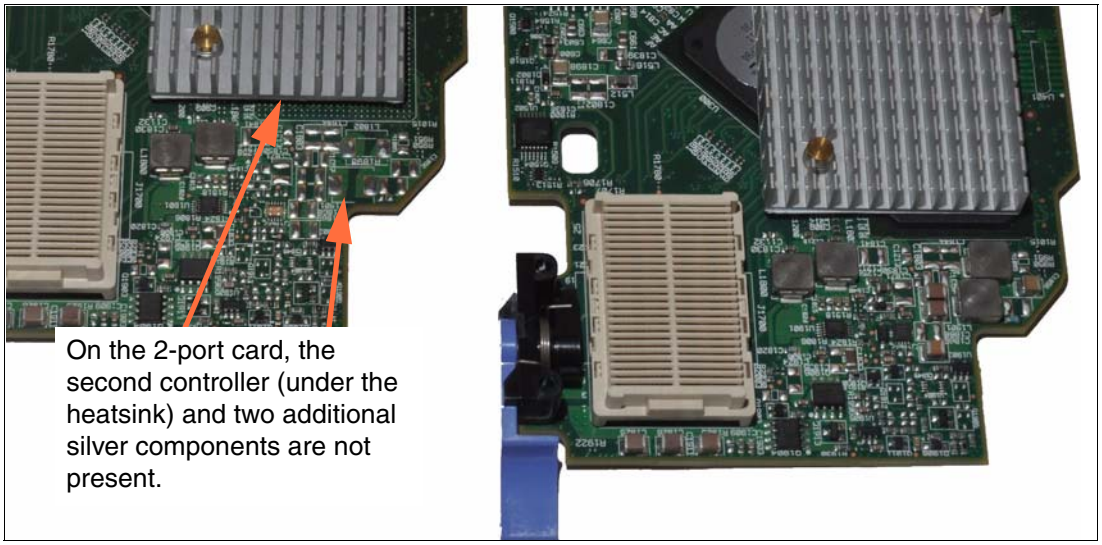

Figure 3-12 Broadcom 2-port (left) and 4-port (right) 10 Gb Ethernet Expansion Cards

3.1.8 QLogic Ethernet & 4Gb Fibre Channel Expansion Card (CFFh)

Part number: 39Y9306

Form factor: CFFh

This is a blade server expansion card with two 4 Gb FC ports and two 1 Gb Ethernet ports. It has the following chipsets:

- A QLogic 2432M PCI-Express x4 ASIC with two 4 Gb Fibre Channel ports
- A Broadcom 5715S PCI-Express x4 ASIC with two 1 Gb Ethernet ports

This expansion card can be combined with a CFFv card on the same blade server.

The QLogic Ethernet and 4 Gb FC Expansion Card is used in conjunction with the Multi-Switch Interconnect Module (MSIM), and requires in the MSIM either or both of the following:

- A supported Ethernet switch module installed in the left bay (for the Ethernet connections)
- A supported Fibre Channel switch module installed in the right bay (for the Fibre Channel connections)

Both switches do not need to be installed because the Fibre Channel and Ethernet networks are separate and distinct. For example, if you do not wish to enable the Fibre Channel connections of the expansion card, you do not need to install a FC switch module in the MSIM.

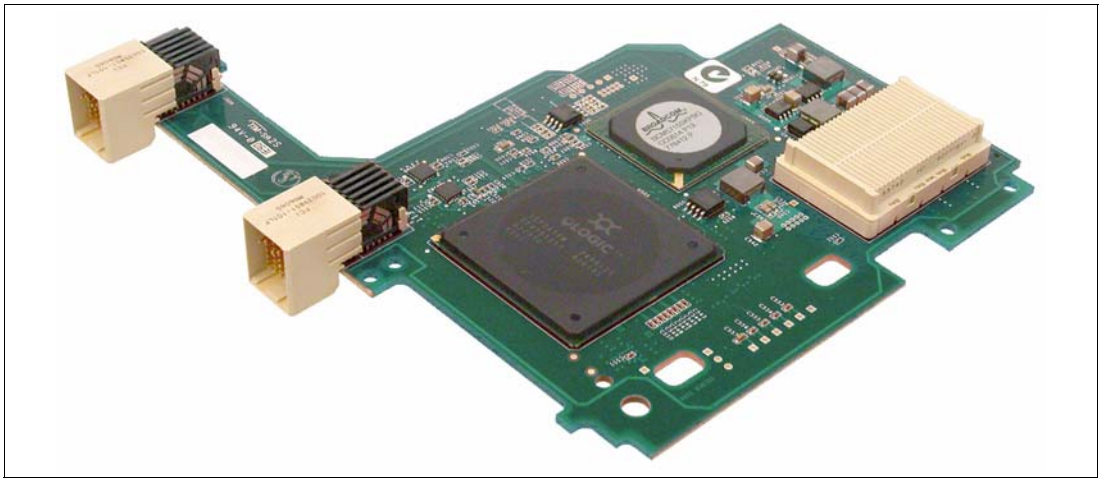

Figure 3-13 QLogic Ethernet & 4Gb Fibre Channel Expansion Card (CFFh)

3.1.9 QLogic Ethernet and 8Gb Fibre Channel Expansion Card (CFFh)

Part number: 44X1940

Form factor: CFFh

This is a blade server expansion card with two 8 Gb FC ports and two 1 Gb Ethernet ports. It has the following chipsets:

- A QLogic 2532 ASIC with two 8 Gb Fibre Channel ports
- A Broadcom 5709S ASIC with two 1 Gb Ethernet ports

This expansion card can be combined with a CFFv card on the same blade server.

The QLogic Ethernet and 8Gb FC Expansion Card is used in conjunction with the Multi-Switch Interconnect Module (MSIM), and requires, in the MSIM, either or both of the following:

- A supported Ethernet switch module installed in the left bay (for the Ethernet connections)
- A supported Fibre Channel switch module installed in the right bay (for the Fibre Channel connections)

Both switches do not need to be installed because the Fibre Channel and Ethernet networks are separate and distinct. For example, if you do not wish to enable the Fibre Channel connections of the expansion card, you do not need to install a FC switch module in the MSIM.

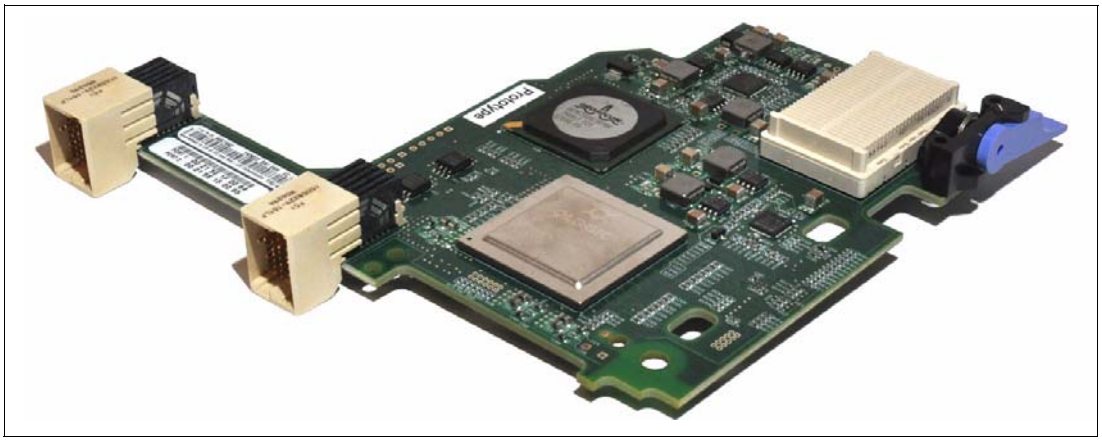

Figure 3-14 QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)

3.1.10 QLogic 4 Gb Fibre Channel Expansion Card (CFFv)

Part number: 41Y8527

Form factor: CFFv

This is a blade server expansion card with two 4 Gb Fibre Channel ports. It provides a PCI-X 1.0 interface for two 4 Gb Fibre Channel ports via a QLogic 2422M ASIC. This expansion card can be combined with a CFFh expansion card on the same blade server.

A Fibre Channel Switch Module, Optical Pass-thru Module or Intelligent Pass-thru Module needs to be installed in switch bays 3 and 4 of the BladeCenter chassis.

Note: When using an OPM, the speed of the Fibre Channel link will be limited to 2 GB. This is a limitation of the OPM.

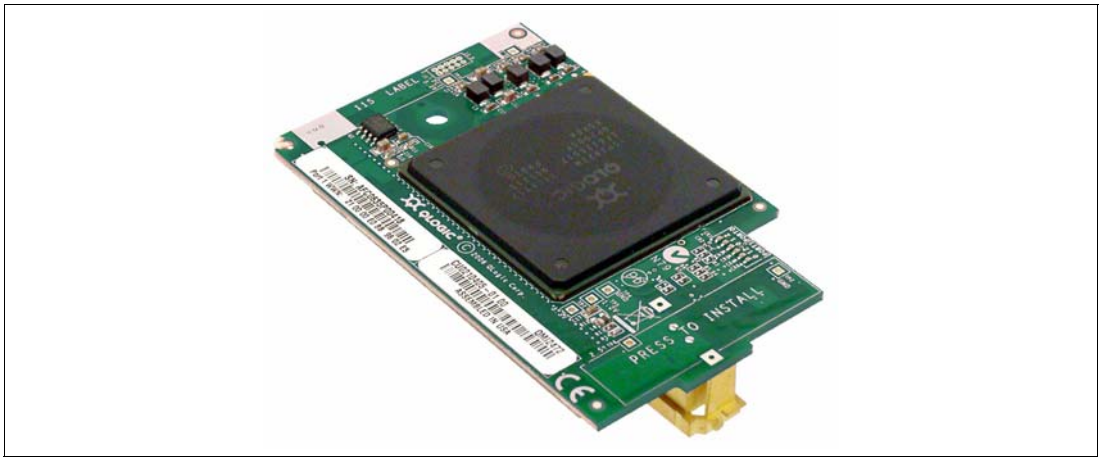

Figure 3-15 QLogic 4 Gb Fibre Channel Expansion Card (CFFv)

3.1.11 BladeCenter 2 Gb Fibre Channel Expansion Card

Part number: 26K4841

Form factor: Small Form Factor

The IBM BladeCenter 2 Gb Fibre Channel Expansion Card contains two Fibre Channel controllers supporting up to 2 Gbps transmissions. It provides functions equivalent to the IBM TotalStorage® FC2-133 Host Bus Adapter.

A FCSM, OPM, or IPM needs to be installed in switch bays 3 and 4 of the BladeCenter chassis.

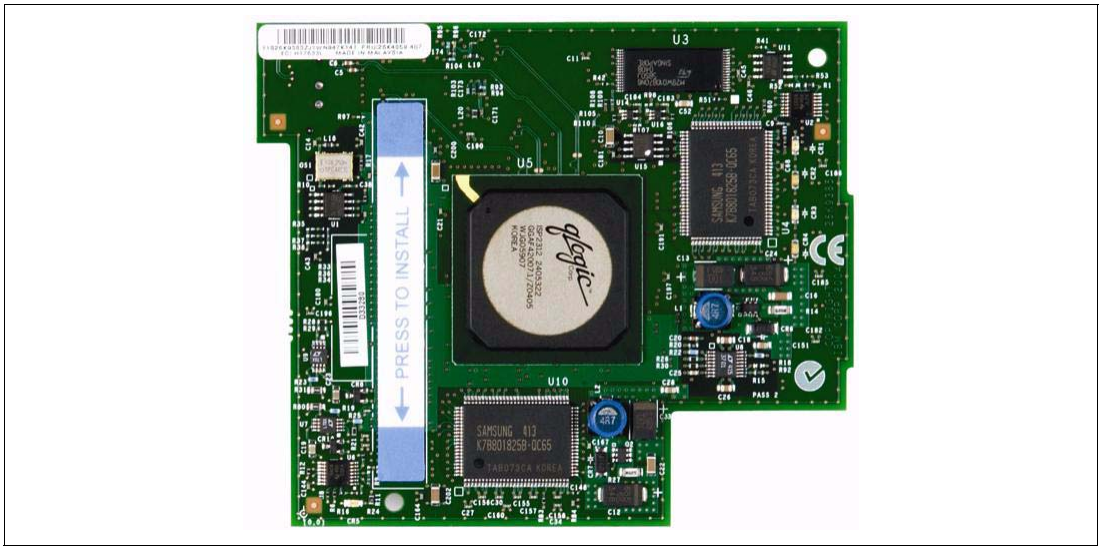

Figure 3-16 BladeCenter 2 Gb Fibre Channel Expansion Card

3.1.12 QLogic 4 Gb Fibre Channel Expansion Card (SFF)

Part number: 26R0890

Form factor: Small Form Factor (SFF)

The QLogic 4 Gb Fibre Channel expansion card provides a high-bandwidth connection between BladeCenter server blades and SAN switches. This 4 Gb expansion cards allow for a complete end-to-end 4 Gbps FC blade server to SAN solution.

A FCSM, OPM, or IPM needs to be installed in switch bays 3 and 4 of the BladeCenter chassis.

Note: When using an OPM, the speed of the Fibre Channel link will be limited to 2 GB. This is a limitation of the OPM.

3.1.13 Emulex 4 GB SFF Fibre Channel Expansion Card

Part number: 39Y9186

Form factor: Small Form Factor

The Emulex 4 Gb SFF Fibre Channel expansion card is an addition to the family of IBM BladeCenter expansion cards and delivers a high bandwidth connection between BladeCenter server blades and SAN switches, and allows for a true end-to-end 4 Gb Fibre Channel blade server to a storage area network (SAN) solution. It is built based on the Emulex Helios Fibre Channel controller.

A FCSM, OPM, or IPM needs to be installed in switch bays 3 and 4 of the BladeCenter chassis.

Note: When using an OPM, the speed of the Fibre Channel link will be limited to 2 GB. This is a limitation of the OPM.

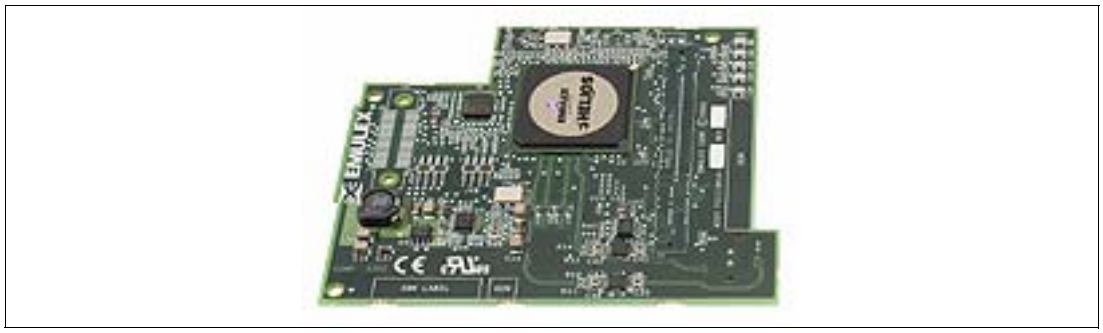

Figure 3-17 Emulex 4 GB SFF Fibre Channel Expansion Card

3.1.14 Emulex 4 GB Fibre Channel Expansion Card (CFFv)

Part number: 43W6859

Form factor: CFFv

The Emulex 4 Gb Fibre Channel Expansion Card (CFFv) is an addition to the family of IBM BladeCenter expansion cards, which due to its CFFv form factor, allows it to be installed at the same time as a CFFh type expansion card. This is a PCI-X based host bus adapter with two 4 Gb Fibre Channel ports that can also auto-negotiate to support 1 Gb and 2 Gb data rates. Part of the Emulex Light Pulse family, the equivalent Emulex part number is LP1105-BCv.

A FCSM, OPM, or IPM needs to be installed in switch bays 3 and 4 of the BladeCenter chassis.

Note: When using an OPM, the speed of the Fibre Channel link will be limited to 2 GB. This is a limitation of the OPM.

Figure 3-18 Emulex 4 GB Fibre Channel Expansion Card (CFFv)

3.1.15 InfiniBand DDR Host Channel Adapters

Form factor: CFFh

Two 4X DDR Host Channel Adapters (HCAs) are available, resulting from a partnership between IBM and Mellanox. Software differentiates the two offerings, in that they use a common CFFh base HCA, and then depending on what part number are ordered, will come with support for different software.

These HCAs require that the InfiniBand Pass-thru Module (or other supported InfiniBand switch module) are installed in bay 7 or bay 9 of the BladeCenter H and HT chassis or both bays 7 and 9 for redundancy.

The details on these options are as follows:

- 4X InfiniBand DDR Expansion Card (43W4423)

  For this HCA the user must independently obtain drivers. This will initially offer limited support from IBM, with OFED support being available from the OFED Web site.

  http://www.openfabrics.org/

- Voltaire 4X InfiniBand DDR Expansion Card (43W4420)

  This version of the HCA will come with software and support from Voltaire, and comes with a license to use the Voltaire GridStack product. The software is also based on the OFED stack with enhancements and improved support to interface with Voltaire switches.

Key features include:

- 1.2µs MPI ping latency
- Two 4x DDR capable (20 Gbps) InfiniBand ports; support auto-negotiation to SDR
- CPU offload of transport operations
- End-to-end QoS and congestion control
- Hardware-based I/O virtualization
- TCP/UDP/IP stateless offload
- Based on Mellanox's new ConnectX InfiniBand to PCI-express bridge chip
- Supports the OpenFabrics Enterprise Distribution (OFED) drivers and protocol support
- Support for Cisco's and Voltaire's drivers stacks

Additional options include:

- InfiniBand 3m DDR Cable (43W6742)
- InfiniBand 8m DDR Cable (43W7243)

Compatibility with other InfiniBand options for BladeCenter.

The DDR expansion cards are supported only with the InfiniBand Pass-thru Module. Check IBM ServerProven for the latest support information:

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

3.1.16 QLogic iSCSI Expansion Card

Part number: 32R1923

Form factor: Standard Form Factor (StFF)

Figure 3-19 QLogic iSCSI Expansion Card

The QLogic iSCSI Expansion Card is a hardware initiator that provides iSCSI communication from the blade server to an iSCSI storage device (target).

With the iSCSI Expansion Card, you gain:

- Full hardware-based iSCSI initiator and TOE for storage traffic only (TCP/IP offload engine) function
- Blades enabled to run diskless in a non-Fibre Channel SAN environment
- Dual port card using the QLogic QMC4052 controller
- NEBS-compliant for use in NGN/Telco environments

This card requires a supported Ethernet switch module in bays 3 and 4 of the BladeCenter chassis.

3.1.17 SAS Expansion Card

Part number: 39Y9190

Form factor: CFFv

The SAS Expansion Card (CFFv) is a configurable SAS device that allows data to be sent and received through its connector on the Blade Server to I/O bay modules 3 and 4. The SAS expansion card must be installed in each Blade Server that requires access to the SAS I/O modules and SAS storage devices. Once you have assigned disks to a blade using the SAS Connectivity Modules predefined or user defined configurations, you will be able create RAID arrays across those disks using the SAS expansion cards LSI configuration utility.

The SAS Expansion Card (CFFv) enables a SAS connection from a Blade Server to the SAS connectivity module and external SAS storage fabric. The SAS Expansion Card is required in each Blade Server installed in a BladeCenter S chassis to enable connection to the internal

disk drives in the BladeCenter S Disk Storage Modules (DSM). The SAS Expansion Card is also required in blade servers installed in a BladeCenter E or BladeCenter H chassis that are required to boot from the IBM BladeCenter Boot Disk System.

Key features are:

- PCI-X bus adapter
- Uses LSI 1064 SAS Controller ASIC
- Provides blade with two 1X SAS ports for connecting to the SAS Connectivity Module (see 2.5.1, “SAS Connectivity Module” on page 164).
- CFFv form factor; CFFv expansion cards when combined with a CFFh expansion card allow blade server I/O connectivity to both standard and high-speed I/O module bays in the BladeCenter H chassis.

Figure 3-20 shows the SAS Expansion Card.

Figure 3-20 SAS Expansion Card (CFFv)

The following RAID levels are supported by the SAS Expansion Card:

- RAID-0 (Integrated Striping - IS)
- RAID-1 (Integrated Mirroring - IM)
- RAID-1E (Integrated Mirroring Enhanced - IME)

The maximum number of drives supported by an IM volume is two, plus one optional global hot spare. An IME volume supports up to ten HDDs, plus two optional hot spares. An IS volume supports up to ten HDDs. The IS volume does not support hot spare drives.

Mixing HDDs of different capacities in a single volume is supported. However, total volume size is aligned with the size of the smallest HDD. Excess space on larger HDDs is not used.

Supported combinations of volumes include:

- Two IM or IME volume per blade server
- One IM or IME volume and IS volume per blade server
- Two IS volumes per blade server

Both SAS and SATA HDDs are supported, as is intermixing SAS/SATA drives. However, each volume must have hard disks of the same type, SAS or SATA.

3.1.18 SAS Connectivity Card

Part number: 43W3974

Form factor: CFFv

The SAS Connectivity Card (CFFv) for IBM BladeCenter is compatible with the HS21 XM and HS12 blades. It is a low-cost alternative to the SAS Expansion Card (CFFv) option (part number 39Y9190).

Unlike the SAS Expansion Card, the SAS Connectivity Card does not contain a SAS controller. Instead, it uses the unused SAS controller ports from the blade for enabling a SAS connection from the blade server to the SAS Connectivity Module and the external SAS storage fabric. The SAS Connectivity Card requires SAS I/O modules to be installed in bay 3 and bay 4 of the BladeCenter chassis.

The SAS Connectivity Card is shown in Figure 3-21.

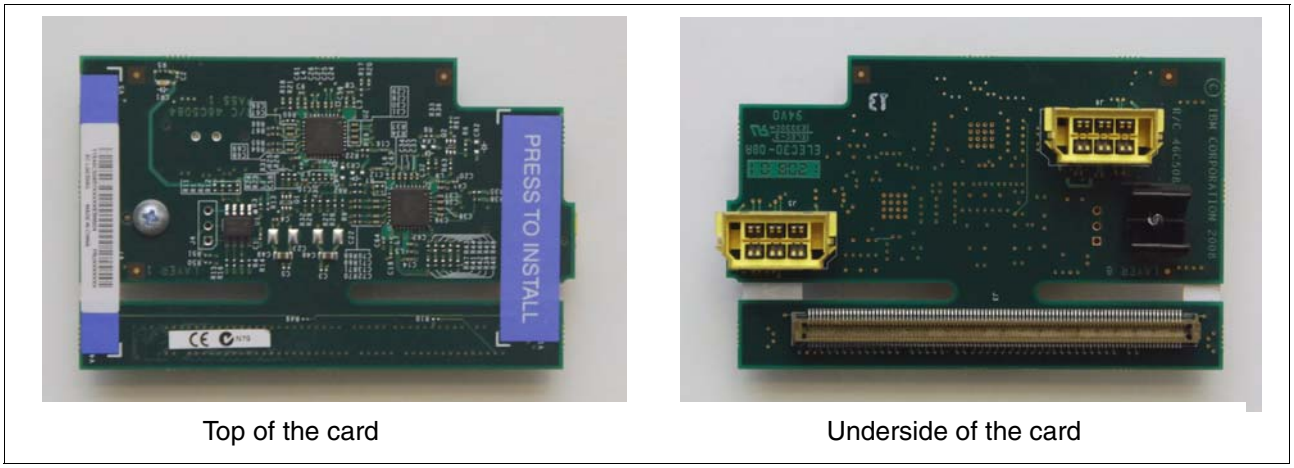

Top of the card Underside of the card

Figure 3-21 SAS Connectivity Card

The RAID functionality is supplied by the onboard SAS controller in the blade server. The following RAID levels are supported by onboard SAS controller:

- RAID-0 (Integrated Striping - IS)
- RAID-1 (Integrated Mirroring - IM)
- RAID-1E (Integrated Mirroring Enhanced - IME)

The maximum number of drives supported by an IM volume is two, plus one optional global hot spare. An IME volume supports up to ten HDDs, plus two optional hot spares. An IS volume supports up to ten HDDs. The IS volume does not support hot spare drives.

Mixing HDDs of different capacities in a single volume is supported. However, total volume size is aligned with the size of the smallest HDD. Excess space on larger HDDs is not used.

Supported combinations of volumes include:

- Two IM or IME volume per blade server
- One IM or IME volume and IS volume per blade server
- Two IS volumes per blade server

Both SAS and SATA HDDs are supported, as is intermixing SAS/SATA drives. However, each volume must have hard disks of the same type, SAS or SATA.

The SAS Connectivity Card is currently only supported in the HS12 and the HS21 XM blades.

3.1.19 ServeRAID MR10ie (CIOv) Controller

Part number: 46C7167

Form factor: CIOv

The ServeRAID MR10ie is an SAS RAID Controller I/O Expansion Card for a blade server. It provides RAID support with battery-backed cache for onboard and external disk drives. Currently, the ServeRAID MR10ie (CIOv) Controller is only supported in the HS12 type 8028 blade server.

Key features of the ServeRAID MR10ie controller are:

- Provides 256 MB of ECC DDR-2 battery-backed cache
- Utilizes an LSI 1078 controller
- Supports up to 26 disk drives
- Support for onboard and external disk drives
- Includes battery backup module
- Supports RAID levels 0, 1, 10, 5, 50, 6 and 60
- Support for both SAS & SATA disks
- Support for OS Boot over SAS

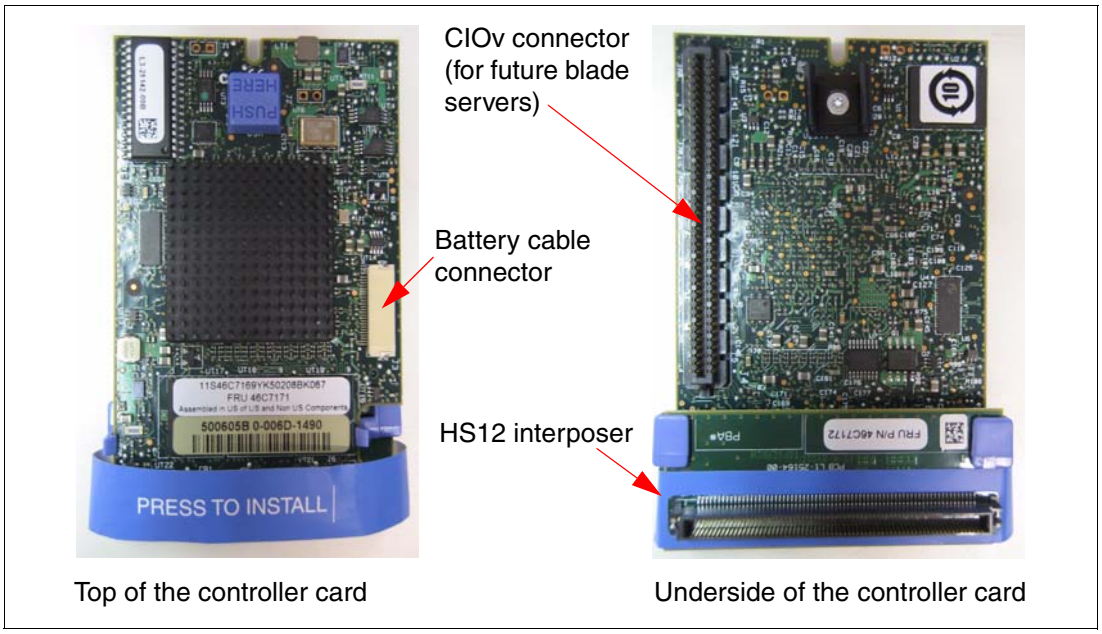

Figure 3-22 ServeRAID MR10ie (CIOv) Controller

The battery is in the form of a DIMM, as shown in Figure 3-23.

Figure 3-23 The Battery Backup unit for the ServeRAID-MR10ie (CIOv) Controller

The ServeRAID MR10ie expansion card has a connector to attach it to the paddle card socket on a HS12 type 8028 blade server. The ServeRAID MR10ie comes with an interposer that allows the connector to reach the socket on the HS12 blade server. Figure 3-22 on page 255 shows the location of the interposer.

The ServeRAID MR10ie ships:

- Battery backup unit, in the form of a DIMM (see Figure 3-23 on page 255)
- Interposer for use with the HS12 (see Figure 3-22 on page 255)
- Two ribbon cables to connect the card to the battery, one for the HS12 and one for future blade server types

IBM recommends that you replace the battery unit annually. The replacement battery unit is part number 46M0800.

When the ServeRAID MR10ie is used in the HS12 to connect to external disk drives, it also requires the installation of the SAS Connectivity Card (part number 43W3974). The SAS Connectivity Card passes the signal from the ServeRAID MR10ie via SAS switch modules in bays 3 and 4 and onto external disk drives. These external disk drives may reside in the BladeCenter S chassis or in an EXP3000 expansion unit attached to the external ports of the SAS switch module. See 5.10.2, “Advanced local storage using ServeRAID MR10ie” on page 426 for more information.

The following tools are provided to configure the ServeRAID controller:

- WebBIOS Configuration Utility
- MegaRAID Command Tool
- MegaRAID Storage Manager
- Storage Configuration Manager

The WebBIOS Configuration Utility is provided by the ServeRAID controller BIOS. It operates independently of the operating system. The MegaRAID Storage Manager and MegaRAID Command Tool are provided on the ServeRAID controller Support CD.

3.2 HC10 workstation blade

The IBM BladeCenter HC10, machine type 7996, provides server-based computing technology for workstation applications by offering high security and manageability, while delivering outstanding graphics performance and full USB capability. Built-in features are ideal for engineering design applications, trading floor and other financial applications, Geographic Information Systems applications, distance collaboration, and more.

The HC10 is designed to provide hardware-based graphics compression and encryption that delivers performance and security. Encrypted USB signals are transmitted transparently over the TCP/IP network to a workstation device. No additional software or device drivers are required by customers for encryption and transmission of graphics and USB information over the network. This workstation solution is designed for Microsoft Windows XP and Vista.

The IBM CP20 Workstation Connection Device workstation connection device is a stateless device with USB, video, and audio connectors. It connects to the HC10 via a standard Ethernet network over TCP/IP, as shown in Figure 3-24.

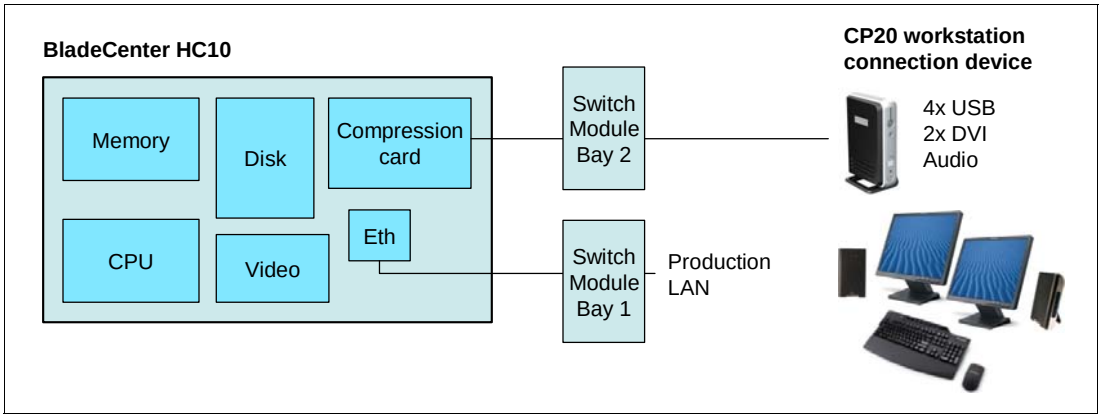

Figure 3-24 HC10 connectivity

The HC10 is an ideal solution for customers that have space or heat restrictions at the end-user work area, or end-user data security requirements. The workstation connection device on the end-user's desk has no moving parts, consumes less than 25 Watts of power, and is about the size of a book (232 mm H x 94 mm W x 174 mm D). The USB ports can be secured if required to prevent the use of data transfer using *thumb drives* or other storage media. The HC10 is also suitable for customers that require physical centralization of PC resources, but need more CPU and graphics power than a terminal server-based solution.

The HC10 contains the hardware typically found in a PC: processor, memory, hard disk, video card, and Ethernet connections, whereas the keyboard, mouse, video, speakers, and USB devices are connected to the IBM CP20 Workstation Connection Device workstation connection device. The IBM CP20 Workstation Connection Device and HC10 communicate using a special protocol called PC-over-IP™, which runs over TCP/IP over Ethernet.

The connection between these two devices is via an Ethernet Switch Module that is installed in bay 2 of the BladeCenter chassis. All ESMs are supported and the IBM Server Connectivity Module is recommended if you only plan to install HC10 blades in the chassis. See 2.2.6, "IBM Server Connectivity Module" on page 113, for details on this switch module.

The HC10 has a second Ethernet port for use by the installed operating system. This Ethernet connection is routed through the Ethernet Switch Module in bay 1 of the chassis. The operating system installed on the HC10 only *sees* this one Ethernet connection.

From a networking perspective, you can isolate the HC10-CP20 communications. However, because this communications is encrypted, it is not necessary for security purposes. Both switch modules can be connected to the same production network if required.

The HC10 has the following features:

- One dual-core Intel Core™ 2 Duo, up to 2.66 GHz
- 1066 MHz front-side bus
- 2 MB or 4 MB L2 cache
- Up to 8 GB of DDR2 SDRAM (PC2-5300 667 MHz)
- One SATA disk bay with an 80 GB SATA drive standard

- nVidia graphics adapter standard; choice of:
  - Quadro NVS 120M Professional 2D Graphics adapter
  - Quadro FX 1600M Advanced 3D Graphics adapter
- Integrated baseboard management controller (BMC)
- Preloaded Microsoft Windows Vista Business Blade PC Edition (64-bit)
- Supported operating systems:
  - Microsoft Windows XP Professional
  - Microsoft Windows XP Professional x64 Edition
  - Microsoft Windows Vista Business (32-bit/64-bit)
  - Microsoft Windows Vista Enterprise (32-bit/64-bit)
  - Microsoft Windows Vista Ultimate (32-bit/64-bit)
- One-year customer replaceable unit and onsite limited warranty

Table 3-8 lists the available models.

Table 3-8 BladeCenter HC10 models

Model	Intel Core 2 Duo	Speed	L2 cache	Mem (std/max)	Graphics	Disk
7996-35x	E6400	2.13 GHz	2 MB	1 GB/8 GB	Quadra NVS 120M 2D	80 GB SATA
7996-55x	E6700	2.66 GHz	4 MB	2 GB/8 GB	Quadra NVS 120M 2D	80 GB SATA
7996-5Rx	E6700	2.66 GHz	4 MB	2 GB/8 GB	Quadra FX 1600M 3D	80 GB SATA
7996-5Sx	E6700	2.66 GHz	4 MB	4 GB/8 GB	Quadra FX 1600M 3D	80 GB SATA

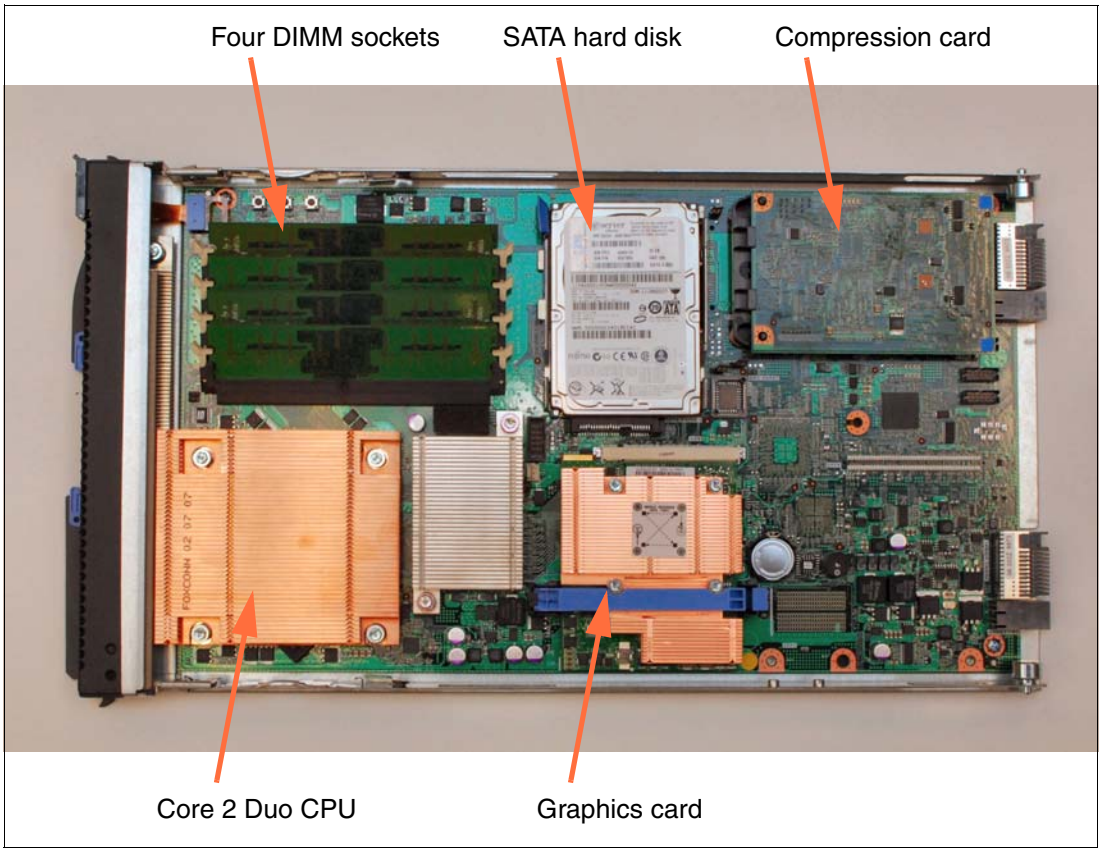

Figure 3-25 The BladeCenter HC10

3.2.1 Processor

The HC10 has one processor standard and maximum. No upgrades are available. There are two processors available, depending on the model. The only differences between the two processors are the internal clock speed and the size of the L2 cache, as explained here and listed in Table 3-9.

- Intel Core 2 Duo E6500: 2.13 GHz clock speed
- Intel Core 2 Duo E6700: 2.66 GHz clock speed

Table 3-9 Processor choices for the HC10

Processor	Clock speed	Front-side bus	L2 cache size	CPU power consumption
Intel Core 2 Duo E6400	2.13 GHz	1066 MHz	2 MB	65 W
Intel Core 2 Duo E6700	2.66 GHz	1066 MHz	4 MB	65 W

Both processors have the following features:

- Dual core architecture with shared L2 cache
- Intel 64 Architecture (EM64T)
- 1066 MHz front-side bus
- Support for Intel Virtualization Technology
- 65 nm technology

3.2.2 Memory

The memory used in the HC10 is unbuffered DDR2 SDRAM running at 667 MHz (PC2-5300). DIMMs are non-ECC and must be installed in matched pairs (sockets 1 and 3, and sockets 2 and 4).

See Table 3-10 for blade server HS20 type 7981 memory options.

Table 3-10 Blade server HS20 type 7981 memory options

Part number	Memory description
41Y2822	1 GB (2 x 512 MB kit) PC2-5300 CL5 DDR2 SDRAM RDIMM
41Y2825	2 GB (2 x 1 GB kit) PC2-5300 CL5 DDR2 SDRAM RDIMM
41Y2828	4 GB (2 x 2 GB kit) PC2-5300 CL5 DDR2 SDRAM RDIMM

3.2.3 Network controllers

The HC10 uses a Broadcom 5708S dual-port Gigabit Ethernet controller. This chip has the following features:

- PCI Express-based
- TCP/IP Offload Engine (TOE) support
- Wake on LAN
- Serial over LAN (SOL)
- PXE 2.0 boot agent
- Alert Standard Format (ASF) 2.0

Note: The iSCSI initiator in the Broadcom 5708S is not enabled.

The 5708S has two ports, as follows (see Figure 3-24 on page 257):

- Connectivity to bay 1 of the chassis for use by the HC10 operating system
- Connectivity to the Compression Card and IBM CP20 Workstation Connection Device routing through bay 2 in the chassis

3.2.4 Graphics adapters

There one of two graphics adapters installed as standard depending on the HC10 model (see Table 3-8 on page 258). The adapters are installed in an MXM3 type connector on the planar. Both are PCI Express x16 adapters.

- The NVIDIA Quadro NVS 120M graphics board are targeted as the professional 2D workstation graphics solution. It has 128 MB of onboard video memory. It has a 64-bit interface.
- The NVIDIA Quadro FX 1600M graphics is targeted at customers needing 3D graphics capabilities. It has 256 MB of onboard video memory. It has a 256-bit memory interface and supports Shader Model 4.0

Both support up to two displays. Resolution support is as follows:

- One monitor attached: maximum resolution is 1920 x 1200 at 75 Hz
- Two monitors attached: maximum resolution is 1600 x 1200 at 60 Hz on each monitor

3.2.5 Integrated systems management processor

The HC10 has an integrated Baseboard Management Controller (BMC) that provides:

- Intelligent Platform Management Interface (IPMI) 1.5 compliant
- Serial over LAN (SOL)
- Wake on LAN (WOL)
- Power management functions
- KVM and media tray ownership
- Light path diagnostics support
- Automatic BIOS recovery
- Automatic server restart
- Predictive failure analysis (PFA) on system components (processor, memory, and drives)
- Inventory
- Error logging
- Environmental monitoring (voltages and temperature on the system board)

The BMC is not directly accessible by the administrator. The BMC is only used as an interface to the management module in the BladeCenter chassis. The management module is a single point of control for hardware management of all server and switch components contained within the chassis.

3.2.6 Local storage

The HC10 has an 80 GB SATA drive as standard with the following features:

- 80 GB capacity
- 2.5” SFF form factor
- SATA interface, 300 MBps
- 7200 RPM rotation speed
- Non-hot swap

Note: The 60 GB drive option has been withdrawn.

3.2.7 I/O expansion options

There are no PCI-X or PCI Express connectors on the HC10 to support I/O expansion cards.

3.2.8 IBM CP20 Workstation Connection Device features

The IBM CP20 Workstation Connection Device is the end-user device that has connectors for keyboard, mouse, displays, and audio system. Figure 3-26 shows the rear of the device with the power switch, USB 1.1 ports, Gigabit Ethernet port, and two DVI display ports.

Figure 3-26 IBM CP20 Workstation Connection Device end-user device (left: front; right: rear)

The front of the device has two additional USB 1.1 ports, a disconnect button, a remote power button, HD audio sockets (microphone and speaker), and status LEDs.

For more information about the HC10, see *Implementing the IBM BladeCenter HC10 Workstation Blade*, REDP-4331, which is available from:

http://www.redbooks.ibm.com/abstracts/redp4331.html

3.3 HS12

The BladeCenter HS12 server (machine types 8028 and 8014) is the first single CPU socket blade server offered by IBM. With the increase in performance of dual and quad core CPUs, single socket servers today have enough processing power to run most single application server workloads.

The HS12 is a single-socket Intel CPU-based blade with the RAS features that clients expect from IBM BladeCenter, at a lower price point than a two-socket blade server. This makes the HS12 an affordable blade server option for customers looking to replace aging tower and rack servers. The HS12 is supported in all of the current IBM BladeCenter chassis, but when combined with the BladeCenter S chassis, it provides a true, cost-effective blade server platform. When combined with the BladeCenter E chassis, it provides the most power-efficient blade solution in the industry.

The internal layout of the HS12 type 8028 (with hot-swap drives) is shown in Figure 3-27.

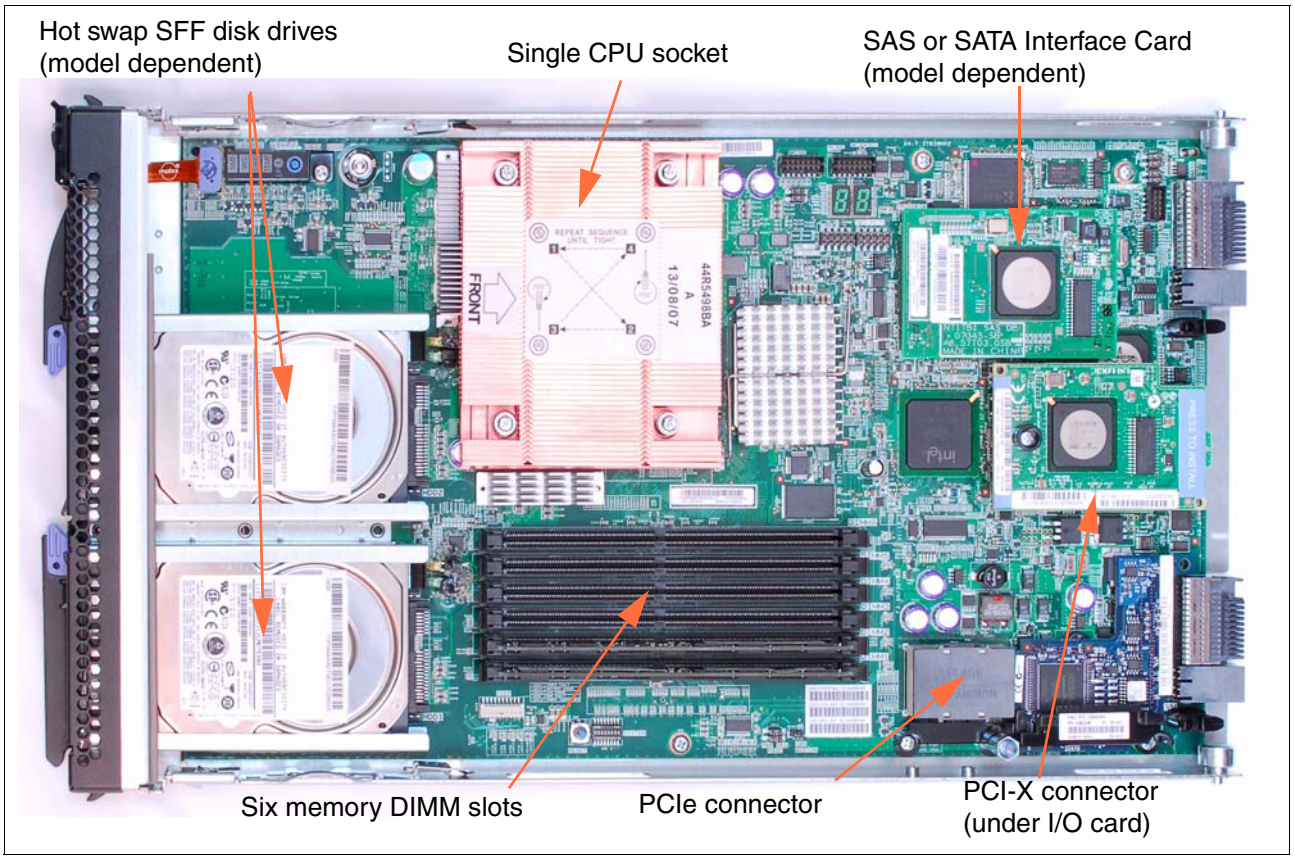

Figure 3-27 The HS12 blade server (machine type 8028)

The HS12 type 8014 (without hot-swap drives) is shown in Figure 3-28 on page 263.

Figure 3-28 The HS12 blade server (machine type 8014)

Features of the HS12 include:

- Single processor socket supporting either dual-core or quad-core Intel Xeon 3000 series processors in type 8028 models.
- Up to 24 GB of system memory in 6 DIMM sockets.
- Dual Gigabit Ethernet connections.
- Support for PCI-X or PCI-Express I/O expansion cards.
- Two 2.5” small form factor (SFF) drive bays
  - Machine type 8028: hot-swap SAS or solid state drive bays.
  - Machine type 8014: non-hot-swap SATA or solid state drive bays.
- Support for RAID-0 or RAID-1 on internal disks (machine type 8028).
- Support for the Storage and I/O (SIO) expansion unit blade with an additional three 2.5-inch SAS HDD bays and RAID support.
- Optional concurrent keyboard, video, mouse support with the addition of the IBM BladeCenter Concurrent KVM (cKVM) feature card. See 3.17, “Concurrent KVM Feature Card (cKVM)” on page 356.
- Integrated systems management processor.
- Supported in all IBM BladeCenter chassis.
- Warranty
  - Three years, customer replaceable unit (CRU) and on-site limited warranty on the type 8028 models.
  - One year, customer replaceable unit (CRU) and on-site warranty on type 8014 models.

Table 3-11 on page 264 lists the features of the HS 12 Blade Server.

Table 3-11 Features of the HS12

Feature	8028	8014
Processor options	Intel Xeon 5400 series quad-core Intel Xeon 3300 series quad-core Intel Xeon 3100 series dual-core Intel Core 2 Duo E6405	Intel Celeron® 445 Intel Core 2 Duo E6305
Number of processors (std/max)	1/1	
Front-side bus	1066 MHz or 1333 MHz	1066 MHz
Cache	2 x 4 MB or 2 x 6 MB or 2 MB	512 K or 2 MB
Memory	6 DIMM slots/24 GB maximum	
Memory Type	PC2-5300 VLP 667 MHz high-performance double data rate (DDR2) ECC memory	
Internal hard disk drives (std/max)	0/2 (SAS)	
Internal storage type	► Up to two 2.5” hot swap SAS SFF drives ► Up to two 2.5” hot-swap SFF solid state drives (SSD)	► Up to two 2.5” non-hot-swap SATA drives ► Up to two 2.5” non-hot-swap SFF solid state drives (SSD)
	Optional: SIO Expansion Unit supports 3 additional 2.5” hot-swap SAS drives	
Network	Two ports, Integrated Dual Gigabit Ethernet (Broadcom 5714),	
I/O upgrade	One PCI-X expansion connector and one PCI-Express expansion connector Optional: SIO Expansion Unit: one or two PCI-X slots	
Warranty	3-year, customer replaceable unit (CRU) and on-site limited warranty	1-year, customer replaceable unit (CRU) and on-site limited warranty

Table 3-12 lists the HS12 models that are currently available. Models are available with single core, dual core or quad core processors, and with different clock speeds.

Table 3-12 Models of the HS12

Model	Intel CPU	Cores per socket	CPU/FSB Speed (GHz/MHz)	L2 cache	CPU std/max	Memory std/max	CPU power	SAS or SATA card
Machine type 8028 (3-year warranty)								
8028-23Y	Core2Duo E6405	Dual	2.136/1066	2 MB	1/1	2 GB/24 GB	65 W	SAS
8028-27Y	Xeon E3113	Dual	3.0/1333	6 MB	1/1	2 GB/24 GB	65 W	SAS
8028-44Y	Xeon X3323	Quad	2.5/1333	6 MB	1/1	2 GB/24 GB	80 W	SAS
8028-45Y	Xeon X3353	Quad	2.66/1333	12 MB	1/1	2 GB/24 GB	80 W	SAS
8028-46Y	Xeon X3363	Quad	2.83/1333	12 MB	1/1	2 GB/24 GB	80 W	SAS
8028-4Lx	Xeon L5420	Quad	2.5/1333	12 MB	1/1	2 GB/24 GB	50 W	SAS
Machine type 8014 (1-year warranty)								
8014-1AY	Celeron 445	Single	1.86/1066	512 KB	1/1	512 MB/24 GB	65 W	SATA
8014-2BY	Core2Duo E6305	Dual	1.86/1066	2 MB	1/1	2 GB/24 GB	65 W	SATA

Note: For further information regarding Intel Xeon processor features, refer to:

http://www.intel.com/products/processor/xeon3000

3.3.1 Processor

The blade server HS12 supports the following processors depending on the model:

- One Intel Xeon 5400 series (Harpertown) quad-core processor
- One Intel Xeon 3300 series (Yorkfield) quad-core processor
- One Intel Xeon 3100 series (Wolfdale) dual-core processor
- One Intel Core 2 Duo dual-core processor
- One Celeron single-core processor

See Table 3-12 on page 264 for specifics. The HS12 has only one processor socket.

3.3.2 Memory

There are a total of six DIMM slots in the HS12. The memory used in the blade server HS12 has the following features:

- PC2-5300 667 MHz high-performance double data rate (DDR2)
- Very low profile (VLP) form factor
- ECC
- Chipkill (on DIMM modules 1 GB or larger)

Figure 3-29 shows the physical layout of the memory DIMMs.

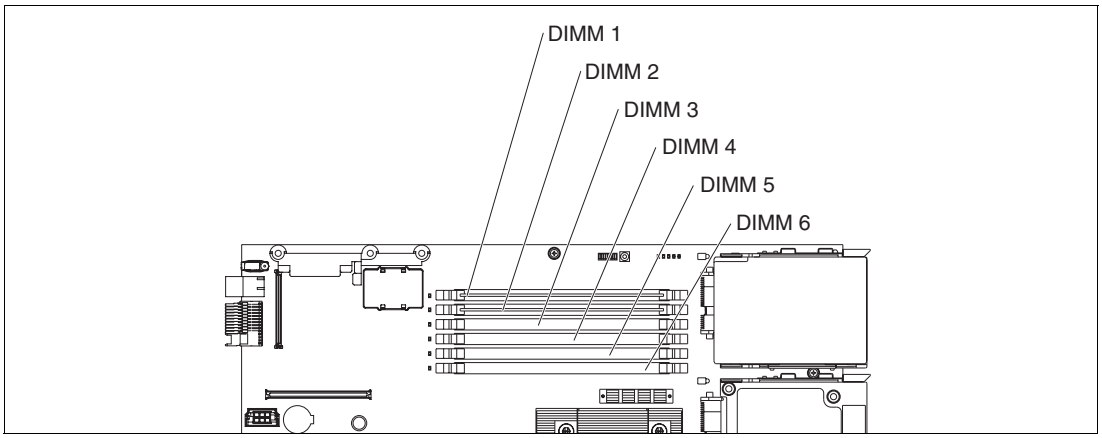

Figure 3-29 DIMM layout of the HS12

Memory configuration must follow these guidelines:

- Memory is 2-way interleaved, so DIMMs should be installed in matched pairs to maximize performance of the memory subsystem.
- All installed DIMMs must be the same speed.
- If you have the ServeRAID MR10ie RAID controller installed then it's battery backup unit is installed in DIMM slot 6. Therefore only five DIMM slots are available for memory. See 3.1.19, "ServeRAID MR10ie (CIOv) Controller" on page 255.

Table 3-13 on page 266 lists the HS12 memory options.

Table 3-13 HS12 memory options

Part number	Feature code[a]	Memory description	Rank
39M5861	0548	1 GB (2x 512 MB) DDR2-RDIMM 667 MHz	Single
39M5864	0549	2 GB (2x 1 GB) DDR2-RDIMM 667 MHz	Single
46C0512	3946	4 GB (2x 2 GB) DDR2-RDIMM 667 MHz	Single
46C0513	3948	8 GB (2x 4 GB) DDR2-RDIMM 667 MHz	Dual

a. Feature codes are for one memory DIMM, not for a pair of DIMMs.

Note: The 4 GB memory DIMM modules in the 8 GB memory kit (46C0513) are dual rank. The other memory DIMMs modules are all single rank.

The HS12 supports two mutually exclusive memory configurations, normal and sparing. The configuration is enabled in BIOS under Advanced Option.

Normal (Flat Mode)

The full capacity of all DIMMs is available to the operating system. In this mode, there are no hot-spare or memory mirroring capabilities.

The configuration rules are as follows:

- DIMMs must be installed in matched pairs
  - DIMMs 1 and 2 must match.
  - DIMMs 3 and 4 must match.
  - DIMMs 5 and 6 must match.
- All DIMMs in a pair or group must be the same size, speed, type, technology, and physical design.

Sparing (online hot spare memory)

This mode provides a degree of memory redundancy. In sparing mode, the trigger for failover is a preset threshold of correctable errors. When this threshold is reached for an active memory rank, the content of that rank is copied to the spare rank. The failed rank of memory is taken offline and the spare rank is put online and used as active memory in place of the failed rank.

The six DIMM slots of the HS12 are divided into two channels of three DIMMs. When online hot-spare memory is enabled using single-rank DIMMs, dual-rank DIMMs, or a combination of both, one rank is set aside per channel as an online spare in case one of the other ranks in that channel fails.

The spare rank must have the capacity of at least that of the largest active rank. (In other words, if a combination of 2 GB and 4 GB DIMMs are used in a channel, then one rank on one 4 GB DIMM per channel will be used for sparing.) In an HS12 with 24 GB installed, up to 20 GB (using 6 x dual rank 4GB DIMMs) of memory is available when the hot-spare feature is active. Sparing is handled at the hardware level; no operating system support is required.

The memory set aside for the hot spare memory is one rank per channel. The size of the rank (and therefore the amount set aside for sparing) varies depending on the DIMMs used.

BladeCenter Memory and I/O (MIO) expansion unit

The HS12 does not support the attachment of the BladeCenter Memory and I/O (MIO) expansion unit.

3.3.3 Onboard network adapters

The HS12 blade server includes an integrated Broadcom BCM5714S Gigabit Ethernet controller. This controller has two Gigabit Ethernet interfaces. These two Ethernet interfaces are routed through to the redundant midplane, providing a foundation for fault-tolerant network setups.

Note: When installed in a BladeCenter E, H, T or HT chassis, Ethernet Port 1 on the blade server is routed to the switch module in Bay 1 and Ethernet port 2 is routed to the switch module in Bay 2.

When installed in BladeCenter S, both Ethernet ports on the blade server are routed to the switch module in Bay 1.

The following features are supported by the Broadcom (BCM5714S) onboard controller:

- iSCSI initiator
- PXE 2.0 boot support
- Wake on LAN
- Load-balancing
- Jumbo frame support (9 KB)

The Broadcom Advanced Server Program (BASP) is management software provided by Broadcom for managing your network cards. This includes a command line interface for scripting purposes.

3.3.4 Integrated systems management processor

The blade server HS12 has an integrated baseboard management controller (BMC) that provides:

- Intelligent Platform Management Interface (IPMI) 1.5 compliance
- Monitoring of system and battery voltage, system temperature, fans, power supplies, and processor
- DIMM status
- Fan speed control
- Product ID and Family ID detection
- Highly secure remote power on/off
- System reset control
- NMI/SMI detection and generation
- System diagnostic LED control (power, HDD, activity, alerts, heartbeat)
- IPMI over LAN
- LAN messaging and alerting

Note: You cannot access this onboard BMC externally through a dedicated port or daisy-chain it to other servers or blades. The BMC is only used for the forwarding of information to the management module within the BladeCenter chassis. The management module is a single point of control for hardware management of all server and switch components contained within the chassis.

3.3.5 Local storage options

Depending on the model of HS12, it will have either a SAS interface card or a SATA interface card installed on a dedicated connector in the server.

- HS12 type 8028 has a SAS interface card installed.
- HS12 type 8014 has a SATA interface card installed.

See Table 3-12 on page 264 for specifics. The interface card comes fitted from the factory and cannot be ordered separately.

Figure 3-30 shows the SAS interface card installed on the HS12 type 8028.

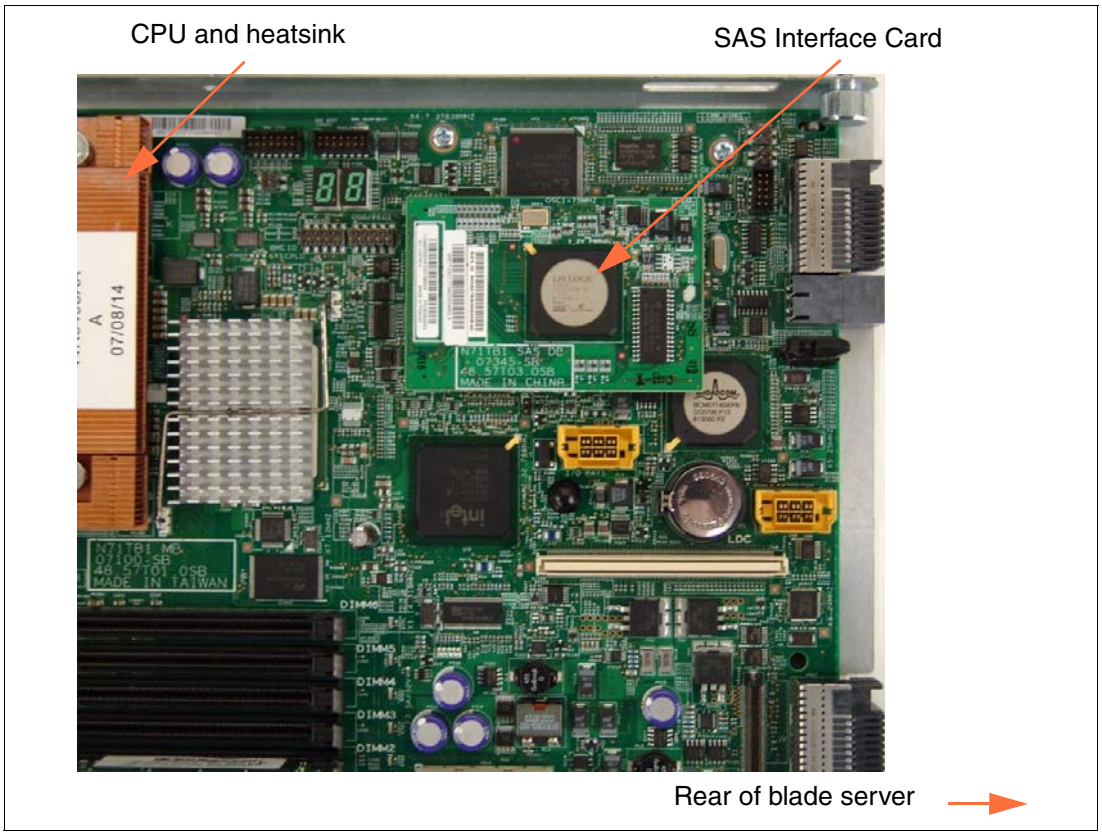

Figure 3-30 SAS Interface Card on a HS12 type 8028

Local drives are 2.5” small form factor (SFF) drives. Drive support is as shown in Table 3-14 on page 269.

Table 3-14 Hard disk drive support on the HS12

Machine type	SAS	Solid state	SATA
8028	Supported (hot-swap)	Supported (hot-swap)	Supported (hot-swap)
8014	No	Supported (non-HS)	Supported (non-HS)

Storage controllers

The SAS interface card (used in machine type 8028) features an onboard LSI Logic 53C1064E SAS controller with two disk drive interfaces. This provides connection for either up to two hot-swap SAS drives or up to two hot-swap solid state drives. The LSI Logic 53C1064E SAS controller provides RAID 0 or RAID 1 capability. The SAS interface card can interface with SAS or SATA type disk drives.

The SATA interface card (supplied in machine type 8014) does not have a controller but instead passes through the SATA interface from the ICH9 Southbridge chip on the motherboard of the blade server. This provides connection for up to two non-hot-swap 2.5-inch drives, either two SATA drives or two solid state drives. The SATA interface card does not provide any RAID capability. The SATA interface card can interface only with SATA or SSD disk drives.

The HS12 type 8028 supports the installation of the ServeRAID MR10ie (CIOv) controller. The ServeRAID MR10ie is a SAS RAID Controller I/O Expansion Card for a blade server. It provides RAID support with battery-backed cache for onboard and external disk drives. The ServeRAID MR10ie (CIOv) Controller I/O Expansion Card has a connector to attach it to the paddle card socket on a HS12 type 8028 blade server. The ServeRAID MR10ie comes with an interposer that allows the connector to reach the socket on the HS12 blade server. Figure 3-22 on page 255 shows the interposer location.

The ServeRAID MR10ie ships with a battery backup unit that installs in the sixth DIMM slot on the HS12 type 8028. When the battery backup unit is installed in the sixth DIMM slot, only the five remaining DIMM slots can be used for memory. A thin ribbon cable is included that connects the battery to the ServeRAID MR10ie (CIOv) Controller.

When the ServeRAID MR10ie is used in the HS12 to connect to external disk drives, it also requires the installation of the SAS Connectivity Card (part number 43W3974). The SAS Connectivity Card passes the signal from the ServeRAID MR10ie via SAS switch modules in bays 3 and 4 and onto external disk drives. These external disk drives may reside in the BladeCenter S chassis or in an EXP3000 expansion unit attached to the external ports of the SAS switch modules.

Figure 3-31 on page 270 shows the ServeRAID MR10ie (CIOv) Controller installed on the HS12 type 8028.

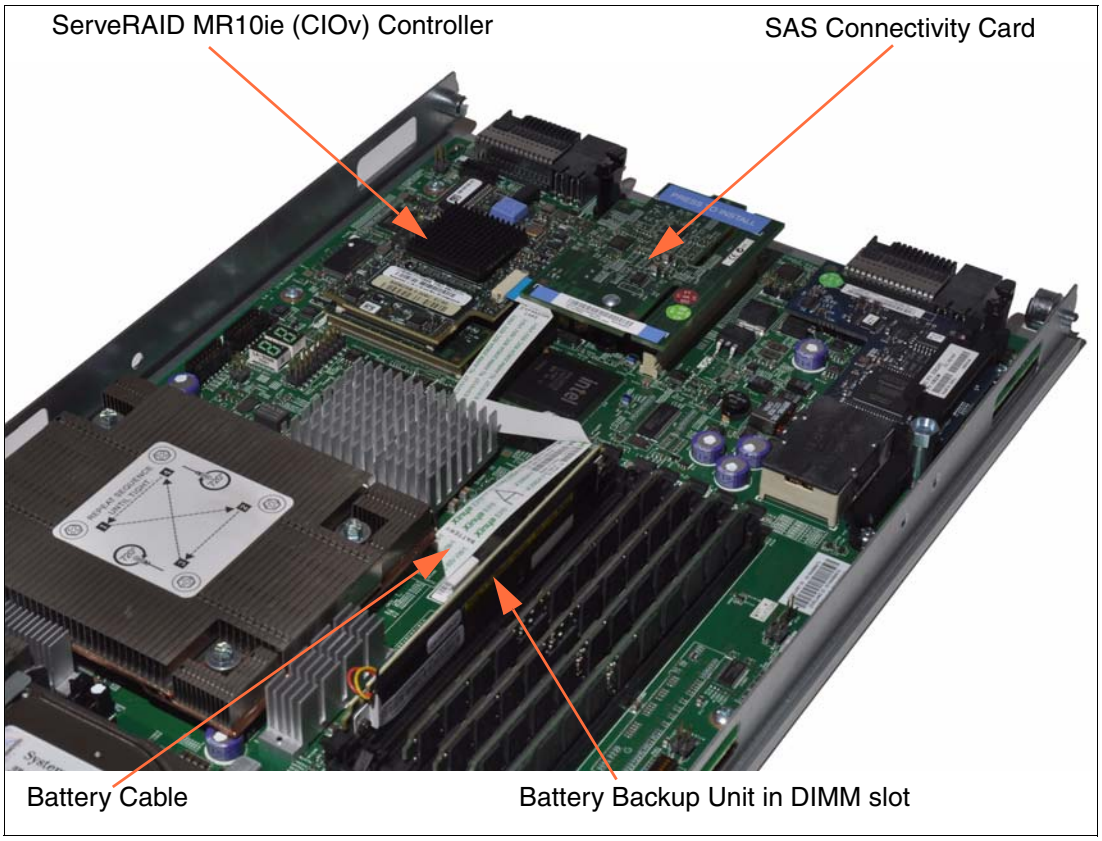

Figure 3-31 The HS12 type 8028 showing the ServeRAID MR10ie installed

For further information, refer to 3.1.19, “ServeRAID MR10ie (CIOv) Controller” on page 255.

Local disk drives for HS12 type 8028

Table 3-15 lists the disk drive options for the HS12 server, machine type 8028. The listed SAS and SATA hard disk drives are mounted on a small form factor (SFF) hot-swap tray and they are installed into the front of the blade server.

Note: A SAS or SATA hard disk drive cannot be part of a RAID array with a solid state drive.

Table 3-15 Small form factor hard disk drive options for HS12

Part number	Feature code	Description	Compatible with
43W7535	5577	IBM 73 GB 10K SAS 2.5” SFF HS	8028
43W7536	5578	IBM 146 GB 10K SAS 2.5” SFF HS	8028
43W7545	5579	IBM 73 GB 15K SAS 2.5” SFF HS	8028
43W7670	5595	300 GB 10K SATA 2.5" Slim-HS HDD	8028

Figure 3-32 on page 271 shows the hot-swap disk drive bays on the front of the HS12 type 8028.

Figure 3-32 Showing the hot swap disk drive bays on the front of the HS12 type 8028

Local disk drives for HS12 type 8014

Table 3-16 lists the hard disk drive options for the HS12 server type 8014. These are all non-hot-swap because the 8014 does not support hot-swap drives.

Table 3-16 Hard disk drive options for HS12 type 8014

Part number	Feature code	Description	Compatible with
43W7640	5575	IBM 80 GB 7200 SATA 2.5” non-hot-swap	8014

Solid state drives

The HS12 supports the installation of solid state drives (SSD). The solid state drive supported can be either hot-swap or non-hot-swap, depending on the type of HS12.

- HS12 type 8028 supports hot-swap SSD drives only.
- HS12 type 8014 supports non-hot-swap SSD drives only.

These SSD devices are designed for operating system boot and low-bandwidth data transmission. SSD devices use less power than a standard SAS drive and therefore also generate less heat.

Table 3-17 on page 272 lists the SSD options for the HS12.

Table 3-17 Solid state drive options for the HS12

Part number	Feature code	Description	Compatible with
43W7648	5576	IBM 31.4 GB SATA 2.5” hot-swap SSD[a]	8028
43W7614	5557	IBM 15.8 GB 2.5” non-hot-swap SSD	8014
43W7618	5563	IBM 31.4 GB 2.5” non-hot-swap SSD	8014

a. This drive is a solid state drive with a SATA interface. However, because it is in a hot-swap carrier, it is only supported on the 8028.

Figure 3-33 shows the 2.5” non-hot swap 15.8 GB solid state drive.

Figure 3-33 A 2.5” non-hot swap 15.8 GB Solid State Drive option

IBM Modular Flash Drives

The HS12 does not have a connector on the blade server motherboard to attach the Modular Flash Drive.

BladeCenter Storage and I/O Expansion Unit (SIO)

The HS12 supports the attachment of one BladeCenter Storage and I/O expansion unit (SIO). The Storage and I/O (SIO) expansion unit provides the HS12 blade server with an additional three 2.5” SAS hot swap HDD bays and battery-backed RAID support, as well as two extra PCI-X I/O expansion sockets.

With the SIO installed, the ServeRAID 8k-l (or ServeRAID 8k, if installed) cannot see the SAS disk drives on the attached blade server. This means that you can only form RAID arrays using the three disk drives in the SIO with the ServeRAID 8k-l or ServeRAID 8k. The disk drives on the attached blade server can still be installed, but they will be controlled separately by the blade server controller.

Note: The drive trays on the disk drives that can be installed in the SIO are different from the drive trays that can be installed in the HS12. Therefore, the disk drives in the SIO unit cannot be interchanged with the disk drives in the HS12.

The HS12 does not support the attachment of any of the previous BSE-I or BSE-II expansion units.

For more information about the SIO, see 3.18.2, “Storage and I/O Expansion Unit” on page 363.

3.3.6 I/O expansion options

The blade server HS12 has two onboard connectors for various I/O expansion card options, a PCI-X connector, and a PCI Express connector.

The PCI-X connector allows the connection of *one* of these expansion cards:

- A Small Form Factor expansion card
- A CFFv form factor expansion card

The installed expansion card is connected to module bays 3 and 4 at the rear of the BladeCenter chassis via the internal midplane. The interface of this I/O connector is a PCI-X 133 MHz bus. The HS12 does not support the installation of a Standard Form Factor (StFF) expansion card because this type of expansion card would interfere with the SAS or SATA Interface Card.

The HS12 also has a PCI Express connector that allows for the attachment of *one* of the following expansion cards:

- A CFFh expansion card
- A HSFF expansion card
- The Storage and I/O (SIO) Expansion Unit

See 3.1, “Expansion cards” on page 234, for further information about this topic.

Tip: If a CFFh or HSFF expansion card is installed in a blade server and you want to install an SIO Expansion Unit, you need to first remove the expansion card, attach the SIO Expansion Unit, and then install the CFFh or HSFF expansion card in the SIO Expansion Unit instead.

Table 3-18 summarizes the available I/O expansion card options for the HS12.

Table 3-18 Blade server HS12 I/O expansion card options

Part number	Feature code	Description	Form factor	Page
39R8624	1548	BladeCenter SFF Gigabit Ethernet Expansion Card	SFF	243
43W6859	2994	Emulex 4 Gb Fibre Channel Expansion Card	CFFv	250
32R1760	1466	Cisco Systems InfiniBand 4X HCA Expansion Card	HSFF	251
39Y9271	2967	NetXen 10 Gb Ethernet Expansion Card	CFFh	243
44W4465	5479	Broadcom 10 Gb Ethernet CFFh Expansion Card	CFFh	245
44W4466	5489	Broadcom 2-Port 10 Gb Ethernet CFFh Expansion Card	CFFh	244
39Y9306	2968	QLogic Ethernet and 4 Gb FC Expansion Card	CFFh	246
44X1940	5485	QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)	CFFh	247
39Y9310	2969	Ethernet Expansion Card	CFFv	242
44W4479	5476	2/4 Port Ethernet Expansion Card (CFFh)	CFFh	242
41Y8527	2970	QLogic 4 Gb Fibre Channel Expansion Card	CFFv	248
39Y9190	2979	SAS Expansion Card	CFFv	252
43W3974	1591	SAS Connectivity Card	CFFv	254

Part number	Feature code	Description	Form factor	Page
43W4420	2993	Voltaire InfiniBand 4X DDR Expansion Card	CFFh	251
43W4421	None	Cisco InfiniBand 4X DDR Expansion Card	CFFh	251
43W4423	None	InfiniBand 4X DDR Expansion Card	CFFh	251
46C7167[a]	5752	ServeRAID-MR10ie (CIOv) Controller	CIOv	255

a. This expansion card is only supported on the HS12 type 8028.

PCI Express I/O Expansion Unit (PEU3e)

The HS12 type 8028 supports the attachment of one PCI Express I/O Expansion Unit. The PEU3e allows standard PCI-Express type adapters to be attached to the HS12 type 8028 only. The HS12 type 8014 does not support the attachment of the PEU3e.

For more information about the PEU3e, see 3.18.5, “PCI Express I/O Expansion Unit (PEU3e)” on page 368.

3.4 HS21 XM type 7995

The BladeCenter HS21 XM (XM stands for extended memory) servers are positioned as high-density servers. These servers can have up to two dual-core or quad-core Intel processors. Plus, with 8 DIMM sockets for up to 64 GB of DDR2 FB-DIMM memory, this makes the HS21 XM the perfect blade server platform for processor and memory-intensive workloads such as virtualization.

Integrated dual Gigabit Ethernet controllers are standard, providing high-speed data transfers and offering TCP Offload Engine (TOE) support, load-balancing, and failover capabilities. Each blade can also connect to additional Ethernet, Myrinet, Fibre Channel, iSCSI, InfiniBand, and other high-speed communication switches housed in the chassis with the use of optional expansion cards.

This blade is designed with power management capability to provide the maximum uptime possible for your systems. In extended thermal conditions or power brownouts, rather than shut down completely, or fail, the HS21 XM automatically reduces the processor frequency to maintain acceptable thermal and power levels.

All HS21 XM models also include support for one SAS hard disk drive or one dual Solid State Disk drive and one modular flash drive. The optional Storage and I/O (SIO) Expansion Unit can connect to a HS21 XM blade server to provide an additional three 2.5 SAS HDDs with hot-swap support, optional RAID-5 with battery-backed cache, and four additional communication ports. See 3.18.2, “Storage and I/O Expansion Unit” on page 363, for details.

There is also a model of HS21 XM type 7995 available with the VMware ESXi 3.5 embedded hypervisor preloaded on a modular flash drive.

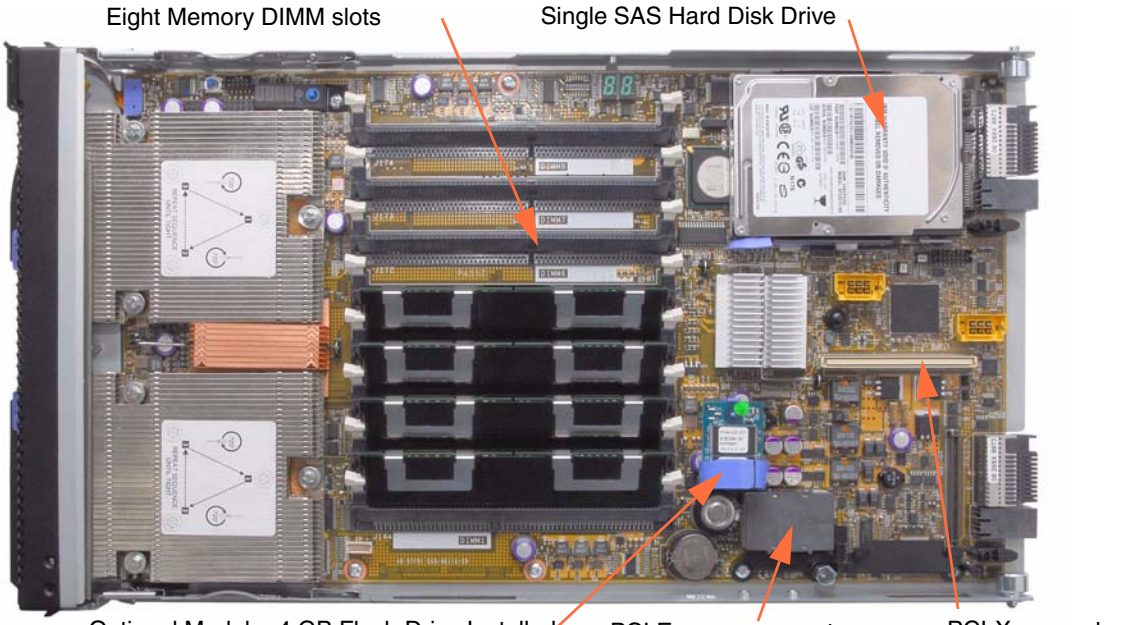

Figure 3-34 The HS21 XM

Features of the HS21 XM include:

- Two processor sockets supporting either dual-core Intel Xeon 5100 & 5200 series processors or quad-core Intel Xeon 5300 and 5400 series processors.
- Up to 64 GB of system memory in 8 DIMM sockets with 8 GB DIMMs.
- One internal 2.5-inch SAS HDD with support for the Storage and I/O (SIO) expansion unit blade with an additional three 2.5-inch SAS HDD bays and RAID support.
- An optional Modular Flash Drive can be used in place of, or in addition to, the internal SAS HDD, as a boot device (only Linux supports the Modular Flash Drive as a boot device).
- Two Gigabit Ethernet ports standard; plus more, using either a 2-port Gigabit Ethernet expansion card or a PCI I/O Expansion Unit II.
- Integrated Baseboard Management Controller (BMC) service processor to monitor server availability, perform Predictive Failure Analysis, and so on, and trigger IBM Director alerts.
- Support for IBM PowerExecutive™ 2.0, software designed to take advantage of new system power management features that monitor actual power usage and provide power consumption capping features.
- Concurrent keyboard, video, mouse (cKVM) support with the addition of the optional IBM BladeCenter Concurrent KVM feature card.
- Model available with the VMware ESXi 3.5 embedded hypervisor (see Table 3-20 on page 276)
- Three-year, on-site limited warranty

Table 3-19 Features of the HS21 XM

Feature	
Processor options	Intel Xeon 5400 Series quad-core processors Intel Xeon 5300 Series quad-core processors Intel Xeon 5200 Series dual-core processors Intel Xeon 5100 Series dual-core processors
Number of processors (std/max)	1 / 2
Front-side bus	1033 MHz or 1333 MHz
Cache	4 MB, 8 MB or 12 MB L2 cache (shared between both cores)
Memory	8 DIMM slots/64 GB maximum
Memory Type	PC2-5300 (667 MHz) Fully Buffered DDR II ECC
Internal hard disk drives (std/max)	0/1 (SAS or SATA)
Maximum internal storage	Optional: One 2.5” Non-Hot Swap SAS or SATA HDD onboard Optional: One modular flash drive onboard Optional: One single or one dual Solid State Drive Optional: SIO Expansion Unit supports 3 additional 2.5” hot-swap SAS drives
Network	Two ports, Integrated Dual Gigabit Ethernet (Broadcom 5708S), TOE
I/O upgrade	One PCI-X expansion connector and one PCI-Express expansion connector Optional: SIO Expansion Unit: one or two PCI-X slots

Table 3-12 on page 264 lists the different models of the HS21 XM that are currently available. Models are available with dual-core or quad-core processors and with different clock speeds. Some models are available with low power processors.

Table 3-20 Models of the HS21 XM

Model	Intel CPU	Cores per socket	CPU/FSB Speed (GHz/MHz)	L2 cache	Processors Std/Max	Memory Std/Max	Processor power
7995-L4Y	Xeon 5140	Dual core	2.33/1333	4 MB	1/2	1 GB/64 GB	65 W
7995-NTY	Xeon 5138 LV	Dual core	2.13/1066	4 MB	1/2	1 GB/64 GB	35 W
7995-H1Y	Xeon 5148 LV	Dual core	2.33/1333	4 MB	1/2	1 GB/64 GB	40 W
7995-R2Y	Xeon X5260	Dual core	3.33/1333	6 MB	1/2	1 GB/64 GB	80 W
7995-RLY	Xeon L5240	Dual core	3.0/1333	6 MB	1/2	1 GB/64 GB	40 W
7995-C2Y	Xeon E5345	Quad core	2.33/1333	8 MB	1/2	1 GB/64 GB	80 W
7995-GNY	Xeon L5408	Quad core	2.13/1066	12 MB	1/2	4 GB/64 GB	40 W
7995-GLY	Xeon L5420	Quad core	2.5/1333	12 MB	1/2	1 GB/64 GB	50 W
7995-G1Y	Xeon E5405	Quad core	2.0/1333	12 MB	1/2	1 GB/64 GB	80 W
7995-G3Y	Xeon E5420	Quad core	2.5/1333	12 MB	1/2	1 GB/64 GB	80 W

Model	Intel CPU	Cores per socket	CPU/FSB Speed (GHz/MHz)	L2 cache	Processors Std/Max	Memory Std/Max	Processor power
7995-G4Y	Xeon E5430	Quad core	2.66/1333	12 MB	1/2	1 GB/64 GB	80 W
7995-G5Y	Xeon E5440	Quad core	2.83/1333	12 MB	1/2	1 GB/64 GB	80 W
7995-G6Y	Xeon E5450	Quad core	3.0/1333	12 MB	1/2	1 GB/64 GB	80 W
7995-G7Y	Xeon X5460	Quad core	3.16/1333	12 MB	1/2	1 GB/64 GB	120 W
7995-HVY[a]	Xeon E5450	Quad core	3.0/1333	12 MB	2/2	4 GB/64 GB	80 W

a. The 7995-HVY is the VMware ESXi 3.5 Embedded Hypervisor model.

Note: For further information regarding Intel Xeon processor features, refer to:

http://www.intel.com/products/processor/xeon5000

3.4.1 Processor

The blade server HS21 XM supports either:

- Up to two Intel Xeon 5400 series (Harpertown) quad-core processors
- Up to two Intel Xeon 5300 series (Clovertown) quad-core processors
- Up to two Intel Xeon 5100 series (Woodcrest) dual-core processors
- Up to two Intel Xeon 5200 series (Wolfdale) dual-core processors

These Intel processors are available in various clock speeds and have standard and low power offerings. The processors support the following Intel technologies:

- Intel 64 Technology (EM64T)
- Intel Virtualization Technology (Intel VT)
- Execute Disable Bit

In a two-socket configuration, both processors must be identical.

Table 3-21 lists the blade server HS21 XM processor options.

Table 3-21 HS21 processor options

Server model	CPU option part number	Feature code	CPU model	CPU specifications (speed/FSB speed/ L2 cache/power)
Intel Xeon Dual Core processors				
7995-L4Y	42C0565	1420	Xeon 5140	2.33 GHz/1333 MHz, 4 MB L2/65w
7995-NTY	42C0568	1423	Xeon 5138 LV	2.13 GHz/1066 MHz/4 MB L2/35w
7995-H1Y	42C0569	1427	Xeon 5148 LV	2.33 GHz/1333 MHz/4 MB L2/40w
7995-R2Y	43W4036	3699	Xeon X5260	3.33 GHz/1333 MHz/6 MB L2/80w
7995-RLY	43C5051	None	Xeon L5240	3.0 GHz/1333 Mhz/6 MB L2/40w
Intel Xeon Quad Core processors				
7995-C2Y	42C0570	1446	Xeon E5345	2.33 GHz/1333 MHz/8 MB L2/80w
7995-G1Y	43W3991	3674	Xeon E5405	2.0 GHz/1333 MHz/12 MB L2/80w

Server model	CPU option part number	Feature code	CPU model	CPU specifications (speed/FSB speed/ L2 cache/power)
7995-G3Y	43W3993	3675	Xeon E5420	2.5 GHz/1333 MHz/12 MB L2/80w
7995-G4Y	43W3994	3365	Xeon E5430	2.66 GHz/1333 MHz/12 MB L2/80w
7995-G5Y	43W3995	3677	Xeon E5440	2.83 GHz/1333 MHz/12 MB L2/80w
7995-G6Y	43W3996	3367	Xeon E5450	3.0 GHz/1333 MHz/12 MB L2/80w
7995-G7Y	43W3997	None	Xeon X5460	3.16 GHz/1333 MHz/12MB L2/120w
7995-GLY	46C5050	3614	Xeon L5420	2.5 GHz/1333 MHz/12 MB L2/50w
7995-GNY	46C5094	6959	Xeon L5408	2.13 GHz/1066 MHz/12 MB L2/40w

3.4.2 Memory

The memory used in the blade server HS21 XM has the following features:

- PC2-5300 (667 MHz) Fully Buffered DDR-2 DIMM
- ECC
- Chipkill

Figure 3-35 shows the physical layout of the memory DIMMs.

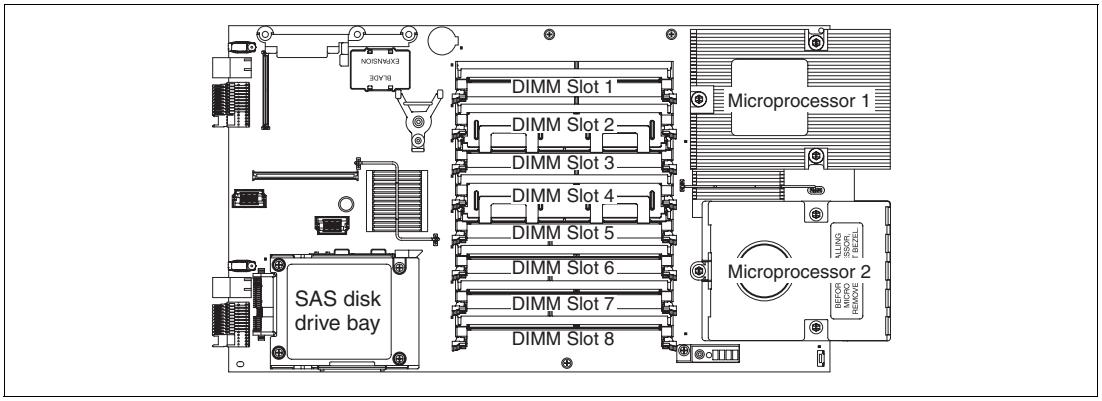

Figure 3-35 DIMM layout of the HS21 XM

Memory configuration has to follow these guidelines:

- There are a total of eight DIMM slots.
- Memory is 2-way interleaved, so DIMMs must be installed in matched pairs.
- A configuration of six DIMMs (three pairs) is not supported.
- All installed DIMMs must be the same speed.

The logical memory DIMM configuration is shown in Figure 3-36 on page 279.

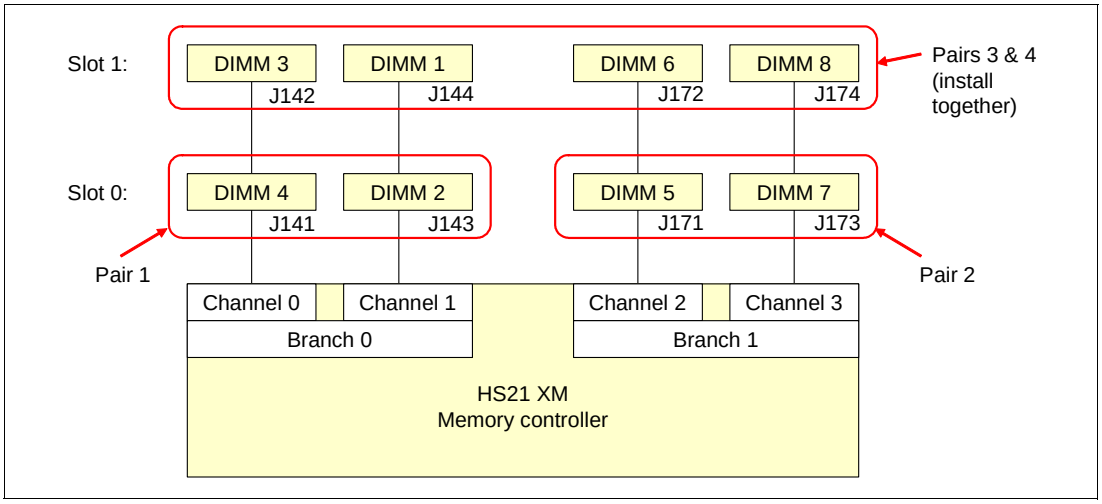

Figure 3-36 HS21 XM memory configuration

The standard HS21 XM models have two 512 MB DIMMs installed into sockets 2 and 4, as shown in Figure 3-36. The installation order is as follows:

1. The first pair is installed in sockets 2 and 4.
2. The second pair of DIMMs is installed in sockets 5 and 7.
3. The third and fourth pairs are installed together at the same time in sockets 1, 3, 6, and 8.

See Table 3-13 on page 266 for blade server HS21 XM memory options. The number of ranks per DIMM is important when determining the space set aside for the hot spare memory in sparing mode.

Table 3-22 HS21 XM memory options

Part number	Memory description	Ranks per DIMM[a]
39M5782	1 GB (2x512 MB) PC2-5300 667 MHz ECC Chipkill DDR2 FB-DIMM	Single (512 MB rank)
39M5785	2 GB (2x1 GB) PC2-5300 667 MHz ECC Chipkill DDR2 FB-DIMM	Dual (512 MB rank)
39M5791	4 GB (2x2 GB) PC2-5300 667 MHz ECC Chipkill DDR2 FB-DIMM	Dual (1GB rank)
39M5797	8 GB (2x4 GB) PC2-5300 667 MHz ECC Chipkill DDR2 FB-DIMM	Dual (2 GB rank)
46C7418	2 GB (2x1GB) Single Rank PC2-5300 CL5 ECC FBD 667MHz Low Power Memory	Single
46C7419	4 GB (2x2GB) Dual Rank PC2-5300 CL5 ECC FBD 667MHz Low Power Memory	Dual
46C7420	8 GB (2x4GB) Quad Rank PC2-5300 CL5 ECC FBD 667MHz Low Power Memory	Quad
46C7577	16 GB (2x8GB) PC2-5300 CL5 ECC FBDIMM 667Mhz AMB+	Quad

a. This is the number of ranks per DIMM. You can mix DIMMs of different rank counts, subject to the configuration rules described here.

The HS21 XM supports three mutually exclusive memory configurations: normal, mirroring, and sparing. The configuration is enabled in BIOS under Advanced Option.

Normal (Flat Mode)

The full capacity of all DIMMs is available to the operating system. However, there are no hot-spare or memory mirroring capabilities.

Configuration rules:

- The above installation sequence applies.
- DIMMs must be installed in matched pairs:
  - DIMMs 2 and 4 must match.
  - DIMMs 5 and 7 must match.
  - DIMMs 3 and 1 must match.
  - DIMMs 6 and 8 must match.

Memory mirroring (Mirroring Mode)

In mirroring mode, the server maintains two identical copies of all data in memory. The contents of branch 0 (see Figure 3-36 on page 279) are duplicated in the DIMMs of branch 1. In the event of an uncorrectable error in one of the copies, the system can retrieve the mirrored copy of the data.

The use of memory mirroring means that only half the installed memory is available to the operating system. Therefore if 8 GB of memory is installed with mirroring enabled (for example, using four 2 Gb DIMMs), only 4 GB is visible to the operating system.

Configuration rules:

- To use memory mirroring, both branches (see Figure 3-36 on page 279) must be populated using identical DIMMs. This means that:
  - DIMMs 2, 4, 5, and 7 (the row of DIMMs in Figure 3-36 on page 279 labelled Slot 0) must match.
  - DIMMs 1, 3, 6, and 8 (top row labelled Slot 1) must match.
  - The top row and bottom row of DIMMs do not have to match.
- If you have eight DIMMs installed, the DIMMs in each channel can be different sizes. For example, you can enable mirroring with 4 GB DIMMs in the row labelled Slot 0 and 2 GB DIMMs in the row labelled Slot 1
- If you have eight DIMMs installed, the rows of DIMMs (slots) can be different rank technology. For example, the four DIMMs in the bottom row (slot 0) can be all single-rank DIMMs (512 MB each), while the four DIMMs in the top row (slot 1) can all be dual-rank DIMMs (1 GB, 2 GB, 4 GB).

Hot spare memory (Sparing Mode)

This mode provides a degree of memory redundancy but not the extent of memory mirroring. It also leaves relatively more memory to the operating system than mirroring.

In sparing mode, the trigger for failover is a preset threshold of correctable errors. When this threshold is reached for an active memory rank, the content of that rank is copied to the spare rank. The failed rank of memory is taken offline and the spare rank is put online and used as active memory in place of the failed rank.

The memory set aside for the spare memory is one rank per channel (see Figure 3-36 on page 279). The size of the rank (and therefore the amount set aside for sparing) varies depending on the DIMMs used, as listed in Table 3-13 on page 266.

Configuration rules:

- DIMMs must be installed in matched pairs and in the same installation order as described above.
- All DIMMs in a row (slot 0 or slot 1, as shown in Figure 3-36 on page 279, must match). For example, if you have four DIMMs installed (in sockets 2, 4, 5, and 7), they must be identical.
- Both DIMMs in a channel must be single-rank, dual-rank, or quad-rank. They cannot be mixed. This effectively means, for example, that you cannot mix dual-rank DIMMs (1, 2, and 4 GB DIMMs, as listed in Table 3-13 on page 266) and single-rank DIMMs (512 MB DIMMs) in the server if you want to enable sparing.
- You can mix different sizes of dual-rank DIMMs (1 GB, 2 GB, and 4 GB) and still enable sparing. The rank set aside for sparing will be the largest sized rank. For example, if you have four 2 GB DIMMs (dual-ranked meaning the rank size is 1 GB) and four 4 GB DIMMs (dual-ranked meaning the rank size is 2 GB), then the space reserved for sparing will be a 2 GB rank per channel (a total of 8 GB).

Figure 3-37 on page 282 shows some example configurations that are supported.

Note: The examples in Figure 3-37 on page 282 show that a specific rank is reserved for sparing. This is for illustrative purposes and is not meant to indicate that the rank furthest away from the controller is the one always used.

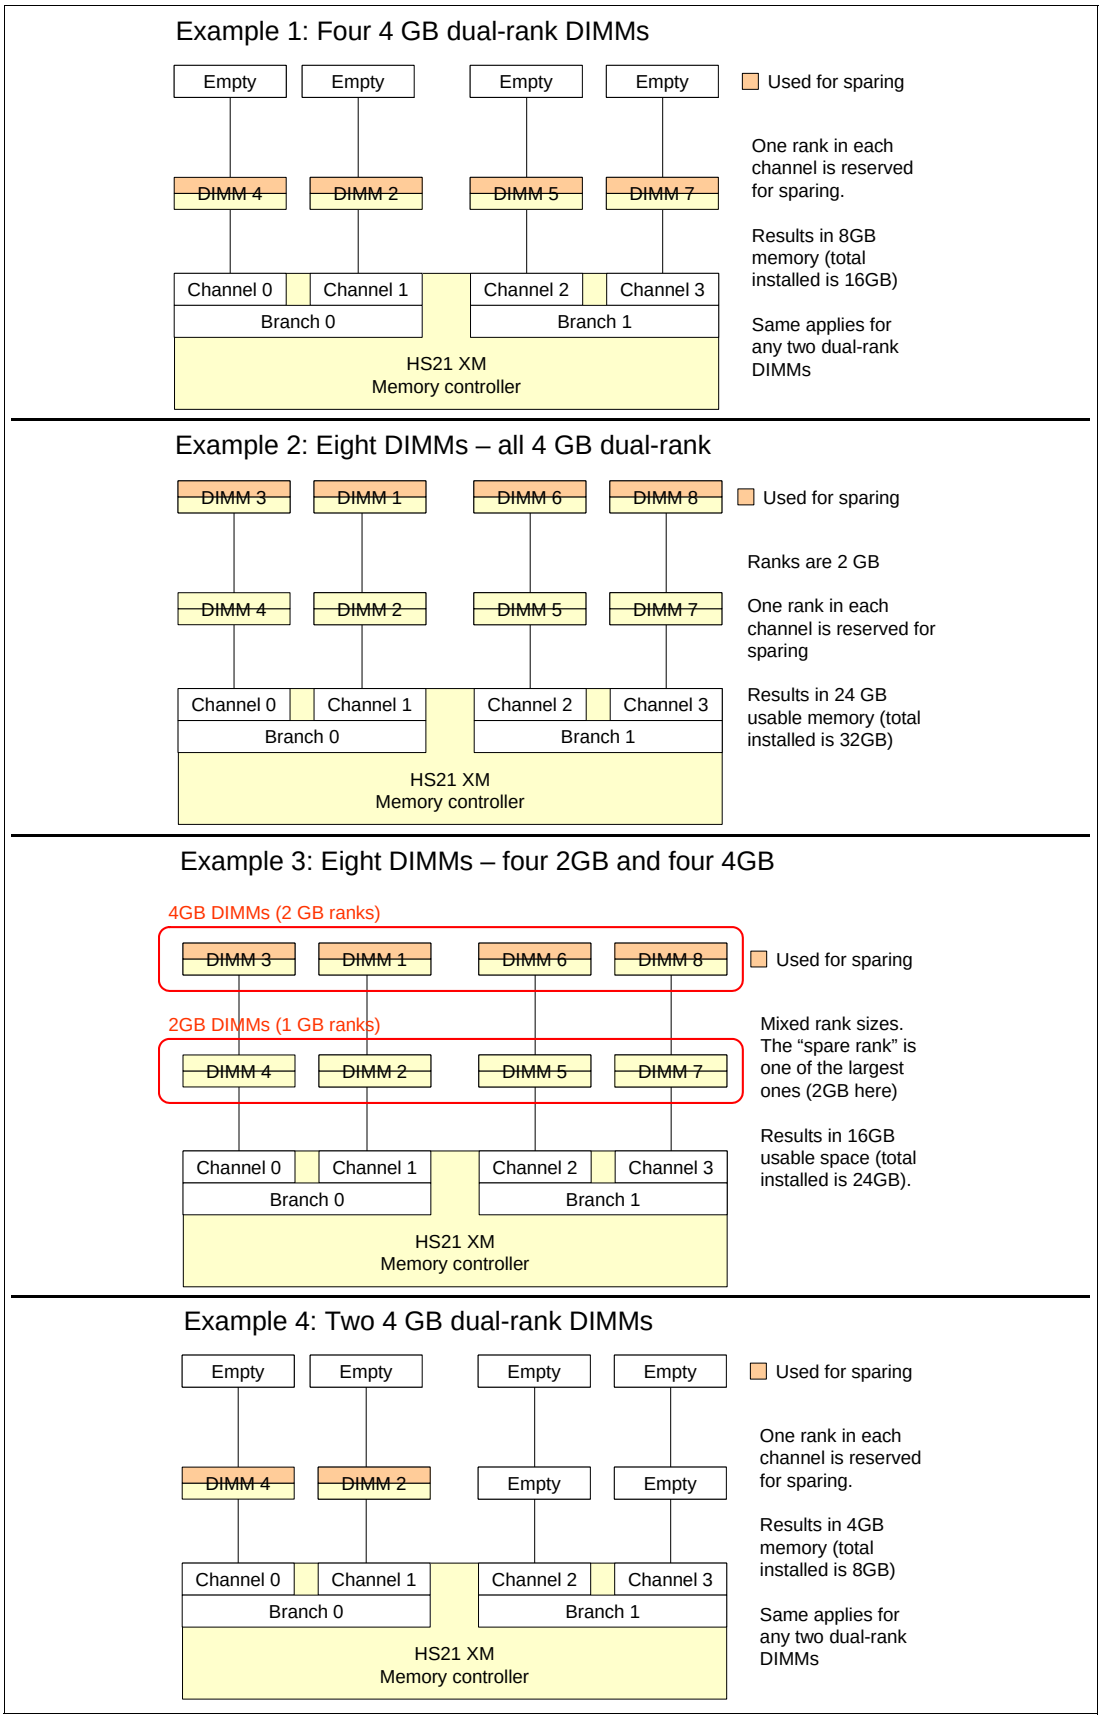

Figure 3-37 Example memory sparing configurations that are supported

Figure 3-38 shows configurations that are *not* supported.

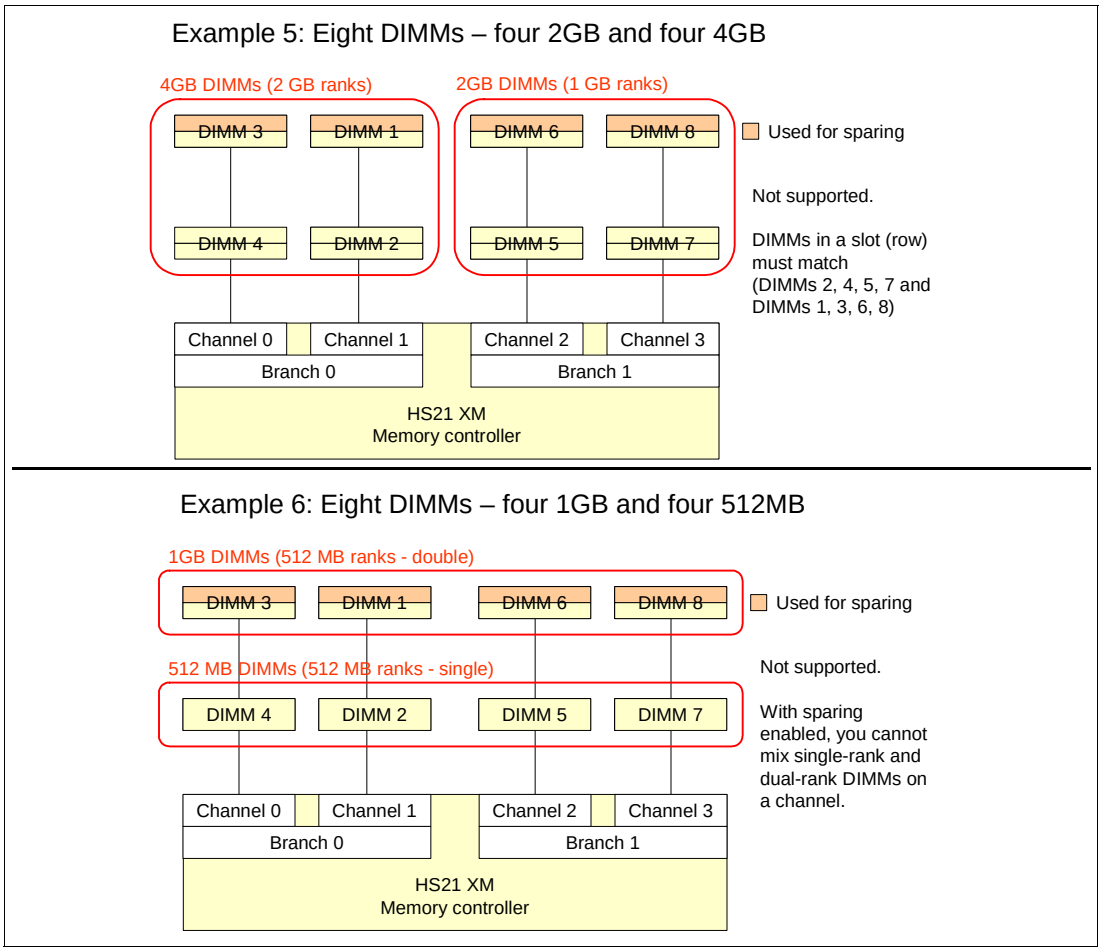

Figure 3-38 Memory configurations not supported

BladeCenter Memory and I/O (MIO) expansion unit

The HS21 XM does not support the attachment of the BladeCenter Memory and I/O (MIO) expansion unit.

3.4.3 Onboard network adapters

The HS21 XM blade server includes two integrated Broadcom BCM5708S Gigabit Ethernet controllers. These two Ethernet interfaces are routed through to the redundant midplane, providing a foundation for fault-tolerant network setups.

Note: When installed in a BladeCenter E, H, T and HT chassis, Ethernet Port 1 on the blade server is routed to the switch in module bay 1 and port 2 is routed to the switch in module bay 2.

When installed in BladeCenter S, both Ethernet ports on the blade server are routed to the switch in module bay 1.

The following features are supported by the Broadcom (BCM5708S) onboard controllers:

- TCP Offload Engine (TOE)
- iSCSI initiator

- PXE 2.0 boot support
- Wake on LAN
- Load-balancing
- VLAN tagging support
- Jumbo frame support (9 KB)

The Broadcom Advanced Server Program (BASP) is management software provided by Broadcom for managing your network cards. This includes a command-line interface for scripting purposes.

3.4.4 Integrated systems management processor

The blade server HS21 XM has an integrated baseboard management controller (BMC) that provides:

- Intelligent Platform Management Interface (IPMI) compliance
- Serial over LAN (SOL)
- Wake on LAN (WOL)
- Power management functions
- KVM and media tray ownership
- Light path diagnostics support
- Automatic BIOS recovery (ABR)
- Automatic server restart (ASR)
- Predictive failure analysis (PFA) on system components (processor, memory, and hard disk drives)
- Inventory
- Error logging
- Environmental monitoring (voltages and temperature on the system board of the blade server)

Note: You cannot access this onboard BMC externally through a dedicated port or daisy-chain it to other servers or blades. The BMC is only for the forwarding of information to the Management Module within the BladeCenter chassis. The management module is a single point of control for hardware management of all server and switch components contained within the chassis.

3.4.5 Local storage options

The HS21 XM blade server features an onboard LSI Logic 53C1064E SAS controller with one dual path SAS connector. This provides connection for either one non-hot swap 2.5-inch SAS SFF hard disk drive, one non-hot swap solid state drive, or one non-hot swap dual solid state drive. The LSI Logic 53C1064E SAS controller provides RAID 0 or RAID 1 capability, which is only useful when using the dual solid state drive (or external storage).

Local SAS or SATA disk drives

The HS21 XM supports the attachment of one optional non-hot swap 2.5-inch Small Form Factor SAS or SATA hard disk drive. It is not possible to create a RAID array with this single SAS or SATA disk drive and any other disk drive. A SAS or SATA hard disk drive cannot be installed at the same time as a Solid State Drive (SSD).

Table 3-23 shows the SFF HDD options for the HS21 XM server.

Table 3-23 Small form factor hard disk drive options for HS21 XM

Part number	Feature code	Description
26K5776	5256	36.4 GB SAS HDD (2.5") SFF
26K5777	5298	73.4 GB SAS HDD (2.5") SFF
43X0845	None	73 GB SAS 15K SFF HDD
42D0421	5529	146 GB SAS 10K SFF HDD
43W7666	5596	300 GB 10K SATA 2.5" NHS HDD

The SAS hard disk is mounted onto the system board of the blade server and is not accessible externally during operation.

Solid state drives

The HS21 XM supports the attachment of a solid state drive (SSD). This device is installed in the existing SAS hard drive bay. If an SSD device is to be used, then a SAS disk drive cannot be installed as they both use the same drive bay.

These SSD devices are designed for operating system boot and low-bandwidth data transmission. SSD devices use less power than a standard SAS drive and therefore also generate less heat.

Table 3-17 on page 272 shows the SSD options for the HS21 XM.

Table 3-24 Solid State Drive options for the HS21 XM

Part number	FC	Description
43W7614	5557	IBM 15.8 GB 2.5” Solid State Drive (single drive with no RAID support)
43W7606	5558	IBM Dual 15.8 GB 2.5” Solid State Drive (two drives in one carrier)[a]
43W7618	5563	IBM 31.4 GB 2.5" Solid State Drive (single drive with no RAID support)

a. With this option installed, it is possible to mirror (RAID-1) these two devices using the onboard LSI Logic SAS controller in the HS21 XM.

Figure 3-39 shows the Dual 15.8 GB Solid State Drive.

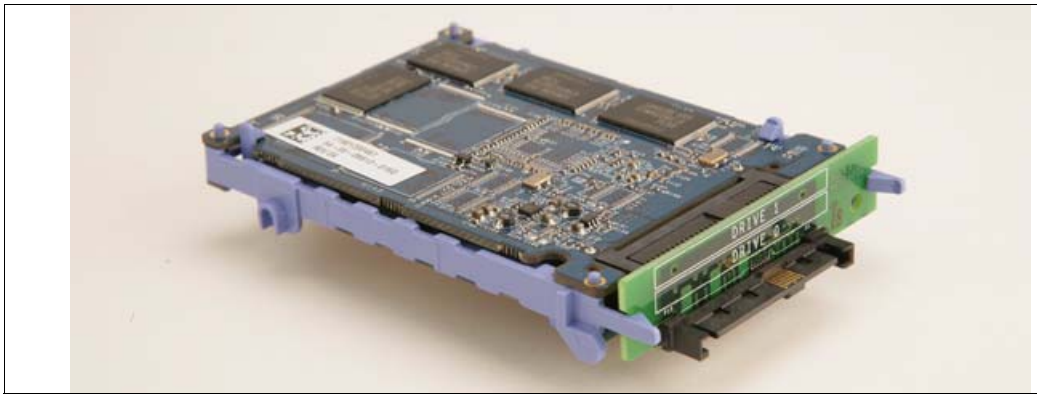

Figure 3-39 A Dual 15.8 GB Solid State Drive option

IBM Modular Flash Drives

The Modular Flash Drive is a USB 2.0 solid state storage device that provides an alternative to the spinning media of traditional hard disk drives for booting and low bandwidth/local storage. No moving parts means potentially less points of failure, lower power consumption, and less heat.

The two Modular Flash Drive options are shown in Table 3-25.

Table 3-25 Modular Flash Drive options for the HS21 XM

Part number	Feature code	Description
39R8697	2042	IBM 4 GB Modular Flash Drive
43W3934	5548	IBM 8 GB Modular Flash Drive

The modular flash drive can be used instead of, or in conjunction with, a traditional hard disk. Like a traditional disk drive, the solid state storage can be used as a boot device (Linux only) and can perform regular storage functions. It is not recommend this for high bandwidth applications such as file serving, Web serving, and mail.

Support as a boot drive is limited to Linux. Under the Windows operating systems, the IBM 4 GB and 8 GB Modular Flash Drive options are limited to usage as a local data drive.

The IBM Modular Flash Drive option is currently only supported on the HS21 XM blade server.

Figure 3-40 shows the IBM 4 GB Modular Flash Drive. Figure 3-27 on page 262 shows the IBM 4 Gb Modular Flash Drive option.

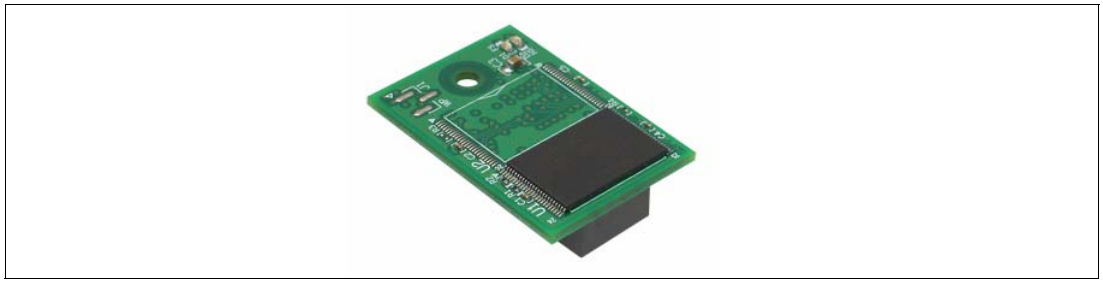

Figure 3-40 The IBM 4 GB Modular Flash Drive (part number 39R8697)

BladeCenter Storage and I/O Expansion Unit (SIO)

The HS21 XM supports the attachment of one BladeCenter Storage and I/O expansion unit.

The Storage and I/O (SIO) expansion unit provides the HS21 XM blade server with an additional three 2.5-inch SAS hot swap HDD bays and battery backed RAID support as well as two extra PCI-X I/O expansion sockets.

The HS21 XM does not support the attachment of any of the previous BSE-I or BSE-II expansion units.

With the SIO installed, the ServeRAID 8k-l (or ServeRAID 8k, if installed) cannot see the single SAS disk drive on the attached blade server. This means that you can only form RAID arrays using the three disk drives in the SIO with the ServeRAID 8k-l or ServeRAID 8k. The SAS disk drive on the attached blade server can still be installed but it will be controlled separately by the blade server SAS controller.

Note: This is different from what the HS21 type 8853 offers, as described in 3.5.5, “Local storage options” on page 293.

For more information about the SIO, see 3.18.2, “Storage and I/O Expansion Unit” on page 363.

3.4.6 I/O expansion options

The blade server HS21 XM has two onboard connectors for various I/O expansion card options, a PCI-X connector and a PCI Express connector.

The PCI-X connectors allows the connection of one of:

- A Small Form Factor expansion card
- A Standard Form Factor expansion
- A CFFv form factor expansion card

The installed expansion card is connected to module bays 3 and 4 at the rear of the BladeCenter chassis via the internal midplane. The interface of this I/O connector is a PCI-X 133 MHz bus.

The HS21 XM also has a PCI Express connector that allows for the attachment of one of the following:

- A High Speed Form Factor (HSFF) expansion card
- A CFFh expansion card
- The Storage and I/O (SIO) Expansion Unit

See 3.1, “Expansion cards” on page 234, for further information.

Tip: If an HSFF or a CFFh expansion card is installed in a blade server and you wish to install an SIO Expansion Unit, you need to first remove the expansion card, then attach the SIO Expansion Unit, then install the HSFF or CFFh expansion card in the SIO Expansion Unit instead.

Table 3-18 on page 273 summarizes the available I/O expansion card options for the HS21 XM.

Table 3-26 Blade server HS21 XM I/O expansion card options

Part number	FC	Description	Form factor	More information
26K4841	None	2 Gb Fibre Channel Expansion Card	SFF	Page 248
39R8624	1548	BladeCenter SFF Gigabit Ethernet Expansion Card	SFF	Page 243
73P6000	1546	Myrinet Cluster Expansion Card[a]	StFF	None
32R1896	1461	Cisco Systems InfiniBand 1X HCA Expansion Card	StFF	None
32R1923	1458	QLogic iSCSI Expansion Card	StFF	Page 252
26R0890	1577	QLogic 4 Gb SFF Fibre Channel Expansion Card	SFF	Page 249
39Y9186	2925	Emulex 4 Gb SFF Fibre Channel Expansion Card	SFF	Page 249
43W6859	2994	Emulex 4 Gb Fibre Channel Expansion Card (CFFv)	CFFv	Page 250

Part number	FC	Description	Form factor	More information
32R1760	1466	Cisco Systems InfiniBand 4X HCA Expansion Card	HSFF	Page 251
39Y9271	2967	NetXen 10 Gb Ethernet Expansion Card (CFFh)	CFFh	Page 243
44W4465	5479	Broadcom 10 Gb Ethernet CFFh Expansion Card	CFFh	Page 245
44W4466	5489	Broadcom 2-Port 10 Gb Ethernet CFFh Expansion Card	CFFh	Page 244
39Y9306	2968	QLogic Ethernet and 4 Gb Fibre Channel Expansion Card (CFFh)	CFFh	Page 246
44X1940	5485	QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)	CFFh	Page 247
39Y9310	2969	Ethernet Expansion Card (CFFv)	CFFv	Page 242
44W4479	5476	2/4 Port Ethernet Expansion Card (CFFh)	CFFh	Page 242
41Y8527	2970	QLogic 4 Gb Fibre Channel Expansion Card (CFFv)	CFFv	Page 248
39Y9190	2979	SAS Expansion Card (CFFv)	CFFv	Page 252
43W3974	1591	SAS Connectivity Card	CFFv	Page 254
43W4420	2993	Voltaire InfiniBand 4X DDR Expansion Card (CFFh)	CFFh	Page 251
43W4421	None	Cisco InfiniBand 4X DDR Expansion Card (CFFh)	CFFh	Page 251
43W4423	None	InfiniBand 4X DDR Expansion Card (CFFh)	CFFh	Page 251

a. The Myrinet card is a single-channel card and requires you to install an OPM in BladeCenter I/O module bay 4.

PCI I/O Expansion Unit 2 (PEU2)

The HS21 XM supports the attachment of one PCI I/O Expansion Unit 2. The PEU2 allows standard PCI-X type adapters to be attached to the HS21 XM type 7995.

For more information about the PEU2, see 3.18.4, “BladeCenter PCI I/O Expansion Unit 2” on page 367

PCI Express I/O Expansion Unit (PEU3e)

The HS21 XM supports the attachment of one PCI Express I/O Expansion Unit. The PEU3e allows standard PCI Express type adapters to be attached to the HS21 XM type 7995.

For more information about the PEU3e, see 3.18.5, “PCI Express I/O Expansion Unit (PEU3e)” on page 368.

3.5 HS21 type 8853

The HS21 type 8853 is a single-slot blade server that can have up to two dual-core or quad-core Intel Xeon processors, up to 32 GB of PC2-5300 ECC DDR2 FBDIMM memory in 4 DIMM slots, and up to two fixed SAS HDD for high performance and reliability. The HS21 is designed to help provide customers with the computing power they require to match their business needs and growth.

The HS21 type 8853 has two Gigabit Ethernet controllers as standard. They provide high speed data connection and offer TOE (TCP Offload Engine) support, load-balancing and failover capabilities.

It is possible to attach optional I/O expansion cards to the HS21. Each blade can connect to additional Ethernet, Fibre Channel, iSCSI, InfiniBand and other high-speed communication switches housed in the chassis. Optional I/O expansion cards add additional fabrics to the HS21 type 8853 as needed.

The HS21 type 8853 blade server has a unique option called the Memory and I/O Expansion Unit (MIO). The MIO attaches to the HS21 blade server making one double width blade server with 8 DIMMs slots and four Ethernet ports in total.

There are various HS21 type 8853 models available. The different models have different Intel Xeon processor types and speeds.

All the models have the following features as standard:

- Standard one processor, upgradable to two
- Models with dual-core or quad-core processors
- Two Gigabit Ethernet (Broadcom 5708S) controllers
- Integrated Baseboard management controller (BMC)
- A LSI Logic 53C1064E SAS controller for onboard SAS hard disk drives
- Onboard RAID-0 or RAID-1 support
- Up to two internal 2.5-inch Small Form Factor (SFF) SAS hard disk drives
- One PCI-X type connector for a StFF, SFF, or CFFv I/O expansion adapter
- One PCI Express connector for the attachment of a CFFh or HSFF expansion card, or an Expansion unit
- Support for the Storage and I/O expansion (SIO) unit, Memory and I/O (MIO) expansion unit and the PCI I/O Expansion Unit 2 (PEU2)
- Light Path Diagnostics on system board
- Predictive failure analysis (PFA)

See Table 3-27 for the available models of the blade server HS21 type 8853.

Table 3-27 Models of the HS21 type 8853 at a glance

Model	Intel CPU	Cores per socket	CPU/FSB speed (GHz/MHz)	L2 cache[a]	CPUs Std/Max	Memory Std/Max	CPU power
Models with dual-core processors							
8853-L4x	Xeon 5140	Dual core	2.33/1333	4 MB	1/2	1 GB/32 GB	65 W
8853-NTx	Xeon 5138 LV	Dual core	2.13/1333	4 MB	1/2	1 GB/32 GB	35 W
8853-H1x	Xeon 5148 LV	Dual core	2.33/1333	4 MB	1/2	1 GB/32 GB	40 W
8853-R2x	Xeon X5260	Dual core	3.33/1333	6 MB	1/2	2 GB/32 GB	80 W
8853-RLx	Xeon L5240	Dual core	3.0/1333	6 MB	1/2	2 GB/32 GB	40 W
Models with quad-core processors							
8853-C2x	Xeon E5345	Quad core	2.33/1333	8 MB	1/2	2 GB/32 GB	80 W
8853-G1x	Xeon E5405	Quad core	2.0/1333	12 MB	1/2	2 GB/32 GB	80 W
8853-G3x	Xeon E5420	Quad core	2.5/1333	12 MB	1/2	2 GB/32 GB	80 W

Model	Intel CPU	Cores per socket	CPU/FSB speed (GHz/MHz)	L2 cache[a]	CPUs Std/Max	Memory Std/Max	CPU power
8853-G4x	Xeon E5430	Quad core	2.66/1333	12 MB	1/2	2 GB/32 GB	80 W
8853-G5x	Xeon E5440	Quad core	2.83/1333	12 MB	1/2	2 GB/32 GB	80 W
8853-G6x	Xeon E5450	Quad core	3.0/1333	12 MB	1/2	2 GB/32 GB	80 W
8853-G7x	Xeon X5460	Quad core	3.16/1333	12 MB	1/2	2 GB/32 GB	120 W
8853-GLx	Xeon L5420	Quad core	2.5/1333	12 MB	1/2	2 GB/32 GB	50 W
8853-GNx	Xeon L5408	Quad core	2.13/1066	12 MB	1/2	4 GB/32 GB	40 W

a. The L2 cache is shared between the cores. Each core has 2 MB, 4 MB, or 6 MB of L2 cache, depending on the model.

Figure 3-41 shows a blade server HS21 type 8853 without its top cover.

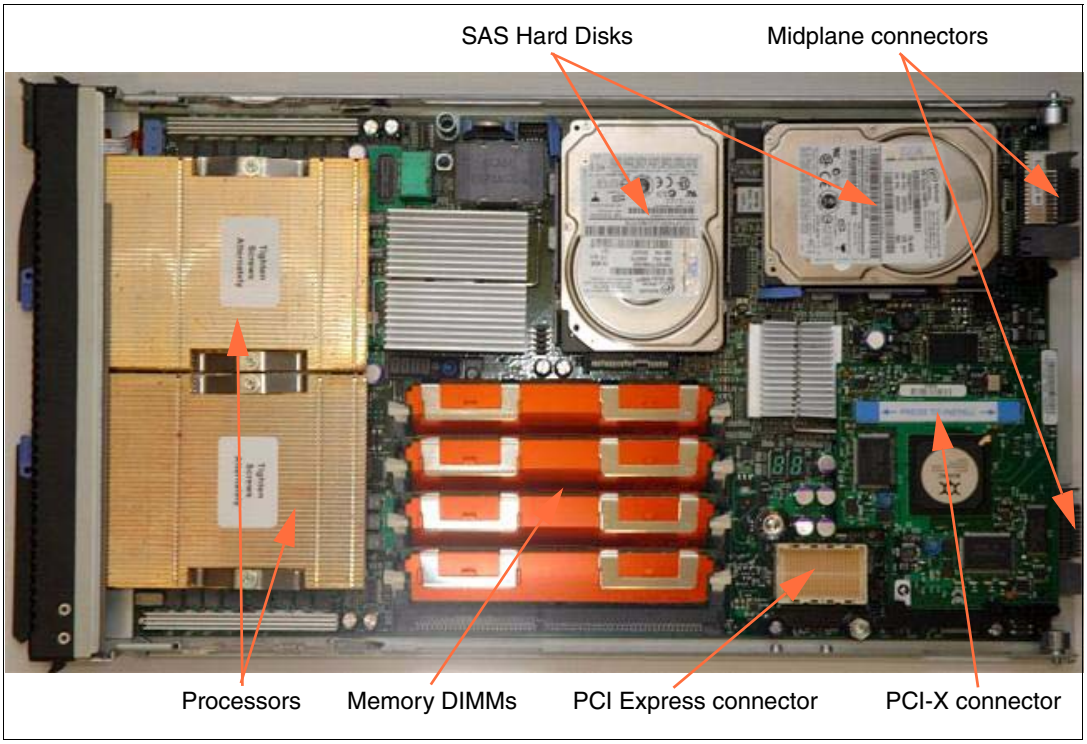

Figure 3-41 Blade server HS21 type 8853 with its top cover removed

3.5.1 Processor

The blade server HS21 supports either:

- Up to two Intel Xeon 5400 series (Harpertown) quad-core processors
- Up to two Intel Xeon 5300 series (Clovertown) quad-core processors
- Up to two Intel Xeon 5200 series (Wolfdale) dual-core processors
- Up to two Intel Xeon 5100 series (Woodcrest) dual-core processors

These processors support Intel 64 Technology (EM64T). In a two socket configuration, both processors must be identical.

Table 3-28 shows the blade server HS21 type 8853 processor options.

Table 3-28 Blade server HS21 type 8853 processor options

Server model	CPU option part number	CPU model	CPU specifications (cores/speed/FSB speed/L2 cache)
Intel Xeon Dual Core processors			
8853-L4x	40K1228	Xeon 5140	2.33 GHz/1333 MHz/4 MB L2
8853-NTx	40K1218	Xeon 5138 LV	2.13 GHz/1333 MHz/4 MB L2
8853-H1x	40K1217	Xeon 5148 LV	2.33 GHz/1333 MHz/4 MB L2
8853-R2x	44T1792	Xeon X5260	3.33 GHz/1333 MHz/6 MB L2/ 80w
8853-RLx	44T1812	Xeon L5240	3.0 GHz/1333 MHz/6 MB L2/ 40w
Intel Xeon Quad Core processors			
8853-C2x	43W1144	Xeon E5345	2.33 GHz/1333 MHz/8 MB L2 80w
8853-G1x	44T1716	Xeon E5405	2.0 GHz/1333 MHz/12 MB L2 /80w
8853-G3x	44T1742	Xeon E5420	2.5 GHz/1333 MHz/12 MB L2 /80w
8853-G4x	44T1724	Xeon E5430	2.66 GHz/1333 MHz/12 MB L2/ 80w
8853-G5x	44T1740	Xeon E5440	2.83 GHz/1333 MHz/12 MB L2/ 80w
8853-G6x	44T1728	Xeon E5450	3.0 GHz/1333 MHz/12 MB L2/ 80w
8853-G7x	44T1738	Xeon X5460	3.16 GHz/1333 MHz/12 MB L2 /120w
8853-GLx	44T1816	Xeon L5420	2.5 GHz/1333 MHz/12 MB L2/50w
8853-GNx	44T1899	Xeon L5408	2.13 GHz/1066 MHz/12 MB L2/40w

3.5.2 Memory

The memory used in the HS21 has the following features:

- PC2-5300 (667 MHz) Fully Buffered DDR-2 DIMM
- ECC
- Chipkill

Memory configuration has to follow these guidelines:

- There are a total of four DIMM slots. Two of these slots (1 and 3) are preconfigured with a pair of DIMMs.
- Because memory is 2-way interleaved, the memory modules must be installed in matched pairs.

See Table 3-29 for HS21 memory options.

Table 3-29 Blade server HS21 type 8853 memory options

Part number	Memory description
39M5782	1 GB (2x 512 MB) PC2-5300 667 MHz ECC Chipkill DDR2 FB-DIMM
39M5785	2 GB (2x 1 GB) PC2-5300 667 MHz ECC Chipkill DDR2 FB-DIMM
39M5791	4 GB (2x 2 GB) PC2-5300 667 MHz ECC Chipkill DDR2 FB-DIMM
39M5797	8 GB (2x 4 GB) PC2-5300 667 MHz ECC Chipkill DDR2 FB-DIMM
46C7418	2 GB (2x 1 GB) Single Rank PC2-5300 CL5 ECC FBD 667MHz Low Power Memory
46C7419	4 GB (2x 2 GB) Dual Rank PC2-5300 CL5 ECC FBD 667MHz Low Power Memory
46C7420	8 GB (2x 4 GB) Quad Rank PC2-5300 CL5 ECC FBD 667MHz Low Power Memory
46C7577	16 GB (2x 8 GB) PC2-5300 CL5 ECC FBDIMM 667Mhz AMB+

Table 3-30 shows the standard DIMM configuration by model.

Table 3-30 HS21 type 8853 standard DIMM configuration

Model	Standard DIMM configuration
L4x, NTx, H1x	2x 512 MB FB-DIMMs
C2x, R2x, RLx, G1x, G3x, G4x, G5x, G6x, G7x, GLx	2x 1 GB FB-DIMMs
GNx	2x 2 GB FB-DIMMs

BladeCenter Memory and I/O (MIO) expansion unit

The HS21 type 8853 also supports the attachment of one BladeCenter Memory and I/O (MIO) expansion unit. The MIO attaches to the HS21 blade server, thus making one double-width blade server with 8 DIMMs slots and four Ethernet ports in total. The MIO is only supported with HS21 type 8853.

For more information about the MIO, see 3.18.1, “BladeCenter Memory and I/O Expansion Unit” on page 361.

3.5.3 Onboard network controllers

The HS21 type 8853 blade server includes two integrated Broadcom BCM5708S Gigabit Ethernet controllers. These two Ethernet interfaces are routed through to the redundant midplane providing a foundation for fault-tolerant network setups.

Note: When installed in a BladeCenter E, H, T, and HT chassis, Ethernet Port 1 on the blade server is routed to the switch in module bay 1 and port 2 is routed to the switch in module bay 2.

When installed in BladeCenter S, both Ethernet ports on the blade server are routed to the switch in module bay 1.

The following features are supported by the Broadcom (BCM5708S) onboard controllers:

- TOE (TCP Offload Engine)
- iSCSI initiator

- PXE 2.0 boot support
- Wake on LAN
- Load-balancing
- VLAN tagging support
- Jumbo frame support (9 KB)

The Broadcom Advanced Server Program (BASP) is management software provided by Broadcom for managing your network cards. This includes a command line interface for scripting purposes.

3.5.4 Integrated systems management processor

The blade server HS21 type 8853 has an integrated baseboard management controller (BMC) that provides:

- IPMI compliance (Intelligent Platform Management Interface)
- Serial over LAN (SOL)
- Wake on LAN (WOL)
- Power management functions
- KVM and media tray ownership
- Light path diagnostics support
- Automatic BIOS recovery (ABR)
- Automatic server restart (ASR)
- Predictive failure analysis (PFA) on system components (processor, memory, and hard disk drives)
- Inventory
- Error logging
- Environmental monitoring (voltages and temperature on the system board of the blade server)

Note: You cannot access this onboard BMC externally through a dedicated port or daisy-chain it to other servers or blades. The BMC is forwarding information to the Management Module within the BladeCenter chassis. The management module is a single point of control for hardware management of all server and switch components contained within the chassis.

3.5.5 Local storage options

The HS21 type 8853 blade server features an onboard LSI Logic 53C1064E SAS controller with two SAS connectors for up to two optional 2.5-inch Small Form Factor (SFF) SAS or SATA hard disk drives. The LSI Logic 53C1064E SAS controller provides RAID 0 or RAID 1 capability.

Table 3-31 shows the SFF HDD options for HS21 type 8853.

Table 3-31 Small form factor hard disk drive options for HS21 type 8853

Part number	Description
26K5776	36.4 GB SAS HDD (2.5") SFF

Part number	Description
26K5777	73.4 GB SAS HDD (2.5") SFF
43X0845	73 GB SAS 15K SFF HDD
42D0421	146 GB SAS 10K SFF HDD
43W7666	300 GB 10K SATA 2.5" NHS HDD

These two hard disks are mounted onto the system board of the blade server and are not accessible externally during operation.

Solid state drives

The HS21 type 8853 supports the attachment of up to two solid state drives (SSD). The SSD device is installed in the disk drive bay. If an SSD device is to be used, then a SAS disk drive cannot be installed as they both use the same drive bay. When two of these SSD devices are installed the onboard LSI Logic 53C1064E SAS controller can create a RAID 0 or RAID 1 array.

These SSD devices are designed for operating system boot and low-bandwidth data transmission. SSD devices use less power than a standard SAS drive and therefore also generate less heat.

Table 3-32 shows the SSD options for the HS21.

Table 3-32 SSD options for the HS21

Part number	Description
43W7614	IBM 15.8 GB 2.5” Solid State Drive
43W7618	IBM 31.4 GB 2.5" Solid State Drive

Note: Unlike the HS21 XM, the HS21 does not support the 15.8 GB dual drive option, part number 43W7606.

BladeCenter Storage and I/O Expansion Unit (SIO)

The HS21 type 8853 supports the attachment of one BladeCenter Storage and I/O expansion unit (SIO). The Storage and I/O (SIO) expansion unit provides the HS21 XM blade server with an additional three 2.5-inch SAS hot swap HDD bays and battery-backed RAID support, as well as two extra PCI-X I/O expansion sockets.

The HS21 does not support the attachment of any of the previous BSE-I or BSE-II expansion units.

With the SIO installed, the ServeRAID 8k-l (or ServeRAID 8k, if installed) can use combinations of both the SAS drives in the SIO as well as in the attached blade server to form arrays. This means that you can form one RAID array of five drives, two in the blade server, and three in the SIO.

Note: This is different to the HS21 XM as described in 3.3.5, “Local storage options” on page 268.

For more information about the SIO, see 3.18.2, “Storage and I/O Expansion Unit” on page 363.

3.5.6 I/O expansion options

The blade server HS21 type 8853 has two onboard connectors for connecting various network and storage I/O expansion cards and expansion units, a PCI-X connector and a PCI Express connector. See 3.1, “Expansion cards” on page 234, for a description of these connectors and the cards they support.

Table 3-33 summarizes the available I/O options.

Table 3-33 Blade server HS21 type 8853 I/O expansion card options

Part number	Description	Form factor	Page
26K4841	2 Gb Fibre Channel Expansion Card	SFF	248
39R8624	BladeCenter SFF Gigabit Ethernet Expansion Card	SFF	243
73P6000	Myrinet Cluster Expansion Card[a]	StFF	None
32R1896	Cisco Systems InfiniBand 1X HCA Expansion Card	StFF	None
32R1923	QLogic iSCSI Expansion Card	StFF	252
26R0890	QLogic 4 Gb SFF Fibre Channel Expansion Card	SFF	249
39Y9186	Emulex 4 Gb SFF Fibre Channel Expansion Card	SFF	249
43W6859	Emulex 4 Gb Fibre Channel Expansion Card (CFFv)	CFFv	250
32R1760	Cisco Systems InfiniBand 4X HCA Expansion Card	HSFF	251
39Y9271	NetXen 10 Gb Ethernet Expansion Card (CFFh)	CFFh	243
44W4465	Broadcom 4-port 10 Gb Ethernet CFFh Expansion Card	CFFh	245
44W4466	Broadcom 2-Port 10 Gb Ethernet CFFh Expansion Card	CFFh	244
39Y9306	QLogic Ethernet and 4 Gb Fibre Channel Expansion Card (CFFh)	CFFh	246
44X1940	QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)	CFFh	247
39Y9310	Ethernet Expansion Card (CFFv)	CFFv	242
44W4479	2/4 Port Ethernet Expansion Card (CFFh)	CFFh	242
41Y8527	QLogic 4 Gb Fibre Channel Expansion Card (CFFv)	CFFv	248
39Y9190	SAS Expansion Card (CFFv)	CFFv	252
43W3974	SAS Connectivity Card	CFFv	254
43W4420	Voltaire InfiniBand 4X DDR Expansion Card (CFFh)	CFFh	251
43W4421	Cisco InfiniBand 4X DDR Expansion Card (CFFh)	CFFh	251
43W4423	InfiniBand 4X DDR Expansion Card (CFFh)	CFFh	251

a. The Myrinet card is a single-channel card and requires you to install an Optical Pass-thru Module (OPM) in BladeCenter I/O module bay 4.

PCI I/O Expansion Unit 2 (PEU2)

The HS21 supports the attachment of one PCI I/O Expansion Unit 2. The PEU2 allows standard PCI-X type adapters to be attached to the HS21 type 8853.

For more information about the PEU2, see 3.18.4, “BladeCenter PCI I/O Expansion Unit 2” on page 367.

PCI Express I/O Expansion Unit (PEU3e)

The HS21 supports the attachment of one PCI Express I/O Expansion Unit. The PEU3e allows standard PCI-Express type adapters to be attached to the HS21 type 8853.

For more information about the PEU3e, see 3.18.5, “PCI Express I/O Expansion Unit (PEU3e)” on page 368.

3.6 HS20 type 8843

There BladeCenter HS20 type 8843 is a general-purpose Intel processor-based blade server suitable for a wide variety of tasks.

Note: The remaining two models of the HS20 type 8843 are available primarily for customers that have standardized on particular models and wish for them to remain available for sale for an extended period.

All these models have these features as standard:

- One standard processor (single core) upgradeable to two processors
- Dual Gigabit Ethernet controller (Broadcom 5704S)
- Baseboard management controller (BMC) service processor
- A single-channel SCSI Ultra320 LSI Logic 53C1020 controller for onboard hard disk drives
- Up to two internal 2.5-inch Ultra320 SCSI disk drives
- One PCI-X slot for an I/O expansion adapter (SFF, StFF or CFFv form factors)
- One expansion connector for the attachment of a Blade Storage Expansion Unit II or a PCI Expansion Unit II
- Light Path Diagnostics on system board
- Predictive failure analysis (PFA)

Table 3-34 lists the available models of blade server HS20 type 8843.

These models have a low profile handle (LPH) for use in a NEBS-3/ETSI environment with the BladeCenter T chassis. These models are denoted by the letter T in the model number, as in 8843-xTx.

Table 3-34 Blade servers HS20 type 8843 at a glance

Model	CPU/FSB Speed (GHz/MHz)	Cache L2	Processors (Standard/Max)	Memory Standard/Maximum
8843-JTY[a]	2.8/800	1 MB L2	1/2	1 GB/16 GB
8843-MTY[a]	3.0/800	2 MB L2	1/2	1 GB/16 GB

a. Models JTY and MTY are telco models with low profile handles.

Figure 3-42 on page 297 shows a blade server 8843 without its top cover.

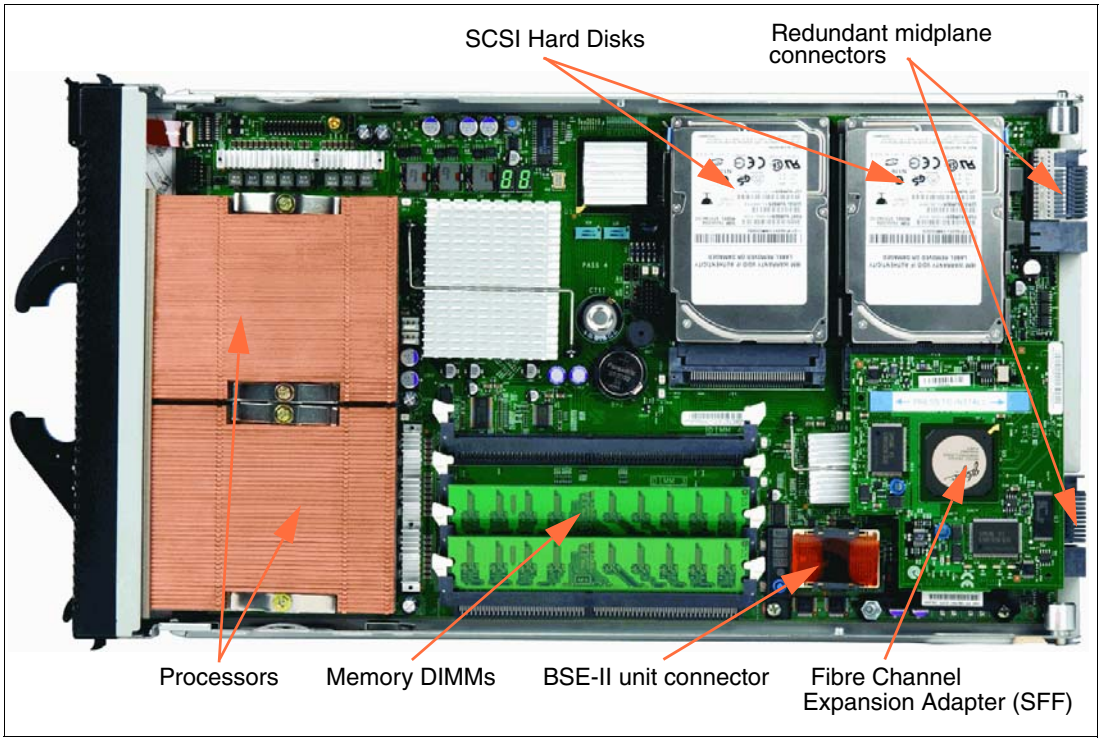

Figure 3-42 Blade server HS20 type 8843 with its top cover removed

3.6.1 Processor

The blade server type HS20 type 8843 supports up to two Intel Xeon DP single-core microprocessors. All these processors support Intel 64 Technology (EM64T). Both processors must have the same cache size and type, the same clock speed, and identical internal and external clock frequencies including system bus speed.

Table 3-35 shows blade server type 8843 processor options.

Table 3-35 Blade server type HS20 type 8843 processor options

Part number	For model	Processor description
13N0699	8843-JTY	Intel Xeon 2.8 GHz/800 MHz, Single Core, 1 MB L2 cache, Low Voltage
25R8877	8843-MTY	Intel Xeon 3.0 GHz/800 MHz, Single Core, 2 MB L2 cache, Low Voltage

3.6.2 Memory

The memory used in the blade server HS20 type 8843 has the following features:

- PC2-3200 400 MHz CL2.5 ECC DDR2-SDRAM
- ECC
- Chipkill

All supported DIMMs listed in Table 3-36 on page 298 support Chipkill.

Memory configuration has to follow these guidelines:

- There are a total of four DIMM slots. Two of these slots (1 and 2) are preconfigured with a pair of DIMMs.

- Because memory is 2-way interleaved, the memory modules must be installed in matched pairs. However, one DIMM pair is not required to match the other in capacity.

See Table 3-36 for blade server 8843 memory options.

Table 3-36 Blade server HS20 type 8843 memory options

Part number	Memory description
39M5809	2 GB (2x 1 GB DIMMs) PC2-3200 CL3 ECC DDR2 SDRAM RDIMM
39M5812	4 GB (2x 2 GB DIMMs) PC2-3200 CL3 ECC DDR2 SDRAM RDIMM
30R5145	8 GB (2x 4 GB DIMMs) PC2-3200 CL3 ECC DDR2 SDRAM RDIMM

Table 3-37 shows the standard DIMM configuration by model.

Table 3-37 Standard DIMM configuration

Model	Standard DIMM configuration	Chipkill
JTY, MTY	2 x 512 MB PC2-3200 ECC DDR2 RDIMM	Yes

3.6.3 Onboard network controller

The HS20 type 8843 blade server uses a dual channel Gigabit Ethernet Broadcom BCM5704S. These two Ethernet interfaces are routed through to the redundant midplane providing a foundation for fault-tolerant network setups.

The following features are supported by the Broadcom onboard adapter:

- PXE boot support
- Wake on LAN
- Load-balancing
- Fault-tolerance
- VLAN support

The Broadcom Advanced Server Program (BASP) is management software provided by Broadcom for managing your network cards. This includes a command line interface for scripting purposes.

3.6.4 Integrated systems management processor

The blade server HS20 type 8843 has an integrated baseboard management controller (BMC) that provides:

- IPMI compliance (Intelligent Platform Management Interface)
- Serial over LAN (SOL)
- Wake on LAN (WOL)
- Power management functions
- KVM and media tray ownership
- Light path diagnostics support
- Automatic BIOS recovery (ABR)
- Automatic server restart (ASR)

- Predictive failure analysis (PFA) on system components (processor, memory, and hard disk drives)
- Inventory
- Error logging
- Environmental monitoring (voltages and temperature on the system board of the blade server)

Note: You cannot access this onboard BMC externally through a dedicated port or daisy-chain it to other servers or blades. The BMC is forwarding information to the Management Module within the BladeCenter chassis. The Management Module is a single point of control for hardware management of all server and switch components contained within the chassis.

3.6.5 Local storage options

The HS20 type 8843 blade server features an onboard LSI Logic 53C1020 SCSI controller with two SCSI connectors for two optional 2.5-inch Ultra320 hard disk drives.

Table 3-38 shows the SFF HDD options for HS20 type 8843.

Table 3-38 Small form factor (SFF) hard disk drive options for HS20 type 8843

Part number	Description
40K1037	36.4 GB Non Hot-Swap 2.5" 10K RPM Ultra320 SCSI HDD
40K1038	73.4 GB Non Hot-Swap 2.5" 10K RPM Ultra320 SCSI HDD

These two hard disks are mounted onto the system board of the blade server and are inaccessible externally during operation.

If you need hot swappable hard disk drives, the BladeCenter SCSI Storage Expansion Unit II provides this capability. See 3.18.3, “BladeCenter SCSI Storage Expansion Unit II” on page 366.

3.6.6 I/O expansion options

The blade server HS20 type 8843 has an onboard connector for various I/O expansion options including network and storage adapters. These adapters are connected to module bays 3 and 4 at the rear of the BladeCenter chassis. The interface of the I/O expansion slot is a PCI-X 133 MHz bus.

Table 3-39 summarizes the available options.

Table 3-39 HS20 type 8843 expansion card options

Part number	FC	Description	Form Factor
13N2203	None	BladeCenter Fibre Channel Expansion Card	StFF
26K4841	None	Fibre Channel Expansion Card	SFF
73P9030	None	BladeCenter Gigabit Ethernet Expansion Card	StFF
39R8624	1548	BladeCenter SFF Gigabit Ethernet Expansion Card	SFF
73P6000	1546	Myrinet Cluster Expansion Card[a]	StFF

Part number	FC	Description	Form Factor
32R1896	1461	Cisco Systems InfiniBand HCA Expansion Card	StFF
32R1923	1458	QLogic iSCSI Expansion Card	StFF
26R0884	None	QLogic 4 Gb Standard Fibre Channel Expansion Card	StFF
26R0890	1577	QLogic 4 Gb SFF Fibre Channel Expansion Card	SFF
39Y9186	2925	Emulex 4 Gb SFF Fibre Channel Expansion Card	SFF
39Y9190	2979	SAS Expansion Card (CFFv)	CFFv

a. The Myrinet card is a single-channel card and requires you to install an Optical Pass-thru Module (OPM) in BladeCenter I/O module bay 4.

The HS20 8843 also supports the PCI Expansion Unit II, part number 25K8373, which provides two full length PCI-X adapters. See 3.18.4, “BladeCenter PCI I/O Expansion Unit 2” on page 367 for details of this expansion unit.

Note: The HS20 type 8843 does not support CFFh or HSFF expansion cards.

3.7 LS21 type 7971

The LS21 type 7971 is a 30 mm blade, capable of supporting up to two dual-core AMD Opteron processors, up to 32 GB of ECC memory and one fixed SAS HDD.

The LS21 blade server supports up to two of the dual-core 2200 series AMD Opteron processors.

Dual Gigabit Ethernet controllers are standard, providing high-speed data transfers and offering TOE (TCP Offload Engine) support, load-balancing and failover capabilities.

With optional expansion cards, the LS21 can connect to additional Ethernet, Myrinet, Fibre Channel, iSCSI, InfiniBand and other high-speed communication switches housed in the chassis. Optional expansion cards add additional fabrics to the LS21 as needed.

This blade is designed with power management capability to provide customers with the maximum uptime possible for their systems. In extended thermal conditions, rather than shut down completely, or fail, the LS21 automatically reduces the processor frequency to maintain acceptable thermal levels.

Figure 3-43 shows the components inside the LS21 type 7971 blade server.

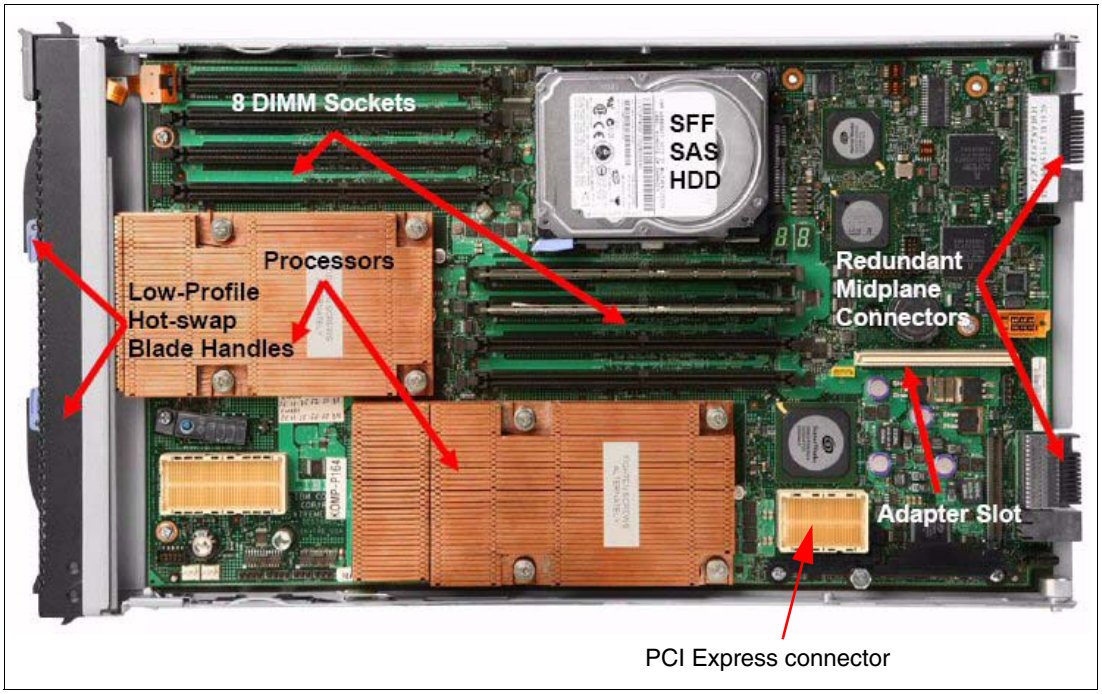

Figure 3-43 Blade server LS21 type 7971 with its top cover removed

All the models of the LS21 type 7971 have the following features as standard:

- One standard processor upgradeable to two
- Two Gigabit Ethernet (Broadcom 5706S) controllers
- Integrated BMC
- A LSI Logic 53C1064 SAS controller for onboard storage
- Support for up to one 2.5-inch internal fixed SAS disk drive
- One PCI-X connector slot for a StFF, SFF or CFFv I/O expansion card
- One PCI Express connector for the attachment of a HSFF expansion card, CFFh expansion card or Expansion unit
- Support for the Storage and I/O expansion (SIO) unit
- Light Path Diagnostics on system board

Table 3-40 lists the available models of the LS21 type 7971 blade server.

Table 3-40 Blade servers LS21 type 7971 at a glance

Type-Model	Processor	CPU speed	L2 cache[a]	Processor s Std/Max	Memory Std/Max	Processor power draw[b]
7971-51x	AMD Opteron 2216 HE	2.4 GHz	2 MB	1/2	2 GB/32 GB	68 W
7971-61x	AMD Opteron 2218 HE	2.6 GHz	2 MB	1/2	4 GB/32 GB	68 W
7971-7Ax	AMD Opteron 2220	2.8 GHz	2 MB	1/2	2 GB/32 GB	95 W
7971-8Ax	AMD Opteron 2222	3.0 GHz	2 MB	1/2	2 GB/32 GB	95 W

a. 1 MB per core for 2200 series.

b. Versions of the LS21 that use the 95 W processors are designed for use in the BladeCenter H chassis. Due to increased power and thermal needs, these 95 W versions will run at 'n-2' processor speed in the BladeCenter and BladeCenter T chassis. In most cases this represents a speed reduction of 400 MHz. The versions of the LS21 that use the High Efficiency (HE) 68 W processors will run at full speed in all BladeCenter chassis.

3.7.1 Processor

The LS21 7971 supports up to two AMD Opteron 2200 series dual-core processors.

Table 3-41 shows blade server LS21 type 7971 processor options.

Table 3-41 Blade server LS21 type 7971 processor options

LS21 model	Part number	Processor model
7971-51x	25R8899	AMD Opteron 2216 HE (2.4 GHz/68w)
7971-61x	44E5246	AMD Opteron 2218 HE (2.6 GHz/68w)
7971-7Ax	40K1264	AMD Opteron 2220 (2.8 GHz/95w)
7971-8Ax	44R4946	AMD Opteron 2222 (3.0 GHz/95w)

When installing two processors, they must be identical (same clock rate, cache size and the same number of cores).

3.7.2 Memory

The memory used in the LS21 has the following features:

- PC2-5300 (667 MHz) DDR2 RDIMM
- PC2-4200 (533 MHz) DDR2 memory (39M5870 only)
- ECC
- Very Low Profile (VLP)
- Chipkill

Memory configuration has to follow these guidelines:

- There are a total of eight DIMM slots (four per processor socket). Two of these slots (1 and 2) are preconfigured with a pair of DIMMs.
- Because memory is 2-way interleaved, the memory modules must be installed in matched pairs. However, one DIMM pair is not required to match the other in capacity.
- The maximum of 32 GB of installed memory is achieved when all DIMM sockets are populated with the 4 GB DIMMs.

For each installed microprocessor, a set of 4 DIMM sockets are enabled. See Table 3-42 for more information.

Table 3-42 Enabled DIMM sockets

Installed Microprocessors	Enabled DIMM sockets
1	1, 2, 3, 4
1 and 2	1, 2, 3, 4, 5, 6, 7, 8

See Table 3-43 on page 303 for blade server LS21 type 7971 memory options.

Table 3-43 Blade server LS21 7971 memory options

Part number	Memory description
39M5861	1 GB (2x512 MB) PC2-5300 DDR2 SDRAM VLP RDIMM
39M5864	2 GB (2x1 GB) PC2-5300 DDR2 SDRAM VLP RDIMM
39M5867	4 GB (2x 2 GB) PC2-5300 DDR2 SDRAM VLP RDIMM
46C0512	4 GB (2 x 2 GB) PC2-5300 DDR2 667 MHz Single Rank VLP RDIMM
39M5870	8 GB (2x 4 GB) PC2-4200 CL4 DDR2 SDRAM VLP RDIMM
46C0513	8 GB (2x 4 GB) PC2-5300 CL5 DDR2 SDRAM VLP RDIMM

Table 3-44 shows the standard LS21 type 7971 DIMM configuration by model.

Table 3-44 Standard LS21 7971 DIMM configuration by model

Model	Standard DIMM configuration	Chipkill
51Y, 7AY, 61Y, 8AY	2 GB (2x 1 GB) PC2-5300 DDR2 SDRAM VLP RDIMM	Yes

3.7.3 Onboard network controllers

The LS21 type 7971 blade server uses two Gigabit Ethernet Broadcom BCM5706S controllers. These two Ethernet interfaces are routed through to the redundant midplane, providing a foundation for fault-tolerant network setups.

Note: When installed in a BladeCenter E, H, T and HT chassis, Ethernet Port 1 on the blade server is routed to the switch module in Bay 1 and Ethernet port 2 is routed to the switch module in Bay 2.

When installed in BladeCenter S, both Ethernet ports on the blade server are routed to the switch module in Bay 1.

The following features are supported by the Broadcom (BCM5706S) onboard adapters:

- TOE (TCP Offload Engine)
- iSCSI initiator
- PXE 2.0 boot support
- Wake on LAN
- Load-balancing
- VLAN tagging support
- Jumbo frame support (9 KB)

The Broadcom Advanced Server Program (BASP) is management software provided by Broadcom for managing your network cards. This includes a command line interface for scripting purposes.

3.7.4 Integrated systems management processor

The LS21 type 7971 has an integrated baseboard management controller (BMC) that provides:

- IPMI compliance (Intelligent Platform Management Interface)
- Serial over LAN (SOL)
- Wake on LAN (WOL)

- Power management functions
- KVM and media tray ownership
- Light path diagnostics support
- Automatic BIOS recovery (ABR)
- Automatic server restart (ASR)
- Predictive failure analysis (PFA) on system components (CPU, memory, and disk drives)
- Inventory
- Error logging
- Environmental monitoring (voltages and temperature on the system board)

Note that you cannot access this onboard BMC externally through a dedicated port or daisy-chain it to other servers or blades. The BMC is forwarding information to the Management Module (MM) within the BladeCenter chassis. The MM is a single point of control for hardware management of all server and switch components contained within the chassis.

3.7.5 Local storage options

The LS21 type 7971 blade server features an onboard LSI Logic 53C1064 SAS controller with a SAS connector for one optional 2.5-inch SAS hard disk drive or one solid state drive (SSD)

Table 3-45 shows the SFF HDD options for LS21 type 7971.

Table 3-45 Hard disk drive options for LS21 type 7971

Part number	Description
26K5777	73 GB SAS HDD (2.5") SFF
42D0421	146 GB SAS 10K SFF
43X0845	73 GB SAS 15K SFF

The hard disk drive is mounted onto the system board of the blade server and are inaccessible from the outside during operation.

Solid state drives

The LS21 supports the attachment of a solid state drive (SSD). This device is installed in the existing SAS hard drive bay. If an SSD device is to be used, then a SAS disk drive cannot be installed as they both use the same drive bay. A single SSD cannot offer RAID support.

These SSD devices are designed for operating system boot and low-bandwidth data transmission. SSD devices use less power than a standard SAS drive and, therefore, also generate less heat.

Table 3-46 shows the SSD options for the LS21.

Table 3-46 Solid State Drive options for the LS21

Part number	Description
43W7614	IBM 15.8 GB 2.5” Solid State Drive
43W7618	IBM 31.4 GB 2.5" Solid State Drive

Note: Unlike the HS21 XM, the LS21 does not support the 15.8 GB dual drive option, part number 43W7606.

BladeCenter Storage and I/O Expansion Unit (SIO)

The LS21 type 7971 supports the attachment of one BladeCenter Storage and I/O expansion unit. For more information about the SIO, see 3.18.2, “Storage and I/O Expansion Unit” on page 363.

When the SIO is used with the LS21, the ServeRAID 8k-l in the SIO (or ServeRAID 8k, if installed) can form arrays using combinations of SAS drives in both the SIO and in the LS21. This means that you can form one RAID array of four drives — three in the SIO and one in the blade server.

Note: The LS21 type 7971 does not support the attachment of any of the previous BSE-I or BSE-II expansion units.

3.7.6 I/O expansion options

The LS21 type 7971 has an onboard PCI-X connector for various I/O expansion card options including network and storage adapters. These adapters are connected via the midplane to module bays 3 and 4 at the rear of the BladeCenter chassis. The interface of the I/O expansion slot is a PCI-X 133 MHz bus.

The LS21 type 7971 also has a PCI Express connector that allows for the attachment of either a HSFF expansion card, a CFFh expansion card or for the attachment of the Storage and I/O (SIO) Expansion Unit.

See 3.1, “Expansion cards” on page 234 for further information about form factors and expansion card options. See 3.18.2, “Storage and I/O Expansion Unit” on page 363 for information about the SIO.

Table 3-47 summarizes the available I/O options.

Table 3-47 Blade server LS21 type 7971 expansion options

Part number	Description	Form Factor
39R8624	IBM BladeCenter Gigabit Ethernet Expansion Card	SFF
26K4841	BladeCenter Fibre Channel Expansion Card	SFF
39Y9310	Ethernet Expansion Card (CFFv)	CFFv
39Y9271	NetXen 10 Gb Ethernet Expansion Card	CFFh
39Y9306	QLogic Ethernet and 4 Gb Fibre Channel Expansion Card	CFFh
41Y8527	QLogic 4 Gb Fibre Channel Expansion Card	CFFv
73P6000	Myrinet Cluster Expansion Card[a]	StFF
32R1896	Cisco Systems InfiniBand 1X HCA Expansion Card	StFF
32R1923	QLogic iSCSI Expansion Card	StFF
26R0890	QLogic 4 Gb SFF Fibre Channel Expansion Card	SFF
39Y9186	Emulex 4 Gb SFF Fibre Channel Expansion Card	SFF
43W6859	Emulex 4 Gb Fibre Channel Expansion Card (CFFv)	CFFv
32R1760	Cisco Systems InfiniBand 4X HCA Expansion Card	HSFF
39Y9190	SAS Expansion Card (CFFv)	CFFv

Part number	Description	Form Factor
43W4420	Voltaire InfiniBand 4X DDR Expansion Card (CFFh)	CFFh
43W4421	Cisco InfiniBand 4X DDR Expansion Card (CFFh)	CFFh
43W4423	InfiniBand 4X DDR Expansion Card (CFFh)	CFFh

a. The Myrinet card is a single-channel card and requires you to install an OPM (Optical Pass-thru module) in BladeCenter I/O module bay 4.

PCI Express I/O Expansion Unit (PEU3e)

The LS21 supports the attachment of one PCI Express I/O Expansion Unit. The PEU3e allows standard PCI-Express type adapters to be attached to the LS21 type 7971.

For more information about the PEU3e, see 3.18.5, “PCI Express I/O Expansion Unit (PEU3e)” on page 368

The LS21 type 7971 does not support the attachment of the PCI I/O Expansion Unit 2 (or any previous PCI I/O expansion unit).

3.8 LS41 type 7972

The LS41 type 7972 is available in two form factors:

- A single slot blade server with two CPU sockets (LS41 base unit)
- A double slot blade server with four CPU sockets (LS41 base unit plus a Multi Processor Expansion unit)

The single slot blade server can also be upgraded to a double slot blade server, with four CPU sockets, by installing a Multi Processor Expansion (MPE) unit.

The Multi Processor expansion (MPE) unit does not include any processors or memory as standard. These need to be added separately. The processors will need to match the existing processors on the LS41 main unit. The memory should be split evenly between the two units.

The double slot LS41 models come standard with two processors. Two additional processors and additional memory need to be ordered separately to make the LS41 a fully configured 4 processor blade server.

The LS41 blade server supports up to four of the dual-core 8200 series AMD Opteron processors.

Note: The single slot LS41 is functionally similar to the LS21, except that the LS21 type 7971 cannot be upgraded to a four CPU-socket blade server like the single slot LS41 can. The LS41 uses the dual core Opteron 8000 series microprocessor while the LS21 uses the Opteron 2000 series microprocessors.

Table 3-48 lists the features of the LS41 type 7972.

Table 3-48 Features of the LS41 type 7972

	LS41 Single Slot	LS41 Double Slot
LS41 Model	51Y	52Y, 6BY, 7BY, 8BY

	LS41 Single Slot	LS41 Double Slot
Form Factor	30 mm blade (expandable to 60 mm using MPE unit)	60 mm blade
Processors (Std/Max)	1/2	2/4
Upgradeable with MPE	Yes	No
Memory (Std/Max)	2 GB/32 GB	4 GB/64 GB
DIMM sockets	8	16
Gigabit Ethernet Controllers Standard	2 x Broadcom BCM5706S	2 x Broadcom BCM5706S and 2 x Broadcom BCM5708S
2.5-inch drive bays	1	2
Adapter slots total	1 PCI-X slot or 1 PCI Express slot	2 PCI-X slots or 1 PCI Express slot

All the models of the LS41 type 7972 have the following features as standard:

- Integrated Baseboard management controller (BMC)
- A LSI Logic 53C1064 SAS controller for onboard hard disk drives
- One PCI Express connector for the attachment of a High Speed Daughter Card (HSDC) or Expansion unit
- Support for the Storage and I/O expansion (SIO) unit
- Light Path Diagnostics on system board
- Predictive failure analysis (PFA)

Note: The single slot LS41 with the MPE unit installed, has the same features as the double slot LS41.

Table 3-49 lists the available models of the LS41 blade server type 7972.

Table 3-49 Models of the LS41 type 7972 at a glance

Type-Model	AMD CPU Name	CPU Speed	Internal L2 cache[a]	Processor Quantity Std/Max	Memory Std/Max	Processor power draw[b]	Blade Width
7972-51Y	Opteron 8216 HE	2.4 GHz	1 MB	1/2	2 GB/32 GB	68 W	Single
7972-52Y	Opteron 8216 HE	2.4 GHz	1 MB	2/4	4 GB/64 GB	68 W	Double
7972-6BY	Opteron 8218	2.6 GHz	1 MB	2/4	4 GB/64 GB	95 W	Double
7972-7BY	Opteron 8220	2.8 GHz	1 MB	2/4	4 GB/64 GB	95 W	Double
7972-8BY	Opteron 8222	3.0 Ghz	1 MB	2/4	4 GB/64 GB	95 W	Double

a. 1 MB per core for 8200 series.

b. Versions of the LS41 that use the 95 W processors (LS41 - 6AY, 3BY, 6BY, 8BY) are designed for use in the BladeCenter H chassis. Due to increased power and thermal needs, these 95W versions will run at 'n-2' processor speed in the BladeCenter and BladeCenter T chassis. In most cases this represents a speed reduction of 400 MHz. The versions of the LS41 that use the High Efficiency (HE) 68W processors (LS41 - 31Y, 51Y, 32Y, 52Y) will run at full speed in all chassis.

Figure 3-44 shows a double slot blade server LS41 type 7972 (base unit and MPE unit).

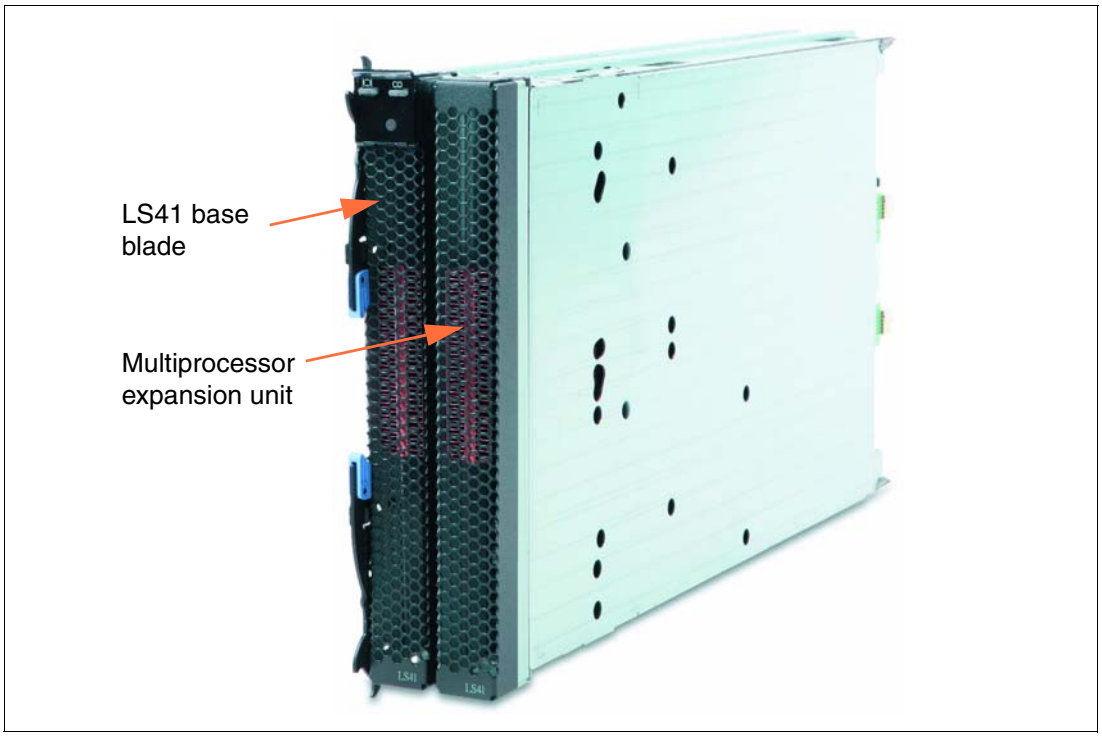

Figure 3-44 Double slot blade server LS41 type 7972 (base unit and MPE unit)

3.8.1 Processor

The LS41 type 7972 supports up to four of the dual-core 8200 series AMD Opteron processors.

In a two-socket and four-socket configuration, all processors must be identical Opteron processors (same clock rate, cache size, and the same number of cores).

The processors used in these blades are standard and low power, full performance Opteron processors. The standard AMD Opteron processors draw a maximum of 95 W. Specially manufactured low-power processors operate at 68 W or less without any performance trade-offs. This savings in power at the processor level combined with the smarter power solution that BladeCenter delivers make these blades very attractive to clients who have limited power and cooling resources.

Table 3-50 shows blade server LS41 type 7972 processor options.

Table 3-50 Blade server LS41 type 7972 processor options

Part number	Processor description	Models
25R8935	AMD Opteron Dual Core Processor 8216 HE 2.4 GHz 1 MB L2 Cache 68W	51Y, 52Y
25R8933	AMD Opteron Dual Core Processor 8218 2.6 GHz, 1 MB L2 Cache, 95W	6BY
40K1266	AMD 2.8 GHz/1 GHz Opteron Dual Core Processor Model 8220, 95W	7BY
44R4951	3.0 GHz / 1 GHz Opteron Dual Core Processor Model 8222, 95W	8BY

3.8.2 Multi Processor Expansion (MPE) Unit

The single slot LS41blade server with two CPU sockets can be upgraded to a double slot blade server with four CPU sockets with the addition of a multi processor expansion (MPE) unit. The LS41 base unit and the Multi Processor Expansion (MPE) unit can be joined by simply clipping the two units together and tightening a thumb screw.

When the MPE unit is connected to the single slot LS41 type 7972, it offers identical features to the double slot LS41 type 7972 blade server.

The Multi Processor Expansion (MPE) unit adds the following features:

- Two additional CPU sockets
- Eight additional memory DIMM sockets
- Two additional Gigabit Ethernet controllers (2 x Broadcom BCM5708S)
- One additional SAS hard disk drive slot
- One additional SFF or StFF I/O expansion slot

Notes:

- Only one Multi Processor Expansion (MPE) unit can be installed on a single slot LS41 type 7972 base unit. The double slot LS41 type 7972 already has a MPE unit installed and cannot have any additional MPE units installed.
- If a Multi Processor Expansion (MPE) unit is installed on an LS41 type 7972 base unit and no processors are installed on the MPE unit, then only the disk drive socket on the MPE will be functional. All other functions on the MPE unit require at least one processor to be installed.

3.8.3 Memory

The memory used in the LS41 type 7972 has the following features:

- PC2-5300 (667 Mhz) DDR2 RDIMM
- also supports PC2-4200 (533 MHz) DDR2 memory
- ECC
- Very Low Profile (VLP)
- Chipkill

Memory configuration has to follow these guidelines:

- There are a total of four DIMM slots per processor socket. Two of these slots are preconfigured with a pair of DIMMs in a single slot LS41 type 7972. In a double slot LS41 type 7972, four DIMM slots are preconfigured with memory DIMMs.
- Because memory is 2-way interleaved, the memory modules must be installed in matched pairs. However, one DIMM pair is not required to match the other in capacity.

For each installed microprocessor, a set of 4 DIMM sockets are enabled. See Table 3-51 for more information.

Table 3-51 Enabled DIMM sockets

Installed processors	Enabled DIMM sockets
1 (in LS41 base unit)	1, 2, 3, 4 (in LS41 base unit)
1 and 2 (in LS41 base unit)	1, 2, 3, 4, 5, 6, 7, 8 (in LS41 base unit)
1 and 2 (in LS41 base unit) 1 (in MPE unit)	1, 2, 3, 4, 5, 6, 7, 8 (in LS41 base unit) 1, 2, 3, 4 (in MPE unit)

Installed processors	Enabled DIMM sockets
1 and 2 (in LS41 base unit) 1 and 2 (in MPE unit)	1, 2, 3, 4, 5, 6, 7, 8 (in LS41 base unit) 1, 2, 3, 4, 5, 6, 7, 8 (in MPE unit)

See Table 3-52 for blade server LS41 type 7972 memory options.

Table 3-52 Blade server LS41 type 7972 memory options

Part number	Memory description
39M5861	1 GB (2x512 MB Kit) PC2-5300 DDR2 SDRAM VLP RDIMM
39M5864	2 GB (2x1GB Kit) PC2-5300 DDR2 SDRAM VLP RDIMM
39M5867	4 GB (2x2 GB Kit) PC2-5300 DDR2 SDRAM VLP RDIMM
46C0512	4 GB (2 x 2 GB Kit) PC2-5300 DDR2 667 MHz Single Rank VLP RDIMM
39M5870	8 GB (2x 4 GB Kit) PC2-4200 CL4 DDR2 SDRAM VLP RDIMM
46C0513	8 GB (2x 4 GB) PC2-5300 CL5 DDR2 SDRAM VLP RDIMM

Table 3-53 shows the standard LS41 type 7972 DIMM configuration by model.

Table 3-53 Standard LS41 type 7972 DIMM configuration by model

Model	Standard DIMM configuration	Chipkill
51Y	2 GB (2x1 GB Kit) PC2-5300 DDR2 SDRAM VLP RDIMM	Yes
52Y, 6BY, 7BY, 8BY	4 GB (4x1 GB Kit) PC2-5300 DDR2 SDRAM VLP RDIMM	Yes

3.8.4 Onboard network controllers

The LS41 type 7972 base unit has two Gigabit Ethernet Broadcom BCM5706S controllers. The MPE unit has two additional Gigabit Ethernet Broadcom BCM5708S controllers. These Ethernet interfaces are routed through to the redundant midplane providing a foundation for fault-tolerant network setups.

Note: When installed in a BladeCenter E, H, T and HT chassis, Ethernet Port 1 on the blade server and Port 1 on the MPE are routed to the switch module in Bay 1 and Ethernet port 2 on the blade server and port 2 on the MPE are routed to the switch module in Bay 2.

When installed in BladeCenter S, both Ethernet ports on the blade server and MPE are routed to the switch module in Bay 1.

The following features are supported by the Broadcom (BCM5706S) onboard controllers and by the Broadcom (BCM5708S) controllers on the MPE unit:

- TOE (TCP Offload Engine)
- iSCSI initiator
- PXE 2.0 boot support
- Wake on LAN
- Load-balancing
- VLAN tagging support
- Jumbo frame support (9 KB)

The Broadcom Advanced Server Program (BASP) is management software provided by Broadcom for managing your network cards. This includes a command line interface for scripting purposes.

3.8.5 Integrated systems management processor

The blade server LS41 type 7972 has an integrated BMC that provides:

- IPMI compliance (Intelligent Platform Management Interface)
- Serial over LAN (SOL)
- Wake on LAN (WOL)
- Power management functions
- KVM and media tray ownership
- Light path diagnostics support
- Automatic BIOS recovery (ABR)
- Automatic server restart (ASR)
- Predictive failure analysis (PFA) on system components (CPUs, memory, and drives)
- Inventory
- Error logging
- Environmental monitoring (voltages and temperature on the system board)

Note that you cannot access this onboard BMC externally through a dedicated port or daisy-chain it to other servers or blades. The BMC is forwarding information to the Management Module (MM) within the BladeCenter chassis. The MM is a single point of control for hardware management of all server and switch components contained within the chassis.

3.8.6 Local storage options

The single slot LS41 type 7972 blade server base unit features an onboard LSI Logic 53C1064 SAS controller with a SAS connector for one optional 2.5-inch Small Form Factor (SFF) SAS hard disk drive. The multi processor expansion (MPE) unit has a SAS connector for an additional 2.5-inch SFF SAS hard disk drive. Therefore in a double slot LS41 type 7972 blade server (LS41 base unit and MPE unit) it is possible to have two SAS hard disk drives in a RAID 0 or RAID 1 array.

Table 3-54 shows the SFF HDD options for LS41 type 7972.

Table 3-54 Small form factor (SFF) hard disk drive options for LS41 type 7972

Part number	Description
26K5777	73 GB SAS HDD (2.5") SFF
42D0421	146 GB SAS 10K SFF NHS
43X0845	73 GB SAS 15K SFF NHS

The hard disk drive is mounted onto the system board of the blade server and are inaccessible from the outside during operation.

Solid state drives

The LS41 supports the attachment of up to two solid state drives (SSD), one on the LS41 base unit and one on the MPE. The SSD device is installed in the disk drive bay. If an SSD device is to be used, then a SAS disk drive cannot be installed because they both use the

same drive bay. When two of these SSD devices are installed, the onboard LSI Logic 53C1064 SAS controller can create a RAID 0 or RAID 1 array.

These SSD devices are designed for operating system boot and low-bandwidth data transmission. SSD devices use less power than a standard SAS drive and therefore also generate less heat.

Table 3-55 shows the SSD options for the LS41.

Table 3-55 Solid State Drive options for the LS41

Part number	Description
43W7614	IBM 15.8 GB 2.5” Solid State Drive
43W7618	IBM 31.4 GB 2.5" Solid State Drive

Note: Unlike the HS21 XM, the LS41 does not support the 15.8 GB dual drive option, part number 43W7606.

BladeCenter Storage and I/O Expansion Unit (SIO)

The LS41 type 7972 supports the attachment of one BladeCenter Storage and I/O expansion unit. For more information about the SIO, see 3.18.2, “Storage and I/O Expansion Unit” on page 363.

When the SIO is used with the LS41, the ServeRAID 8k-l in the SIO (or ServeRAID 8k, if installed) can form arrays using combinations of SAS drives in both the SIO and in the LS41. This means that you can form one RAID array of four or five drives — three in the SIO and one or two in the blade server (depending on whether you have the MPE connected).

Note: The LS41 type 7972 does not support the attachment of any of the previous BSE-I or BSE-II expansion units.

3.8.7 I/O expansion options

The single slot LS41 type 7972 base unit has one onboard connector for various StFF, SFF, or CFFv I/O expansion cards including network and storage adapters. This onboard connector is connected via the midplane to module bays 3 and 4 at the rear of the BladeCenter chassis. The interface of the I/O expansion slot is a PCI-X 133 MHz bus.

The single slot LS41 type 7972 base unit also has a PCI Express connector that allows for the attachment of a HSFF expansion card, a CFFh expansion card, or for the attachment of a supported expansion unit.

The double slot LS41 type 7972 with the MPE unit attached has two onboard connectors (one on the LS41 base unit and one on the MPE unit) for attaching various StFF, SFF, or CFFv I/O expansion card options. It also has one PCI Express connector available on the MPE unit that allows for the attachment of either a HSFF expansion card, a CFFh expansion card, or a supported expansion unit.

Note: With a double slot LS41 with a MPE unit attached, the PCI Express connector on the base unit is used to connect the LS41 base unit to the MPE unit. Therefore it is only possible to install one HSFF or CFFh expansion card, and in this instance it must be installed on the PCI Express connector on the MPE unit.

The LS41 type 7972 supports five different form factors for expansion cards that can connect to either one or the other of these two types of connector. See 3.1, “Expansion cards” on page 234, for further information about form factors and expansion card options.

Table 3-56 summarizes the available I/O options.

Table 3-56 Blade server LS41 type 7972 expansion options

Part number	Description	Form factor
39R8624	IBM BladeCenter Gigabit Ethernet Expansion Card	SFF
26K4841	BladeCenter Fibre Channel Expansion Card	SFF
39Y9310	Ethernet Expansion Card (CFFv) for IBM BladeCenter	CFFv
39Y9271	NetXen 10 Gb Ethernet Expansion Card	CFFh
39Y9306	QLogic Ethernet and 4 Gb Fibre Channel Expansion Card	CFFh
41Y8527	QLogic 4 Gb Fibre Channel Expansion Card	CFFv
73P6000	Myrinet Cluster Expansion Card[a]	StFF
32R1896	Cisco Systems InfiniBand 1X HCA Expansion Card	StFF
32R1923	QLogic iSCSI Expansion Card	StFF
26R0890	QLogic 4 Gb SFF Fibre Channel Expansion Card	SFF
39Y9186	Emulex 4 Gb SFF Fibre Channel Expansion Card	SFF
43W6859	Emulex 4 Gb Fibre Channel Expansion Card (CFFv)	CFFv
32R1760	Cisco Systems InfiniBand 4X HCA Expansion Card	HSFF
39Y9190	SAS Expansion Card (CFFv)	CFFv
43W4420	Voltaire InfiniBand 4X DDR Expansion Card (CFFh)	CFFh
43W4421	Cisco InfiniBand 4X DDR Expansion Card (CFFh)	CFFh
43W4423	InfiniBand 4X DDR Expansion Card (CFFh)	CFFh

a. The Myrinet card is a single-channel card and requires you to install an OPM in BladeCenter I/O module bay 4.

PCI Express I/O Expansion Unit (PEU3e)

The LS41 supports the attachment of one PCI Express I/O Expansion Unit. The PEU3e allows standard PCI-Express type adapters to be attached to the LS41 type 7972.

For more information about the PEU3e, see 3.18.5, “PCI Express I/O Expansion Unit (PEU3e)” on page 368

The LS41 type 7972 does not support the attachment of the PCI I/O Expansion Unit 2 (or any previous PCI I/O expansion unit).

3.9 LS22 type 7901

The LS22 type 7901 is a 30mm blade (single slot blade server), capable of supporting up to two quad core AMD Opteron 2300 series processors and up to 32 GB DDR2 memory (in 8 DIMM slots with 4 GB DIMM modules) with Chipkill protection, for high performance and

reliability. The latest AMD Opteron processors are designed with 2 MB of L2 cache and another 6 MB of L3 cache, 64-bit extensions and HyperTransport technology, to help provide the computing power you require to match your business needs and growth.

The LS22 type 7901 supports up to two non-hot swap 2.5-inch SAS HDDs or two non-hot swap solid state drives. A dual port Gigabit Ethernet controller is standard, providing high-speed data transfers and offering TCP Offload Engine (TOE) support, load-balancing and failover capabilities.

With optional expansion cards, the LS22 can connect to additional Ethernet, Myrinet, Fibre Channel, InfiniBand and other high-speed communication switches housed in the chassis. Optional expansion cards can add additional fabrics to the LS22 as needed.

Figure 3-45 shows the components inside the LS22 type 7901 blade server.

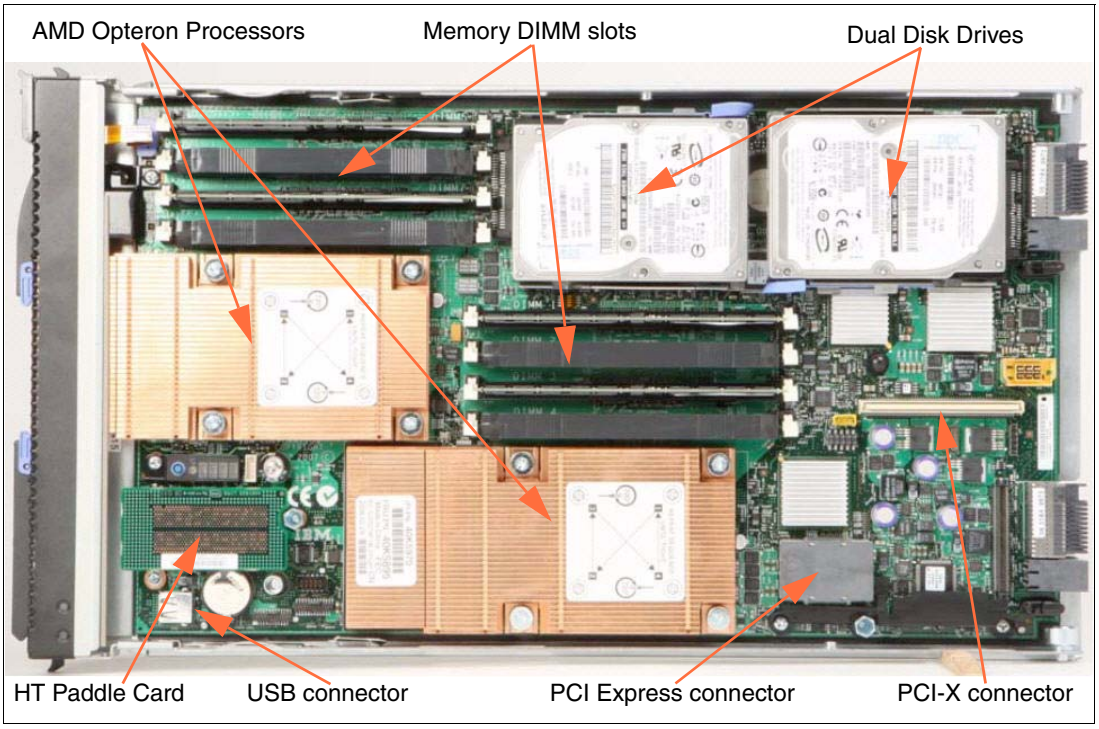

Figure 3-45 Blade server LS22 type 7901 with its top cover removed

All models of the LS22 type 7901 have the following features as standard:

- One standard AMD Opteron 2300 series processor (upgradeable to two)
- One dual port Gigabit Ethernet Broadcom 5709S controller
- An ATI Radeon™ ES1000 video controller
- Internal USB connector
- An integrated BMC
- An LSI Logic 1064e SAS controller for onboard storage
- Support for up to two 2.5-inch internal fixed disk drives
- One PCI-X connector slot for an StFF, SFF, or CFFv I/O expansion card
- One PCI Express connector for the attachment of a HSFF expansion card, CFFh expansion card, or Expansion unit
- Support for the PCI Express I/O expansion (PEU3e) unit and PCI expansion (PEU2) unit II

- Light Path Diagnostics on system board
- IBM exclusive Memory Booster (Hyper-Transport Paddle Card)

Table 3-57 lists the available models of the LS22 type 7901 blade server.

Table 3-57 Blade servers LS22 type 7901 at a glance

Type-Model	Processor[a]	CPU speed	L2 /L3 cache[b]	Processors Std/Max	Memory Std/Max	Processor power draw
7901-3Qx	AMD Opteron 2347 HE	1.9 GHz	2 MB/2 MB	1/2	2 GB/32 GB	55 W ACP
7901-CQx	AMD Opteron 2356	2.3 GHz	2 MB/2 MB	1/2	2 GB/32 GB	75 W ACP
7901-3Sx	AMD Opteron 2384	2.7 GHz	2 MB/6 MB	1/2	2 GB/32 GB	75 W ACP

a. Only versions of the LS22 that use the High Efficiency (HE) processors can operate in the BladeCenter E chassis.
b. L2 cache is 512 KB per core for 2300 series. The L3 cache is shared.

3.9.1 Processor

The LS22 type 7901 supports up to two AMD Opteron 2300 series quad-core processors.

Table 3-58 lists the blade server LS22 type 7901 processor options.

Table 3-58 Blade server LS22 type 7901 processor options

LS21 model	Part number	Processor model
7901-3Qx	44X1542	AMD Opteron 2347 HE (1.9 GHz/79w)
7901-CQx	44R4961	AMD Opteron 2356 (2.3 GHz/115w)
7901-3Sx	46M6808	AMD Opteron 2384 (2.7 GHz/95w)

When installing two processors, they must be identical (same clock rate and cache size, and the same number of cores).

3.9.2 Memory

The memory used in the LS22 type 7901 has the following features:

- PC2-5300 (667 MHz) DDR2 RDIMM
- PC2-6400 (800 MHz) DDR2 RDIMM
- ECC
- Very Low Profile (VLP)
- Chipkill

Memory configuration must follow these guidelines:

- There are a total of eight DIMM slots (four per processor socket). Two of these slots (1 and 2) are preconfigured with a pair of DIMMs.
- Because memory is 2-way interleaved, the memory modules must be installed in matched pairs. However, one DIMM pair is not required to match the other in capacity.
- The maximum of 32 GB of installed memory is achieved when all DIMM sockets are populated with the 4 GB DIMMs.
- All DIMM attached to the one processor must be the same speed.

For each installed microprocessor, a set of 4 DIMM sockets is enabled. Table 3-59 lists the enabled DIMM sockets information.

Table 3-59 Enabled DIMM sockets

Installed Microprocessors	Enabled DIMM sockets
1	1, 2, 3, 4
1 and 2	1, 2, 3, 4, 5, 6, 7, 8

Table 3-60 lists the blade server LS22 type 7901 memory options.

Table 3-60 Blade server LS22 type 7901 memory options

Part number	Memory description
46C0522	2 GB (2 x 1 GB) PC2-5300 CL5 ECC DDR2 667Mhz VLP RDIMM
46C0512	4 GB (2 x 2 GB) PC2-5300 CL5 ECC DDR2 667MHz SR VLP RDIMM
46C7524	8 GB (2 x 4 GB) PC2-5300 CL5 ECC DDR2 667Mhz VLP RDIMM
46C7525	8 GB (2 x 4 GB) PC2-6400 CL6 ECC DDR2 800Mhz VLP RDIMM

Table 3-61 lists the standard LS22 type 7901 DIMM configuration by model.

Table 3-61 Standard LS22 7901 DIMM configuration by model

Model	Standard DIMM configuration	Chipkill
3Qx, CQx, 3Sx	2 GB (2x 1 GB) PC2-5300 DDR2 SDRAM VLP RDIMM	Yes

HyperTransport Paddle Card - *Memory Booster*

The LS22 type 7901 ships standard with the HyperTransport Paddle Card (HT Paddle Card). The HT Paddle Card (which is also known as the Memory Booster) improves remote memory access performance by utilizing HyperTransport connections that would otherwise be unused. Remote memory access refers to when a CPU needs to access memory connected to another CPU in the same server. The HT Paddle Card connects to the HyperTransport Expansion connector on the LS22 motherboard as shown in Figure 3-45 on page 314.

NUMA-aware applications may not see much performance increase with the HT Paddle Card because most memory access in a NUMA-aware application is kept local.

3.9.3 Onboard network controllers

The LS22 type 7901 blade server uses one dual port Gigabit Ethernet Broadcom 5709S controller. The two Ethernet ports are routed through to the redundant midplane, providing a foundation for fault-tolerant network setups.

Note: When installed in a BladeCenter E, H, T, and HT chassis, Ethernet Port 1 on the blade server is routed to the switch module in Bay 1 and Ethernet port 2 is routed to the switch module in Bay 2.

When installed in BladeCenter S, both Ethernet ports on the blade server are routed to the switch module in Bay 1.

The following features are supported by the Broadcom (BCM5709S) onboard controller:

- TCP/IP Offload Engine (TOE) IPv4 and IPv6 support (operating system dependant)
- iSCSI initiator
- PXE 2.0 boot support
- Wake on LAN (WOL)
- Load-balancing
- VLAN 802.1q tagging support
- Jumbo frame support (9 KB)
- Serial over LAN (SOL) support

The Broadcom Advanced Server Program (BASP) is management software provided by Broadcom for managing your network cards. This includes a command-line interface for scripting purposes.

3.9.4 Integrated systems management processor

The LS22 type 7901 has an integrated baseboard management controller (BMC) that provides:

- IPMI compliance (Intelligent Platform Management Interface)
- Serial over LAN (SOL)
- Wake on LAN (WOL)
- Power management functions
- KVM and media tray ownership
- Light path diagnostics support
- Automatic BIOS recovery (ABR)
- Automatic server restart (ASR)
- Predictive failure analysis (PFA) on system components (CPU, memory, and disk drives)
- Inventory
- Error logging
- Environmental monitoring (voltages and temperature on the system board)

Note that you cannot access this onboard BMC externally through a dedicated port or daisy-chain it to other servers or blades. The BMC forwards information to the Management Module (MM) within the BladeCenter chassis. The MM is a single point of control for hardware management of all server and switch components contained within the chassis.

3.9.5 Local storage options

The LS22 type 7901 blade server features an onboard LSI Logic 1064e SAS controller with two SAS connectors for attaching either two non-hot swap 2.5-inch SAS or SATA hard disk drives or two solid state drives (SSD).

Table 3-62 shows the SFF HDD options for LS22 type 7901.

Table 3-62 Hard disk drive options for LS22 type 7901

Part number	Description
26K5777	73 GB SAS HDD (2.5") SFF
42D0421	146 GB SAS 10K SFF
43X0845	73 GB SAS 15K SFF
43W7666	300 GB 10K SATA 2.5" NHS HDD

The hard disk drives are mounted onto the system board of the blade server and are inaccessible from the outside during operation.

Solid state drives

The LS22 type 7901 supports the attachment of up to two solid state drives (SSD). These devices are installed in the existing SAS hard drive bays. They present to the operating system as SATA devices. If a solid state drive is to be used, then an SAS disk drive cannot be installed because they both use the same drive bays. When two SSD devices are installed, a RAID array can be created using the onboard SAS controller. A single solid state drive cannot offer RAID support. A RAID array cannot be created using a solid state drive and a SAS drive together.

These solid state drive devices are designed for operating system boot and low-bandwidth data transmission. SSD devices use less power than a standard SAS drive and, therefore, also generate less heat.

Table 3-63 shows the solid state drive options for the LS22.

Table 3-63 Solid State Drive options for the LS22

Part number	Description
43W7614	IBM 15.8 GB 2.5” Solid State Drive
43W7618	IBM 31.4 GB 2.5" Solid State Drive

Note: Unlike the HS21 XM, the LS22 does not support the 15.8 GB dual drive option, part number 43W7606.

BladeCenter Storage and I/O Expansion Unit (SIO)

The LS22 type 7901 does not currently support the attachment of any storage expansion unit.

3.9.6 I/O expansion options

The LS22 type 7901 has an onboard PCI-X connector for attaching various I/O expansion card options, including network and storage adapters. These adapters are connected via the midplane to module bays 3 and 4 at the rear of the BladeCenter chassis. The interface of the I/O expansion slot is a PCI-X 133 MHz bus.

The LS22 type 7901 also has a PCI Express connector that allows for the attachment of an HSFF expansion card, a CFFh expansion card or a supported expansion unit.

The installation of an StFF adapter in the LS22 type 7901 will mean that the second disk drive bay cannot be used.

See 3.1, “Expansion cards” on page 234 for further information about form factors and expansion card options.

Table 3-64 summarizes the available I/O options.

Table 3-64 Blade server LS22 type 7901 expansion options

Part number	Description	Form Factor
39R8624	IBM BladeCenter Gigabit Ethernet Expansion Card	SFF
39Y9310	Ethernet Expansion Card (CFFv)	CFFv

Part number	Description	Form Factor
39Y9271	NetXen 10 Gb Ethernet Expansion Card	CFFh
39Y9306	QLogic Ethernet and 4 Gb Fibre Channel Expansion Card	CFFh
41Y8527	QLogic 4 Gb Fibre Channel Expansion Card	CFFv
73P6000	Myrinet Cluster Expansion Card[a]	StFF
43W6859	Emulex 4 Gb Fibre Channel Expansion Card (CFFv)	CFFv
32R1760	Cisco Systems InfiniBand 4X HCA Expansion Card	HSFF
39Y9190	SAS Expansion Card (CFFv)	CFFv
43W4420	Voltaire InfiniBand 4X DDR Expansion Card (CFFh)	CFFh
43W4421	Cisco InfiniBand 4X DDR Expansion Card (CFFh)	CFFh
43W4423	InfiniBand 4X DDR Expansion Card (CFFh)	CFFh
44W4465	Broadcom 10 Gb Ethernet CFFh Expansion Card	CFFh
44W4466	Broadcom 2-Port 10 Gb Ethernet CFFh Expansion Card	CFFh
44W4479	2/4 Port Ethernet Expansion Card	CFFh
44X1940	QLogic Ethernet and 8 Gb Fibre Channel Expansion Card	CFFh

a. The Myrinet card is a single-channel card and requires you to install an Optical Pass-thru module (OPM) in BladeCenter I/O module bay 4.

PCI I/O Expansion Unit 2 (PEU2)

The LS22 type 7901 supports the attachment of one PCI I/O Expansion Unit 2. The PEU2 allows supported standard PCI-X type adapters to be attached to the LS22 type 7901.

For more information about the PEU2, see 3.18.4, “BladeCenter PCI I/O Expansion Unit 2” on page 367.

PCI Express I/O Expansion Unit (PEU3e)

The LS22 type 7901 supports the attachment of one PCI Express I/O Expansion Unit. The PEU3e allows supported standard PCI-Express type adapters to be attached to the LS22 type 7901.

For more information about the PEU3e, see 3.18.5, “PCI Express I/O Expansion Unit (PEU3e)” on page 368

3.10 LS42 type 7902

The LS42 type 7902 is a high throughput, highly scalable four-socket-capable AMD Opteron-based blade server. The LS42 type 7902 is available as a four-socket server or as a two-socket server with the ability to upgrade to a four-socket server by purchasing and installing a Multi Processor Expansion (MPE) unit.

When fully configured, the LS42 type 7902 can support up to 64 GB of DDR2 ECC memory in 16 DIMM slots and up to four of the quad core 8300 series AMD Opteron processors.

The LS42 type 7902 supports up to two SAS or SATA disk drives or up to two solid state drives. The disk drives are internally fixed to the motherboard and cannot be removed without removing the LS42 type 7902 from the BladeCenter chassis. The onboard LSI 1064e SAS controller provides mirroring capability for these internal drives.

A dual port Gigabit Ethernet controller is standard in both the base unit and the multi-processor expansion unit for a total of four built-in gigabit Ethernet ports offering TCP/IP Offload Engine (TOE) support, load-balancing, and failover capabilities. Optional I/O Expansion Cards add additional fabrics to the LS42 as needed. Each LS42 type 7902 blade server can connect to additional Ethernet, SAS, Fibre Channel, iSCSI, InfiniBand and other high-speed communication switches housed in the chassis.

The LS42 type 7902 is available in two form factors:

1. A single slot blade server with two CPU sockets (LS42 base unit)
2. A double slot blade server with four CPU sockets (LS42 base unit plus a Multi Processor Expansion unit)

The single slot LS42 blade server models are fully functional two CPU socket blade servers. The single slot LS42 blade server models can also be upgraded to a double slot blade server, with four CPU sockets, by installing a Multi Processor Expansion (MPE) unit. The MPE includes two processors. All that is required for a fully functional four CPU socket blade server is to install memory in the MPE. The single slot LS42 blade server offers investment protection by offering “pay as you grow” flexibility.

The double slot LS42 blade server models come standard with two processors installed on the base unit sockets. The included Multi Processor Expansion (MPE) unit has two empty CPU sockets. Two additional processors and additional memory need to be ordered separately to make the LS42 a fully functional four CPU socket blade server. The processors will need to match the existing processors on the LS42 base unit.

Note: The single blade slot LS42 is functionally similar to the LS22, except that the LS22 type 7901 cannot be upgraded to a four CPU-socket blade server as the single slot LS42 can. The LS42 uses the Opteron 8300 series processors. The LS22 uses the Opteron 2300 series processors.

Table 3-65 lists the features of the LS42 type 7902.

Table 3-65 Features of the LS42 type 7902

	LS42 single slot (30mm)	**LS42 double slot (60mm)**
LS42 Model	1Sx, 3Sx	3Qx, CQx
Form Factor	30 mm blade (expandable to 60 mm using MPE unit)	60 mm blade
Processors (Std/Max)	2/2[a]	2/4
Upgradeable with MPE	Yes	No
Memory (Std/Max)	4 GB/32 GB[b]	4 GB/64 GB
DIMM sockets	8[c]	16
Gigabit Ethernet Controllers Standard	1 x Broadcom BCM5709S[d]	2 x Broadcom BCM5709S
2.5-inch drive bays	2	2

	LS42 single slot (30mm)	LS42 double slot (60mm)
Adapter slots total	1 PCI-X slot[e] or 1 PCI Express slot	2 PCI-X slots or 1 PCI Express slot

a. Maximum value can be increased with the installation of the MPE unit. The Max Processor quantity with the MPE installed is four.
b. Maximum value can be increased with the installation of the MPE unit. The Max Memory with the MPE installed is 64 GB.
c. Maximum values can be increased with the installation of the MPE unit. The Max DIMM slots with the MPE install is 16.
d. Maximum values can be increased with the installation of the MPE unit. The number of Ethernet controllers with the MPE installed is two.
e. Maximum values can be increased with the installation of the MPE unit. The number of PCI-X adapter slots with the MPE installed is two.

All models of the LS42 type 7902 have the following features as standard:

- Integrated Baseboard management controller (BMC)
- An LSI Logic 1064e SAS controller for onboard disk drives
- One PCI Express connector for the attachment of a CFFh expansion card or supported expansion unit
- One PCI-X connector for the attachment of I/O expansion cards
- Light Path Diagnostics on system board
- Predictive failure analysis (PFA)
- Internal USB connector

Note: The single slot LS42 with the MPE unit installed has the same features as the double slot LS42.

Table 3-66 lists the available models of the LS42 blade server type 7902.

Table 3-66 Models of the LS42 type 7902 at a glance

Type-Model	AMD CPU Name[a]	CPU Speed	Internal L2/L3 cache[b]	Processor Quantity Std/Max	Memory Std/Max	Processor power draw	Blade Width
7902-3Qx	Opteron 8347 HE	1.9 GHz	2MB/2MB	2/4	4 GB/64 GB	55 W ACP	Double
7902-CQx	Opteron 8356	2.3 GHz	2MB/2MB	2/4	4 GB/64 GB	75 W ACP	Double
7902-1Sx	Opteron 8380	2.5 GHz	2MB/6MB	2/2[c]	4 GB/32 GB[c]	75 W ACP	Single
7972-3Sx	Opteron 8384	2.7 GHz	2MB/6MB	2/2[c]	4 GB/32 GB[c]	75 W ACP	Single

a. Only versions of the LS42 that use the High Efficiency (HE) processors can operate in the BladeCenter E chassis.
b. L2 cache is 512 KB per core for 8300 series. The L3 cache is shared.
c. Maximum values can be increased with the installation of the MPE unit. The Max Processor quantity with the MPE installed is four. Max Memory with the MPE installed is 64 GB.

Figure 3-46 shows LS42 type 7902 blade server base unit.

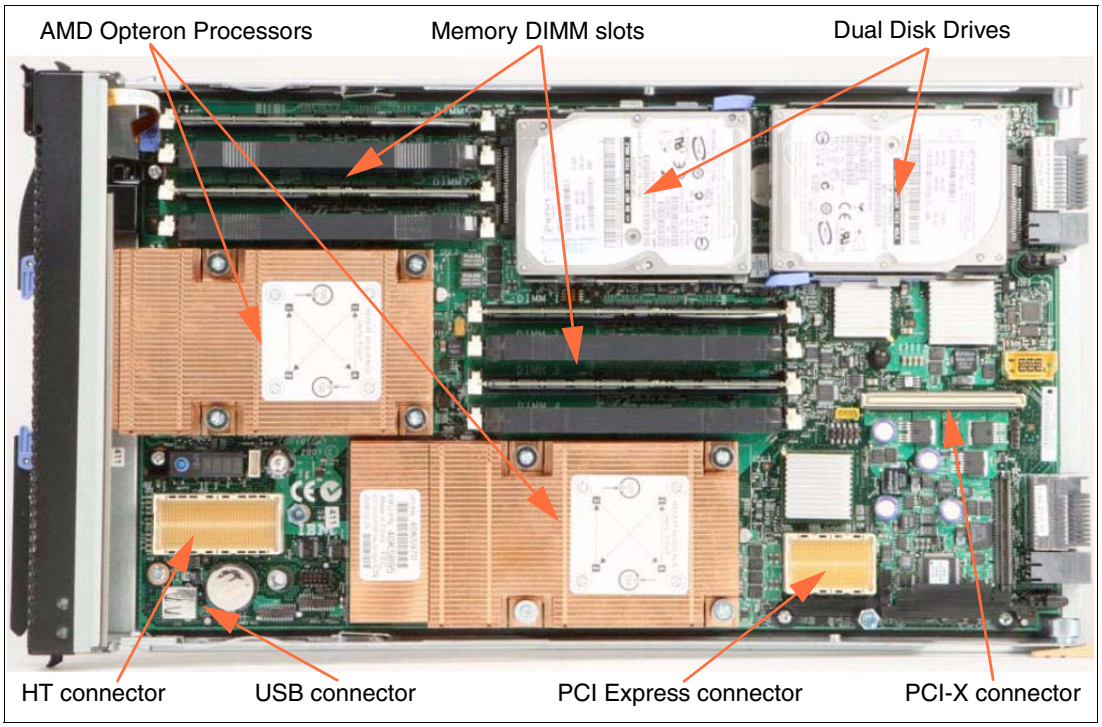

Figure 3-46 The LS42 type 7902 base unit

3.10.1 Processor

The LS42 type 7902 supports up to four of the quad core AMD Opteron 8300 series processors. In a two-socket and four-socket configuration, all processors must be identical Opteron processors (same clock rate and cache size, and the same number of cores).

AMD released the first of their quad core 8300 series Opteron processors using a 65nm manufacturing process. The internal AMD codename given to these 65nm 8300 series processors was *Barcelona*. AMD then released further versions of the quad core 8300 series Opteron processors, but this time they where based on a 45nm manufacturing process. The internal AMD codename given to these 45nm 8300 series processors was *Shanghai*.

The AMD Opteron *Shanghai* processors are designed with 2MB of L2 cache (512KB per processor core) and 6 MB of shared L3 cache. The AMD Opteron *Barcelona* processors are designed with 2 MB of L2 cache (512KB per processor core) and 2MB of shared L3 cache.

Note: The *Shanghai* and *Barcelona* processors have a slightly different feature set. For further information regarding AMD Opteron processor features, refer to:

`http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_8796,00.html`

Table 3-67 on page 323 shows blade server LS42 type 7902 processor options.

Table 3-67 Blade server LS42 type 7902 processor options

Part number	Processor description	Model
44X1552[a]	AMD Opteron Quad Core Processor 8347 HE 1.9 GHz 2MB L2 2MB L3 Cache	3Qx
44R4971[a]	AMD Opteron Quad Core Processor 8356 2.3 GHz 2MB L2 2MB L3 Cache	CQx

a. AMD Opteron 8300 series CPU using a 65nm manufacturing process.

Only the LS42 type 7902 models listed in the third column of Table 3-67 can be upgraded using the individual processors specified in the table. In contrast, the LS42 type 7902 models 1Sx and 3Sx can only be upgraded by installing a MPE unit that also includes the appropriate processors as listed in Table 3-68 on page 324.

3.10.2 Multi Processor Expansion (MPE) Unit

The single slot LS42 type 7902 blade server with two CPU sockets can be upgraded to a double slot blade server with four CPU sockets with the addition of a Multi Processor Expansion (MPE) unit. The LS42 base unit and the MPE unit can be joined by simply clipping the two units together and tightening a thumb screw.

The LS42 type 7902 models 3Qx and CQx are double-wide blade servers and come standard with a MPE unit already attached. These models simply require additional processors and memory to create a functional four-socket server.

The LS42 type 7902 models 1Sx and 3Sx are single-wide two-socket blade servers and can be upgraded to a four-socket server with the attachment of the MPE unit. The MPE unit for these models comes standard with processors already installed. The correct MPE unit needs be installed for each model because the processors on the MPE unit need to match the processors on the LS42 base unit. All that is required to create a functional four-socket server, after the MPE unit is installed, is to add additional memory to the MPE unit.

When the MPE unit is connected to the single slot LS42 type 7902, it offers identical features to the double slot LS42 type 7902 blade server.

The MPE unit has the following features:

- Two additional CPU sockets
- Eight additional memory DIMM sockets
- One additional dual port Gigabit Ethernet controllers (Broadcom BCM5709S)
- One additional CFFv I/O expansion slot

Figure 3-47 on page 324 shows LS42 type 7902 MPE unit.

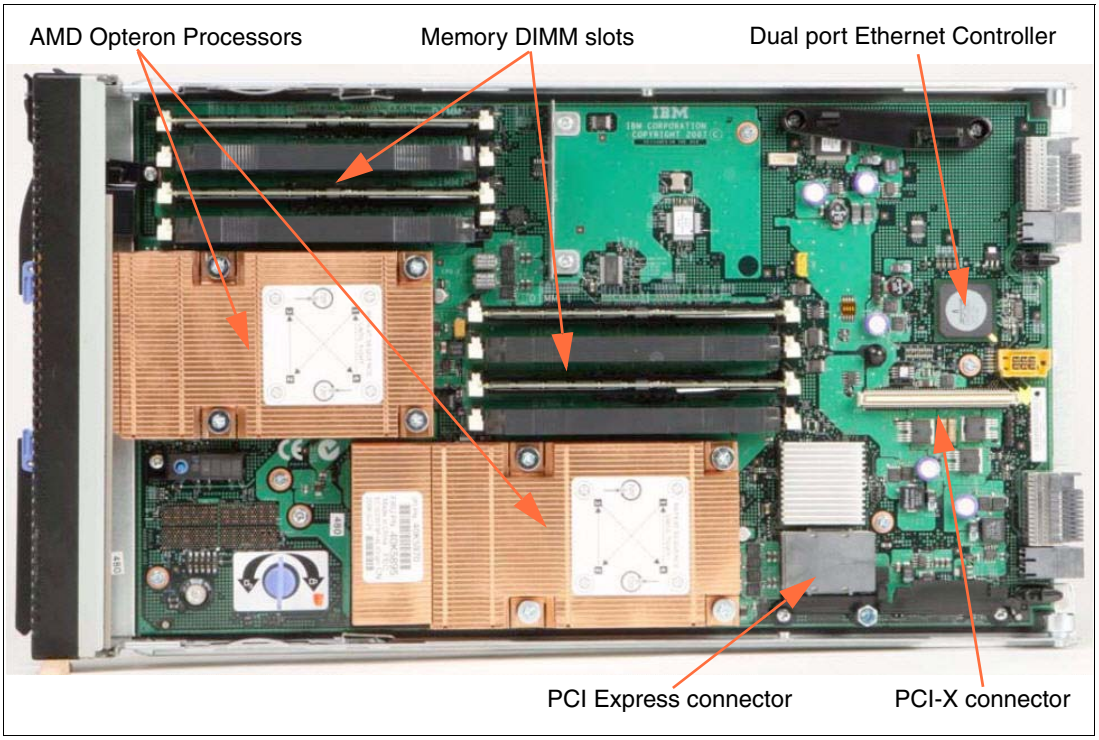

Figure 3-47 The LS42 Multi Processor Expansion (MPE) unit

Table 3-68 Blade server LS42 type 7902 Multi Processor Expansion unit options

Part number	Processor description	Supported model
46M6817	LS42 Multi Processor Expansion Unit: Processor Model 8380[a] (2 x 2.5GHz)	1Sx
44X1685	LS42 Multi Processor Expansion Unit: Processor Model 8384[a] (2 x 2.7GHz)	3Sx

a. AMD Opteron 8300 series CPU using a 45nm manufacturing process.

Only the LS42 type 7902 models listed in the Supported model column of Table 3-68 can be upgraded by installing an MPE unit. In contrast, the LS42 type 7902 models 3Qx and CQx come standard with an MPE already installed that requires installation of individual processors and memory, as listed in Table 3-67 on page 323.

Note: Only one Multi Processor Expansion (MPE) unit can be installed on a single slot LS42 type 7902 base unit. The double slot LS42 type 7902 already has a MPE unit installed and cannot have any additional MPE units installed.

3.10.3 Memory

The memory used in the LS42 type 7902 has the following features:

- PC2-5300 (667 MHz) DDR2 RDIMM
- PC2-6400 (800 MHz) DDR2 RDIMM
- ECC
- Very Low Profile (VLP)
- Chipkill

Memory configuration must follow these guidelines:

- There are a total of four DIMM slots per processor socket. A single slot LS42 type 7902 will have four of these slots preconfigured with two pairs of memory DIMMs. In a double slot LS42 type 7902, four DIMM slots are also preconfigured with two pairs of memory DIMMs, but the attached MPE has no DIMMs installed.
- Because memory is 2-way interleaved, the memory modules must be installed in matched pairs. However, one DIMM pair is not required to match the other in capacity.
- The maximum of 64 GB of installed memory is achieved when all 16 DIMM sockets of a double slot LS42 are populated with the 4 GB DIMMs.
- All DIMM attached to the one processor must be the same speed.

For each installed microprocessor, a set of 4 DIMM sockets are enabled. See Table 3-69 for more information.

Table 3-69 Enabled DIMM sockets

Installed processors	Enabled DIMM sockets
1 (in LS42 base unit)	1, 2, 3, 4 (in LS42 base unit)
1 and 2 (in LS42 base unit)	1, 2, 3, 4, 5, 6, 7, 8 (in LS42 base unit)
1 and 2 (in LS42 base unit) 1 (in MPE unit)	1, 2, 3, 4, 5, 6, 7, 8 (in LS42 base unit) 1, 2, 3, 4 (in MPE unit)
1 and 2 (in LS42 base unit) 1 and 2 (in MPE unit)	1, 2, 3, 4, 5, 6, 7, 8 (in LS42 base unit) 1, 2, 3, 4, 5, 6, 7, 8 (in MPE unit)

See Table 3-70 for blade server LS42 type 7902 memory options.

Table 3-70 Blade server LS42 type 7902 memory options

Part number	Memory description
46C0522	2 GB (2 x 1 GB) PC2-5300 CL5 ECC DDR2 667Mhz VLP RDIMM
46C0512	4 GB (2 x 2 GB) PC2-5300 CL5 ECC DDR2 667MHz SR VLP RDIMM
46C7524	8 GB (2 x 4 GB) PC2-5300 CL5 ECC DDR2 667Mhz VLP RDIMM
46C7525	8 GB (2 x 4 GB) PC2-6400 CL6 ECC DDR2 800Mhz VLP RDIMM

Table 3-71 shows the standard LS42 type 7902 DIMM configuration by model.

Table 3-71 Standard LS42 type 7902 DIMM configuration by model

Model	Standard DIMM configuration	Chipkill
3Qx, CQx, 1Sx, 3Sx	4 GB (4 x 1 GB) PC2-5300 CL5 ECC DDR2 667Mhz VLP RDIMM	Yes

HyperTransport Paddle Card - *Memory Booster*

Like the LS22, the single-slot (30 mm) LS42 models ship standard with the HyperTransport Paddle Card (HT Paddle Card). The HT Paddle Card (which is also known as the Memory Booster) improves remote memory access performance by utilizing HyperTransport connections that would otherwise be unused. Remote memory access refers to when a CPU needs to access memory connected to another CPU in the same server. The HT Paddle Card connects to the HyperTransport Expansion connector (similar to the LS22 as shown in per Figure 3-45 on page 314).

NUMA-aware applications may not see much performance increase with the HT Paddle Card because most memory access in a NUMA-aware application is kept local.

3.10.4 Onboard network controllers

The LS42 type 7902 base unit has one dual port Gigabit Ethernet Broadcom 5709S controller. The MPE unit has an additional dual port Gigabit Ethernet Broadcom 5709S controller. These Ethernet ports are routed through to the redundant midplane, providing a foundation for fault-tolerant network setups.

Note: When installed in a BladeCenter E, H, T, or HT chassis, Ethernet Port 1 on the blade server and Port 1 on the MPE are routed to the switch module in Bay 1, and Ethernet port 2 on the blade server and port 2 on the MPE are routed to the switch module in Bay 2.

When installed in BladeCenter S, both Ethernet ports on the blade server and MPE are routed to the switch module in Bay 1.

The following features are supported by the Broadcom (BCM5709S) onboard controller:

- TCP/IP Offload Engine (TOE) IPv4 and IPv6 support (operating system-dependent)
- iSCSI initiator
- PXE 2.0 boot support
- Wake on LAN (WOL)
- Load-balancing
- VLAN 802.1q tagging support
- Jumbo frame support (9 KB)
- Serial over LAN (SOL) support

The Broadcom Advanced Server Program (BASP) is management software provided by Broadcom for managing your network cards. This includes a command-line interface for scripting purposes.

3.10.5 Integrated systems management processor

The blade server LS42 type 7902 has an integrated BMC that provides:

- Intelligent Platform Management Interface (IPMI) compliance
- Serial over LAN (SOL)
- Wake on LAN (WOL)
- Power management functions
- KVM and media tray ownership
- Light path diagnostics support
- Automatic BIOS recovery (ABR)
- Automatic server restart (ASR)
- Predictive failure analysis (PFA) on system components (CPUs, memory, and drives)
- Inventory
- Error logging
- Environmental monitoring (voltages and temperature on the system board)

Note that you cannot access this onboard BMC externally through a dedicated port or daisy-chain it to other servers or blades. The BMC forwards information to the Management Module (MM) within the BladeCenter chassis. The MM is a single point of control for hardware management of all server and switch components contained within the chassis.

3.10.6 Local storage options

The LS42 type 7902 blade server base unit features an onboard LSI Logic 1064e SAS controller with two SAS connectors for attaching either two non-hot swap 2.5-inch SAS or SATA hard disk drives or two non-hot swap solid state drives (SSD). The multi-processor expansion (MPE) unit has no disk drive connectors and does not support local storage attachment. The onboard LSI Logic 1064e SAS controller provides RAID support for the attached local storage.

Table 3-72 shows the SFF HDD options for LS42 type 7902.

Table 3-72 Small form factor (SFF) hard disk drive options for LS42 type 7902

Part number	Description
26K5777	73 GB SAS HDD (2.5") SFF
42D0421	146 GB SAS 10K SFF
43X0845	73 GB SAS 15K SFF
43W7666	300 GB 10K SATA 2.5" NHS HDD

The hard disk drives are mounted onto the system board of the blade server and are inaccessible from the outside during operation.

Solid state drives

The LS42 type 7902 supports the attachment of up to two solid state drives (SSD) on the LS42 blade server base unit. The SSD device is installed in one of the disk drive bays. When two of these SSD devices are installed, the onboard LSI Logic 1064e SAS controller can create a RAID 0 or RAID 1 array.

These SSD devices are designed for operating system boot and low-bandwidth data transmission. SSD devices use less power than a standard SAS drive and therefore also generate less heat.

Table 3-73 shows the SSD options for the LS42.

Table 3-73 Solid State Drive options for the LS42

Part number	Description
43W7614	IBM 15.8 GB 2.5” Solid State Drive
43W7618	IBM 31.4 GB 2.5" Solid State Drive

Note: Unlike the HS21 XM, the LS42 does not support the 15.8 GB dual drive option, part number 43W7606.

BladeCenter Storage and I/O Expansion Unit (SIO)

The LS42 type 7902 does not currently support the attachment of any storage expansion unit.

3.10.7 I/O expansion options

The single slot LS42 type 7902 base unit has one PCI-X onboard connector for various SFF or CFFv I/O expansion cards including network and storage adapters. This onboard connector is connected via the midplane to module bays 3 and 4 at the rear of the BladeCenter chassis. The interface of the I/O expansion slot is a PCI-X 133 MHz bus.

The single slot LS42 type 7902 base unit also has a PCI Express connector that allows for the attachment of either a HSFF expansion card, a CFFh expansion card, or for the attachment of a supported expansion unit.

The double slot LS42 type 7902 with the MPE unit attached has two PCI-X onboard connectors (one on the LS42 base unit and one on the MPE unit) for attaching various SFF or CFFv I/O expansion card options. It also has one PCI Express connector available on the MPE unit that allows for the attachment of either a HSFF expansion card, a CFFh expansion card, or a supported expansion unit.

Note: With a double slot LS42 with an MPE unit attached, the PCI Express connector on the base unit is used to connect the LS42 base unit to the MPE unit. Therefore, it is only possible to install one CFFh expansion card, and in this instance it must be installed on the PCI Express connector on the MPE unit.

The LS42 type 7902 supports five different form factors for expansion cards that can connect to either one or the other of these two types of connector. The installation of a StFF adapter in the LS42 type 7902 will mean that the second disk drive bay cannot be used. See 3.1, “Expansion cards” on page 234, for further information about form factors and expansion card options.

Table 3-74 summarizes the available I/O options.

Table 3-74 Blade server LS42 type 7902 expansion options

Part number	Description	Form factor
39R8624	IBM BladeCenter Gigabit Ethernet Expansion Card	SFF
39Y9310	Ethernet Expansion Card (CFFv) for IBM BladeCenter	CFFv
39Y9271	NetXen 10 Gb Ethernet Expansion Card	CFFh
39Y9306	QLogic Ethernet and 4 Gb Fibre Channel Expansion Card	CFFh
41Y8527	QLogic 4 Gb Fibre Channel Expansion Card	CFFv
73P6000	Myrinet Cluster Expansion Card[a]	StFF
43W6859	Emulex 4 Gb Fibre Channel Expansion Card (CFFv)	CFFv
32R1760	Cisco Systems InfiniBand 4X HCA Expansion Card	HSFF
39Y9190	SAS Expansion Card (CFFv)	CFFv
43W4420	Voltaire InfiniBand 4X DDR Expansion Card (CFFh)	CFFh
43W4421	Cisco InfiniBand 4X DDR Expansion Card (CFFh)	CFFh
43W4423	InfiniBand 4X DDR Expansion Card (CFFh)	CFFh
44W4466	Broadcom 2-Port 10 Gb Ethernet CFFh Expansion Card	CFFh
44X1940	QLogic Ethernet and 8 Gb Fibre Channel Expansion Card	CFFh

a. The Myrinet card is a single-channel card and requires you to install an OPM in BladeCenter I/O module bay 4.

PCI I/O Expansion Unit 2 (PEU2)

The LS42 type 7902 supports the attachment of one PCI I/O Expansion Unit 2. The PEU2 allows supported standard PCI-X type adapters to be attached to the LS42 type 7902.

For more information about the PEU2, see 3.18.4, “BladeCenter PCI I/O Expansion Unit 2” on page 367

PCI Express I/O Expansion Unit (PEU3e)

The LS42 type 7902 supports the attachment of one PCI Express I/O Expansion Unit (PEU3e). The PEU3e allows supported standard PCI-Express type adapters to be attached to the LS42 type 7902.

For more information about the PEU3e, see 3.18.5, “PCI Express I/O Expansion Unit (PEU3e)” on page 368

3.11 JS12 Type 7998-60x

The IBM BladeCenter JS12 Express is the latest high-density blade server with POWER6 technology with EnergyScale™ technology. The JS12 is a single socket with two cores per socket, 3.8 GHz POWER6 SCM processor. The processor is directly mounted to the planar board and includes 64 KB I-cache and 64 KB D-cache L1 cache per core with 4 MB L2 cache per core. The integrated memory controller supports eight pluggable registered DIMMs, which must be installed in pairs. The minimum memory that can be installed is 2 GB (2 X 1 GB) and the maximum is 64 GB (8 x 8 GB). The JS12 supports Memory Scrubbing, ECC, Chipkill, and Bit Steering.

As a heterogeneous infrastructure consolidation platform, all the BladeCenter Chassis support running the JS12 alongside your Intel processor-based HS blades, AMD Opteron processor-based LS blades, and Cell/B.E. processor-based QS21 blades in the same chassis with independent monitoring, security, power, and systems management. And with several gigabit Ethernet, Fibre Channel, iSCSI, and InfiniBand high-performance expansion cards to choose from, the JS12 can be tailored to the demands of your data center network and application workloads.

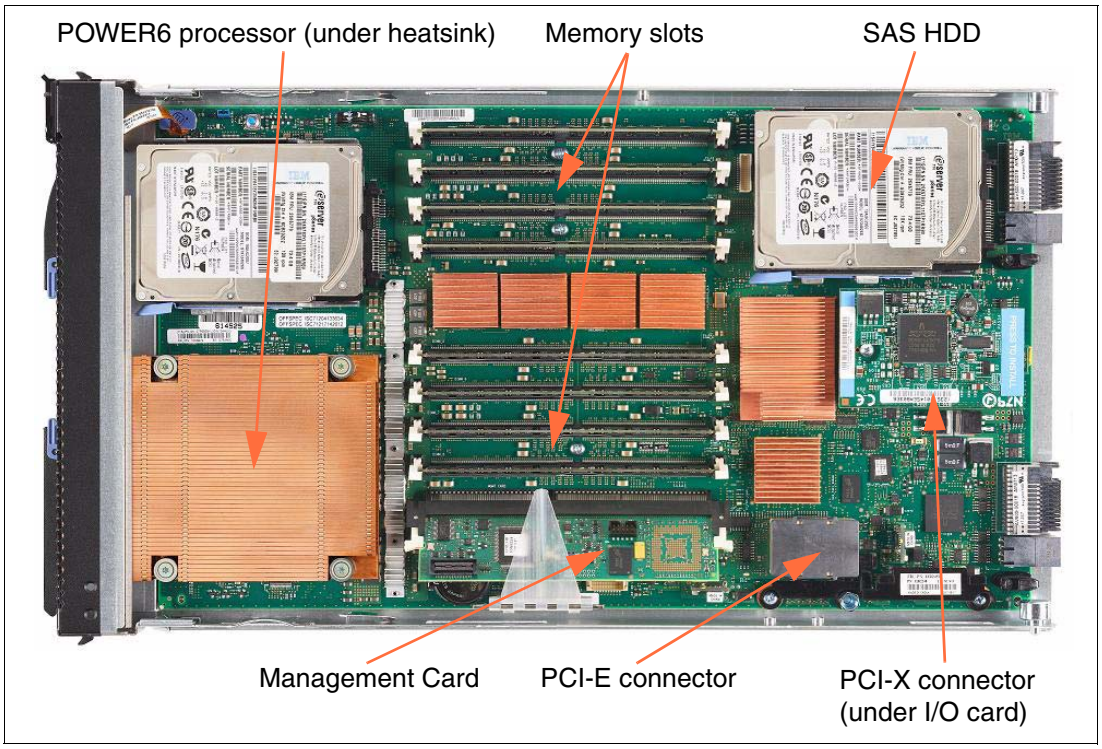

Figure 3-48 JS12 internals

The JS12 has autonomic power and thermal controls designed to maintain operation in out-of-specification situations. If thermal or power conditions become unacceptable, it may reduce the processor or memory performance capability to maintain acceptable conditions. When the conditions are restored to an acceptable value, the blade will again automatically resume normal operation. Some applications may be sensitive to these capability changes. Contact your application vendors to determine any possible impacts.

Full keyboard, video, and mouse support can be used at installation instead of Serial over LAN (SOL), which requires a network switch module in the chassis. Multiple console connections are also supported by simultaneously using the BladeCenter keyboard, video, and mouse, and a SOL connection. The USB bus on the BladeCenter chassis is used for keyboard support on the JS12. The Advanced Management Module (aMM) on the chassis controls the routing of keyboard inputs to one of up to 14 blade servers. An embedded video controller is used to support 2D graphics on an SVGA video display. The JS12 has four USB ports. Two of these ports interface with the keyboard and mouse through the aMM, and two interface with a removable media tray also through the aMM.

As a robust blade solution for 64-bit UNIX® applications, the JS12 also delivers leadership single instruction multiple-data (SIMD) capabilities with the Linux operating system for scientific research and HPC.

3.11.1 JS12 models

Table 3-83 on page 341 lists the models of the BladeCenter JS12.

Table 3-75 JS12 models

7998-60X	Processor	L2 cache	Memory	Ethernet	HDD
#8442	Single socket dual-core 3.8 GHz POWER6	4 MB/core	2 GB - 64GB	Dual Gigabit Ethernet	Open (two bays for fixed SAS drives)

3.11.2 Minimum system configuration

Each 7998-60X configuration must contain a minimum of:

- One IBM BladeCenter JS12 2-core, 3.8 GHz (#8442)
- One Memory DIMMs (at least one of these):
  - 2 GB (2 x 1 GB DIMMS) DDR2 667 MHz (#8220)
  - 4 GB (2 x 2 GB DIMMS) DDR2 667 MHz (#8229)
  - 8 GB (2 x 4 GB DIMMs) DDR2 667 MHz (#8239)
  - 16 GB (2 x 8 GB DIMMs) DDR2 400 MHz (#8245)
- Two Processor Entitlements (one of these)
  - 2 x #8444
  - 1 x #8444 and 1 x #8443 (Express Product Offering)

Notes:

- For the IBM i operating system on the JS12, one SAS Expansion Card (#8250) is required.
- A minimum of one copy of *Service Warranty Publications* (#8737) and one copy of *JS12 Installation and User's Guide* (#8742-8752, #8756) is required at each customer installation.

3.11.3 Express configurations

If you order a JS12 server Express Product Offering, as defined here, you may qualify for a processor entitlement at no additional charge. A minimum of 4 GB of memory and one SAS 10 K SFF HDD is required to determine whether a customer is entitled to a processor entitlement at no additional charge.

You can make changes to the standard features as needed and still qualify for processor entitlements at no additional charge. However, selection of total memory or DASD smaller than the totals defined as the minimums disqualifies the order as an Express Product Offering.

You must purchase PowerVM Standard Edition (#5406) or PowerVM Enterprise Edition (#5606) for each Express Product Offering.

- One IBM BladeCenter JS12 2-core, 3.8 GHz (#8442)
- 4 GB memory DIMMs with at least 73 GB SAS HDD
- One of these:
  - Two 2 GB (2 x 1 GB DIMMS) DDR2 667 MHz (#8220)
  - One 4 GB (2 x 2 GB DIMMS) DDR2 667 MHz (#8229)
  - One 8 GB (2 x 4 GB DIMMS) DDR2 667 MHz (#8239)
  - One 16 GB (2 x 8 GB DIMMS) DDR2 400 MHz (#8245)

Additional memory DIMMs of any size might be added.

- One of these:
  - Two PowerVM Standard Edition (#5406)
  - Two PowerVM Enterprise Edition (#5606)
- One of these:
  - One IBM 73 GB SAS 10 K SFF HDD (#8237)
  - One IBM 146 GB SAS 10 K SFF HDD (#8236)
- One Processor Entitlement (#8444)
- One Zero-priced Processor Entitlement (#8443); receive one processor entitlement (#8443) at no additional charge. Two processors are entitled.

Note: A minimum of one copy of *Service Warranty Publications* (#8737) and one copy of *JS12 Installation and User's Guide* (#8742-8752, #8756) is required at each customer installation.

3.11.4 Memory and hard disk drives

- Up to 64 GB of system memory
  - The JS12 system board contains eight RDIMM memory connections.
  - Memory must be installed in matched pairs with a minimum of 2 GB and maximum of 64 GB.
  - Memory sizes can be mixed.
  - JS12 supports Memory Scrubbing, ECC, Chipkill, and Bit Steering.
  - 2 GB (2 x 1 GB DIMMs), 4 GB (2 x 2 GB DIMMs), and 8 GB (2 x 4 GB DIMMS) memory features operate at 667 MHz. The 16 GB (2 x 8 GB DIMMs) memory feature operates at 400 MHz.
- SAS Controller, supporting up to 146 GB
  - Internal disk drive, SFF (2.5 in) 73 GB SAS (#8237)
  - Internal disk drive, SFF (2.5 in) 146 GB SAS (#8236)

3.11.5 Advanced Power Virtualization

The JS12 has support for PowerVM technology built into the microprocessor. The JS12 is designed to make it more affordable to consolidate multiple independent applications on a single blade using the same proven virtualization technologies offered on IBM Power servers.

The JS12 supports two leading-edge virtualization technologies:

- PowerVM Standard Edition (feature #5406)

  PowerVM Standard Edition makes the JS12 an ideal platform for consolidation of AIX, Linux, and IBM i operating system applications, thus helping clients reduce infrastructure complexity and cost. Offering an intuitive, Web-based interface for managing virtualization within a single blade, the Integrated Virtualization Manager (IVM) component of Virtual I/O Server (VIOS) allows the small business IT manager to quickly and easily set up and manage logical partitions (LPARs). It also enables Virtual I/O and Virtual Ethernet so that storage and communications adapters can be shared among all the LPARs running on the JS12. Ultimately, IBM Micro-Partitioning™ technology allows each processor core to be subdivided into as many as 10 virtual servers. And because the JS12 is built with POWER6 technology, other advanced virtualization functions such as Shared Dedicated Capacity may be exploited.

- PowerVM Enterprise Edition (feature #5606)

  PowerVM Enterprise Edition includes all the features of PowerVM Standard Edition plus a new capability called Live Partition Mobility. Live Partition Mobility allows for the movement of a running AIX or Linux partition from one POWER6 processor-based server to another with no application downtime, resulting in better system utilization, improved application availability, and potential energy savings. With Live Partition Mobility, planned application downtime due to regular server maintenance can be a thing of the past. Software Maintenance for Virtual I/O Server (577x-PVE) must be purchased with VIOS (5765-PVE). PowerVM Enterprise Edition must be purchased separately.

Note: PowerVM Standard Edition and PowerVM Enterprise Edition are optional when running AIX or Linux. Power VM Standard or Enterprise Edition is required when running the IBM i operating system on the JS12.

3.11.6 Expansion cards

The form factor of the expansion card determines whether an additional expansion card or SAS internal HDD can be installed in the rear drive bay in the JS12.

- The QLogic iSCSI Expansion Card for IBM BladeCenter (#8238) is a standard form factor (StFF) expansion card. The installation of this card precludes support for the SAS HDD. The maximum allowed is one.
- One high-speed daughter card (HSFF) expansion card may be installed on the JS12 along with one SAS HDD in the rear drive bay. The Cisco Systems 4X InfiniBand HCA Expansion Card for IBM BladeCenter (#8256) is an HSFF expansion card.
- One combination form factor vertical (CFFv) may be installed with a combination form factor horizontal (CFFh) expansion card and one SAS HDD in the rear drive bay.

  Supported CFFv expansion cards include:

  – QLogic 4 GB 2 Port Fibre Channel Expansion Card (#8248)
  – Ethernet Expansion Card (CFFv) (#8249)
  – SAS Expansion Card (#8250)
  – Emulex 4 GB Fibre Channel Expansion Card (#8251)

  Supported CFFh expansion cards include:

  – QLogic Ethernet and 4 GB Fibre Channel Expansion Card (#8252)
  – 4X InfiniBand DDR Expansion Card (#8258)
  – Voltaire 4X InfiniBand DDR Expansion Card (#8298)

Not all of the I/O expansion cards available for the JS12 are supported by all of the operating systems that run on the JS12. See Table 3-76 for a list of I/O expansion cards supported by each operating system.

Table 3-76 JS12 Expansion Card compatibility

Expansion Card name	FC	IBM i V6R1	VIOS 1.5	AIX 5.3	AIX 6.1	SLES 10 SP2	RHEL 5.1	RHEL 4.6
SAS Expansion Card, CFFv	8250	Y	Y	Y	Y	Y	Y	Y
Emulex 4 GB FC dual port, CFFv	8251	N	Y	Y	Y	Y	Y	Y
QLogic 4 GB FC and 1 GB Ethernet CFFh	8252	N	Y	Y	Y	Y	Y	Y

Expansion Card name	FC	IBM i V6R1	VIOS 1.5	AIX 5.3	AIX 6.1	SLES 10 SP2	RHEL 5.1	RHEL 4.6
QLogic 2-port 4 GB Fibre Channel CFFv	8248	N	Y	Y	Y	Y	Y	Y
Broadcom 1 GB Ethernet dual port, CFFv	8249	N	Y	Y	Y	Y	Y	Y
QLogic iSCSI dual port, StFF	8238	N	Y	Y	Y	Y	Y	Y
4X IB DDR + Voltaire SW license, CFFh	8298	N	N	N	N	N	Y	N
4X IB Dual Port DDR,CFFh (Mellanox)	8258	N	N	Y	Y	N	Y	N
4X IB Dual Port, HSFF (Cisco)	8256	N	N	Y	Y	N	Y	Y

Note: The IBM i operating system only supports the SAS Expansion Card (CFFv) for IBM BladeCenter (#8250).

3.11.7 Software requirements

The JS12 server supports the AIX, Linux, and IBM i operating systems, providing the flexibility of using applications written for any one of the three. IBM has qualified two popular Linux distributions for use with the JS12 server:

- Red Hat Enterprise Linux for POWER
- Novell SUSE Linux Enterprise Server for POWER

One or more of the following operating systems is required for an operational JS12:

- AIX V5.3 with the 5300-06 Technology Level with Service Pack 7, or later
- AIX V5.3 with the 5300-07 Technology Level with Service Pack 4, or later
- AIX V5.3 with the 5300-08 Technology Level, or later
- AIX V6.1 with the 6100-01 Technology Level, or later (Available 6/10/08)
- AIX V6.1 with the 6100-00 Technology Level with Service Pack 5, or later
- SUSE Linux Enterprise Server 10 Service Pack 2 for POWER, or later
- Red Hat Enterprise Linux V4.6 for POWER, or later
- Red Hat Enterprise Linux V5.1 for POWER, or later
- iV6R1, or later

Note that iSCSI boot is available through an iSCSI software initiator running on the embedded Ethernet controller or the iSCSI expansion card (#8238). The AIX operating system supports iSCSI boot through an iSCSI software initiator on any supported Ethernet controller including Ethernet expansion cards.

3.11.8 Integrated systems management processor

The enhanced BMC for the JS12 blade server is a flexible service processor which provides support for the following functions:

- Alert Standard Format (ASF) and RS-485 interfaces to the management modules

- Intelligent Platform Management Interface (IPMI)
- Power control and advanced power management
- Reliability, availability, and serviceability (RAS) features
- Serial over LAN (SOL)
- Continuous health monitoring and control
- Configurable notification and alerts
- Event logs that are time-stamped and saved in nonvolatile memory and that can be attached to e-mail alerts
- Point-to-point protocol (PPP) support
- Remote power control
- Remote firmware update and access to critical blade server settings

3.11.9 Systems management and administrative tools

Integrated diagnostic and administrative tools, such as IBM Predictive Failure Analysis and light path diagnostics, are designed to simplify administration to help lower costs and improve control of the IT environment. Remote management capabilities allow automation of IT networking tasks. Proven technology like VIOS allows sharing of disk drives, communications, and Fibre Channel adapters.

Following are systems management tools for the JS12:

- AIX 5.3 or AIX 6: IBM Director V5.2 and Cluster Systems Management (CSM) for AIX, V1.7
- SUSE Linux Enterprise Server 10 Service Pack 2 for POWER, or later; IBM Director V5.2 and CSM for Linux on POWER, V1.7
- Red Hat Enterprise Linux V4.6, for POWER, or later: IBM Director V5.2 and CSM for Linux on POWER, V1.7
- Red Hat Enterprise Linux V5.1, for POWER, or later: IBM Director V5.2 and CSM for Linux on POWER, V1.7
- iV6R1, or later: IBM Director V5.2

3.12 JS21 type 8844

The IBM BladeCenter JS21 type 8844 delivers what cost-conscious clients demand with a 64-bit PowerPC 970MP solution. Along with optional support for PowerVM, the JS21 BladeCenter offers a robust blade server solution for high-performance computing (HPC), AIX 5L, Linux on POWER, and server consolidation.

Key features of the JS21 include the following:

- Two 64-bit PowerPC 970MP processors
- Up to 16 GB ECC Chipkill DDR2 memory
- Dual Gigabit Ethernet controller (Broadcom 5780)
- One PCI-X connection for an I/O expansion card
- One PCI Express connection for a high speed expansion card
- Light Path Diagnostics on system board
- Integrated systems management processor
- Integrated SAS controller with RAID 0 and 1

See Table 3-77 for an overview of JS21 type 8844 models.

Table 3-77 Blade servers JS21 type 8844 at a glance

Model	CPU Speed/L2 Cache	CPU Qty (Std/Max)	Memory (Std/Max)	Standard SAS disk
8844-31X	2-way single core 2.7 GHz[a]/1 MB	2/2	1 GB/16 GB	None
8844-41X	2-way dual-core 2.3 GHz/1+1 MB	2/2	2 GB/16 GB	None
8844-51X	2-way dual core 2.5 GHz[a]/1+1 MB	2/2	2 GB/16 GB	None
8844-E3X	2-way single core 2.7 GHz[a]/1 MB	2/2	2 GB/16 GB	1x 73 GB
8844-E5X	2-way dual core 2.5 GHz[a]/1+1 MB	2/2	2 GB/16 GB	1x 73 GB

a. Processor frequency is dependent upon which BladeCenter chassis the JS21 is installed in. See 3.12.1, "Processor" on page 336.

3.12.1 Processor

The BladeCenter JS21 servers use PowerPC 970MP processors either single core or dual core. All JS21 models have two processors installed which are directly mounted to the planar board.

For some models, the processor speed is lowered in the BladeCenter 8677 chassis and BladeCenter T chassis. See Table 3-78.

Table 3-78 JS21 processor frequencies when installed in each BladeCenter chassis

Model	Nominal frequency	Frequency in BladeCenter E	Frequency in BladeCenter H	Frequency in BladeCenter T
8844-31X	2.7 GHz	2.6 GHz	Full speed (2.7 GHz)	2.6 GHz
8844-41X	2.3 GHz	Full speed (2.3 GHz)	Full speed (2.3 GHz)	Full speed (2.3 GHz)
8844-51X	2.5 GHz	2.3 GHz	Full speed (2.5 GHz)	2.3 GHz
8844-E3X	2.7 GHz	2.6 GHz	Full speed (2.7 GHz)	2.6 GHz
8844-E5X	2.5 GHz	2.3 GHz	Full speed (2.5 GHz)	2.3 GHz

3.12.2 Memory

The JS21 type 8844 models support a maximum of 16 GB ECC Chipkill DDR2 memory when using 4 Gb DIMM modules. Each server has four DIMM sockets and memory is 2-way interleaved and must be installed in matched pairs. The base JS21 comes with two DIMM slots populated. Memory sizes can be mixed, but DIMMs must be all PC2-3200 (400 MHz DDR2) or all PC2-4200 (533 MHz DDR2). No mixing of different speed DIMM modules is allowed.

Table 3-79 JS21 memory configuration

Model	Standard memory	Memory speed
8844-31X	1 GB (2x 512 MB	400 MHz (PC2-3200)
8844-41X	2 GB (2x 1 GB)	533 MHz (PC2-4200)
8844-51X	2 GB (2x 1 GB)	400 MHz (PC2-3200)
8844-E3X	2 GB (2x 1 GB)	533 MHz (PC2-4200)

Model	Standard memory	Memory speed
8844-E5X	2 GB (2x 1 GB)	533 MHz (PC2-4200)

Table 3-80 lists the supported memory options.

Table 3-80 JS21 Memory Module Options

Part number	Size	Memory speed
39M5821	1 GB (2x 512 MB)	400 MHz (PC2-3200)
39M5809	2 GB (2x 1 GB)	400 MHz (PC2-3200)
39M5812	4 GB (2x 2 GB)	400 MHz (PC2-3200)
41Y2703	8 GB (2x 4 GB)	400 MHz (PC2-3200)
41Y2707	1 GB (2x 512 MB)	533 MHz (PC2-4200)
41Y2711	2 GB (2x 1 GB)	533 MHz (PC2-4200)
41Y2715	4 GB (2x 2 GB)	533 MHz (PC2-4200)
41Y2723	8 GB (2x 4 GB)	533 MHz (PC2-4200)

3.12.3 Onboard network controllers

The JS21 type 8844 blade server has a dual channel Gigabit Ethernet Broadcom 5780. These two Ethernet interfaces are routed through to the midplane providing a foundation for fault-tolerant network setups.

The Broadcom onboard adapter supports the following features:

- PXE boot support
- Wake on LAN
- Load-balancing
- Fault-tolerance
- VLAN support

3.12.4 Integrated systems management processor

The JS21 type 8844 blade server has an integrated baseboard management controller (BMC) that provides:

- IPMI compliance (Intelligent Platform Management Interface)
- Serial over LAN (SOL)
- Wake on LAN (WOL)
- Power management functions
- KVM and media tray ownership
- Light path diagnostics support
- Automatic BIOS recovery (ABR)
- Automatic server restart (ASR)
- Predictive failure analysis (PFA) on system components (CPUs, memory, and disk drives)
- Inventory
- Error logging
- Environmental monitoring (voltages and temperature on the system board)

Note that this onboard BMC cannot be accessed externally through a dedicated port or be daisy-chained to other servers or blades. The BMC is forwarding information to the

Management Module (MM) within the BladeCenter chassis. The MM is a single point of control for hardware management of all server and switch components contained within the chassis.

3.12.5 Local storage options

The BladeCenter JS21 features an integrated SAS controller with support for up to two fixed 2.5” SFF SAS hard disk drives. The controller can also provide RAID 0 or RAID 1 support.

Table 3-81 shows the SFF HDD options for JS21 type 8844.

Table 3-81 Small form factor hard disk drive options for JS21 type 8844

Part number	Description
26K5776	36.4 GB SAS HDD (2.5") SFF
26K5777	73 GB SAS HDD (2.5") SFF
42D0421	146 GB SAS 10K SFF NHS
43X0845	73 GB SAS 15K SFF NHS

The hard disk drives are mounted onto the system board of the blade server and are not accessible from the outside during operation.

Note: The JS21 type 8844 does not currently support the attachment of any of the BladeCenter storage expansion units.

3.12.6 I/O expansion options

The blade server JS21 type 8844 has two onboard connectors for various I/O expansion card options including network and storage adapters.

One of these connectors is a 133 MHz PCI-X connection, which allows the connection of an SFF, StFF, or CFFv form factor expansion card. The installed expansion card is connected to module bays 3 and 4 at the rear of the BladeCenter chassis via the internal midplane.

The JS21 type 8844 also has a PCI Express connector that allows for the attachment of one of the following:

- A High Speed Form Factor expansion card
- A Combination Form Factor - Horizontal expansion card

The JS21 type 8844 supports five different form factors for expansion cards that can connect to either one or the other of these two connectors. See 3.1, “Expansion cards” on page 234, for further information about form factors and expansion card options.

A StFF I/O expansion card cannot be installed if a hard disk drive is installed in HDD slot 2. Table 3-82 on page 339 summarizes the available I/O expansion card options.

Table 3-82 Blade server JS21 XM type 8844 I/O expansion card options

Part number	Description	Form factor	More information
26K4841	2 Gb Fibre Channel Expansion Card	SFF	248
39R8624	BladeCenter SFF Gigabit Ethernet Expansion Card	SFF	243
73P6000	Myrinet Cluster Expansion Card[a]	StFF	None
32R1896	Cisco Systems InfiniBand 1X HCA Expansion Card	StFF	None
32R1923	QLogic iSCSI Expansion Card for IBM eServer™ BladeCenter	StFF	252
26R0890	QLogic 4 Gb SFF Fibre Channel Expansion Card for BladeCenter	SFF	249
39Y9186	Emulex 4 Gb SFF Fibre Channel Expansion Card	SFF	249
32R1760	Cisco Systems InfiniBand 4X HCA Expansion Card	HSFF	251
39Y9271	NetXen 10 Gb Ethernet Expansion Card (CFFh)	CFFh	243
39Y9306	QLogic Ethernet and 4 Gb Fibre Channel Expansion Card (CFFh)	CFFh	246
39Y9310	Ethernet Expansion Card (CFFv)	CFFv	242
41Y8527	QLogic 4 Gb Fibre Channel Expansion Card (CFFv)	CFFv	248

a. The Myrinet card is a single-channel card and requires you to install an OPM in BladeCenter I/O module bay 4.

PCI Expansion Unit 2

The JS21 type 8844 does not support the attachment of the PCI I/O Expansion Unit 2 (or any previous PCI I/O expansion unit).

3.12.7 Reliability, availability, and serviceability features

Component-level features include:

- Transparent CPU hardware error recovery (for example, for L2 cache errors)
- Memory scrubbing
- Single DIMM failure isolation for CEs; DIMM pair isolation for UEs
- Memory chipkill
- PCI bus parity

Blade-level features include:

- Degraded boot support (memory, CPUs)
- Auto Server Recovery/Restart (automatic reboot after boot hangs or detected checkstop conditions)
- Environmental monitors and alerts
- EPOW support
- System VPD and VPD on all major electronic components
- FRU/CRU-level LEDs
- SMS menu support

The integrated systems management processor combined with light path diagnostics provides true hardware-based monitoring and alerting.

Standard serial console connection is routed to the Ethernet controller for Serial-over-LAN (SoL) as a dedicated port for the primary serial console support.

3.13 JS22 Type 7998-61x

The BladeCenter JS22 is the latest high-density blade server and the first blade to use POWER6 technology with EnergyScale technology. The JS22 is a single-wide, 4-core, 2-socket with two cores per socket, 4.0 GHz POWER6 SCM processor. Each processor is directly mounted to the planar board and includes 64 KB I-cache and 32 KB D-cache L1 cache per core with 4 MB L2 cache per core. The integrated memory controller supports four pluggable registered DIMMs, which must be installed in pairs. The minimum memory that can be installed is 2 GB (2 X 1 GB) and the maximum is 32 GB (2 x 16 GB). The JS22 supports Memory Scrubbing, ECC, Chipkill, and Bit Steering.

As a heterogeneous infrastructure consolidation platform, the BladeCenter H and HT Chassis also support running the JS22 alongside your Intel processor-based HS blades, AMD Opteron processor-based LS blades, and Cell/B.E. processor-based QS21 blades in the same chassis with independent monitoring, security, power, and systems management. And with several gigabit Ethernet, Fibre Channel, iSCSI, and InfiniBand high-performance expansion cards to choose from, the JS22 can be tailored to the demands of your datacenter network and application workloads.

The BladeCenter H Chassis supports a maximum of 14 JS22 blades or a mix of JS22 blades with other blades. For the BladeCenter HT Chassis, the maximum number of blades is ten with five JS22 blades in each power domain.

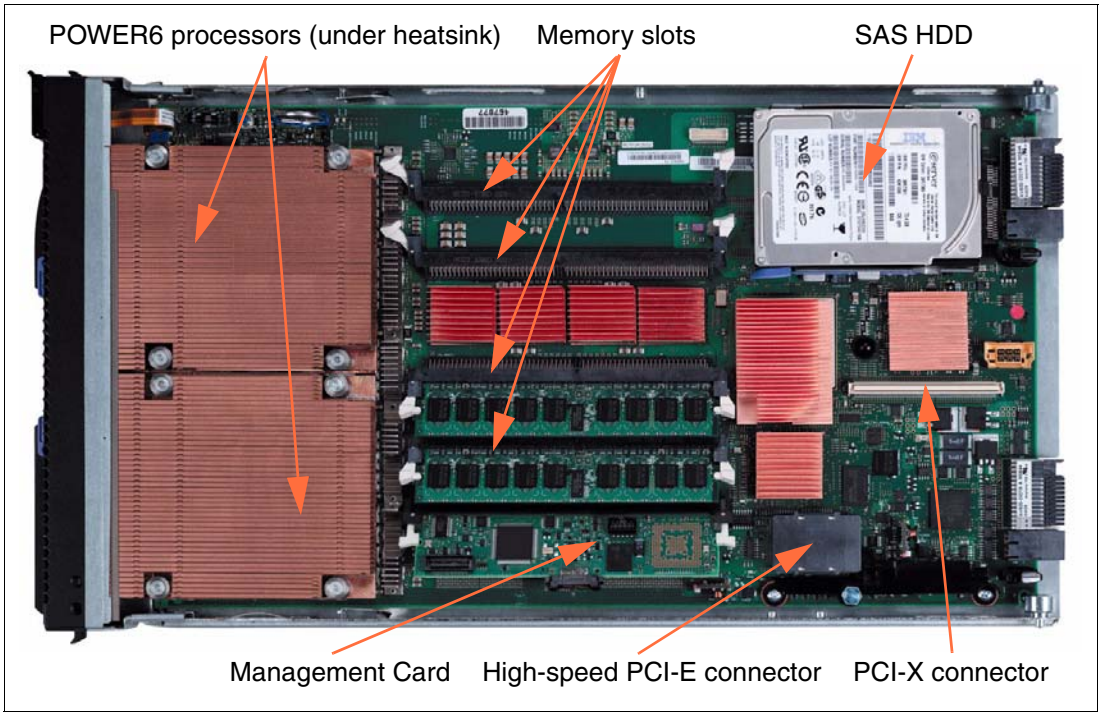

Figure 3-49 JS22 Internals

The JS22 has autonomic power and thermal controls designed to maintain operation in out-of-specification situations. If thermal or power conditions become unacceptable, it might reduce the processor or memory performance capability to maintain acceptable conditions.

When the conditions are restored to an acceptable value, the blade will again automatically resume normal operation. Some applications might be sensitive to these capability changes. Contact your application vendors to determine any possible impacts.

Full keyboard, video, and mouse support can be used at installation instead of Serial over LAN (SOL), which requires a network switch module in the chassis. Multiple console connections are also supported by simultaneously using the BladeCenter keyboard, video, and mouse, a SOL connection, and a direct-serial connection. The USB bus on the BladeCenter chassis is used for keyboard support on the JS22. The Advanced Management Module (aMM) on the chassis controls the routing of keyboard inputs to one of up to 14 blade servers. An embedded video controller is used to support 2D graphics on an SVGA video display. The JS22 has four USB ports; two interface with keyboard and mouse through the aMM and two with a removable media tray also through the aMM.

As a robust blade solution for 64-bit UNIX applications, the JS22 also delivers leadership SIMD (single instruction multiple-data) capabilities with the Linux operating system for scientific research and HPC.

3.13.1 JS22 Models

Table 3-83 lists the models of the BladeCenter JS22.

Table 3-83 JS22 models

7998-61X	Processor	L2 Cache	Memory	Ethernet	HDD
#8400	2-socket dual-core 4 GHz POWER6	4 MB / core	2 GB - 32 GB	Dual Gigabit Ethernet	Open (one bay for SAS drive)

3.13.2 Minimum system configuration

Each 7998-61X configuration must contain a minimum of:

- One - IBM BladeCenter JS22 4-core, 4.0 GHz (feature #8400)
- One - Memory DIMMs (at least one of these):
  - 2 GB (2 x 1 GB) DDR2 667 MHz DIMMs (#8232)
  - 4 GB (2 x 2 GB) DDR2 667 MHz DIMMs (#8233)
  - 8 GB (2 x 4 GB) DDR2 667 MHz DIMMs (#8234)
  - 16 GB (2 x 8 GB) DDR2 533 MHz DIMMs (#8235)
- Processor Entitlement (one of these)
  - Four - Processor Entitlement (#8401)

    or (Express configurations)
  - Two - Processor Entitlement (#8401)
  - Two - Zero-priced Processor Entitlement (#8399)

Note: A minimum of one copy of the Service Warranty Publications (#8259) and one copy of the JS22 Installation and User's Guide (#8260-8263, #8266-8269, or #8278-8281) is required at each customer installation.

3.13.3 Express configurations

If you order a JS22 server Express Product Offering, as defined here, you might qualify for a processor entitlement at no additional charge. A minimum of 4 GB of memory and one SAS 10 K SFF HDD is required to determine if a customer is entitled to a processor entitlement at no additional charge.

You can make changes to the standard features as needed and still qualify for processor entitlements at no additional charge. However, selection of total memory or DASD smaller than the totals defined as the minimums disqualifies the order as an Express Product Offering.

You must purchase PowerVM Standard Edition (#5409) or PowerVM Enterprise Edition (#5649) for each Express Product Offering.

- One - IBM BladeCenter JS22 4-core, 4.0 GHz (#8400)
- 4 GB memory DIMMs with at least 73 GB SAS HDD
- One of these:
  - Two - 2 GB (2 x 1 GB) DDR2 667 MHz DIMMs (#8232)
  - One - 4 GB (2 x 2 GB) DDR2 667 MHz DIMMs (#8233)
  - One - 8 GB (2 x 4 GB) DDR2 667 MHz DIMMs (#8234)
  - One - 16 GB (2 x 8 GB) DDR2 533 MHz DIMMs (#8235)
  - Additional memory DIMMs of any size might be added
- One of these:
  - Four - PowerVM Standard Edition (#5409)
  - Four - PowerVM Enterprise Edition (#5649)
- One of these:
  - One - IBM 73 GB SAS 10 K SFF HDD (#8237)
  - One - IBM 146 GB SAS 10 K SFF HDD (#8236)
- Two - Processor Entitlement (#8401)
- Two - Zero-priced Processor Entitlement (#8399); Receive two processor entitlements (#8399) at no additional charge. Four processors are entitled

Note: A minimum of one copy of the Service Warranty Publications (#8259) and one copy of the JS22 Installation and User's Guide (#8260-8263, #8266-8269, or #8278-8281) is required at each customer installation.

3.13.4 Memory and hard disk drives

- Up to 32 GB of system memory
  - The JS22 system board contains four RDIMM memory connections.
  - Memory must be installed in matched pairs with a minimum of 2 GB and maximum of 32 GB.
  - Memory sizes can be mixed.
  - JS22 supports Memory Scrubbing, ECC, Chipkill, and Bit Steering.
  - 2 GB (2 x 1 GB), 4 GB (2 x 1 GB), and 8 GB (2 x 1 GB) DIMMs operate at 667 MHz while the 16 GB (2 x 1 GB) DIMM operates at 533 MHz.
- SAS Controller, supporting up to 146 GB
  - Internal disk drive, SFF (2.5 in) 73 GB SAS (#8237)
  - Internal disk drive, SFF (2.5 in) 146 GB SAS (#8236)

3.13.5 Advanced Power Virtualization

The JS22 has support for Advanced Power Virtualization (APV) built into the microprocessor. The JS22 is designed to make it more affordable to consolidate multiple independent applications on a single blade using the same proven virtualization technologies offered on IBM System p servers.

The JS22 supports two leading-edge virtualization technologies:

- PowerVM Standard Edition (feature #5409)

  APV Standard Edition makes the JS22 an ideal platform for consolidation of AIX and Linux applications, helping clients reduce infrastructure complexity and cost. Offering an intuitive, Web-based interface for managing virtualization within a single blade, the Integrated Virtualization Manager (IVM) component of Virtual I/O Server (VIOS) allows the small business IT manager to quickly and easily set up and manage logical partitions (LPARs). It also enables Virtual I/O and Virtual Ethernet so that storage and communications adapters can be shared among all the LPARs running on the JS22. Ultimately, IBM Micro-Partitioning technology allows each processor core to be subdivided into as many as 10 virtual servers. Because the JS22 is built with POWER6 technology, other advanced virtualization functions such as Shared Dedicated Capacity can be exploited.

- PowerVM Enterprise Edition (feature #5649)

  APV Enterprise Edition includes all the features of APV Standard Edition plus a new capability called Live Partition Mobility. Live Partition Mobility allows for the movement of a running partition from one POWER6 processor-based server to another with no application downtime, resulting in better system utilization, improved application availability, and potential energy savings. With Live Partition Mobility, planned application downtime due to regular server maintenance can be a thing of the past. Software Maintenance for Virtual I/O Server (577x-VIO) must be purchased with VIOS (5765-G34).

3.13.6 Expansion cards

The form factor of the expansion card determines whether an additional expansion card or SAS internal HDD can be installed in the JS22.

- The QLogic iSCSI Expansion Card for IBM BladeCenter (#8238) is a standard form factor (StFF) expansion card. The installation of this card precludes support for the SAS HDD. The maximum allowed is one.
- One high speed daughter card (HSDC) expansion card can be installed on the JS22 along with one SAS HDD. An example of a supported HSDC is the Cisco Systems 4X InfiniBand HCA Expansion Card (#8256).
- One CFFv expansion card might be installed with a CFFh expansion card and one SAS HDD.

  Supported CFFv expansion cards include:

  - QLogic 4 GB 2 Port Fibre Channel Expansion Card (#8248)
  - Ethernet Expansion Card (CFFv) (#8249)
  - SAS Expansion Card (#8250)
  - Emulex 4 GB Fibre Channel Expansion Card (#8251)

  Supported CFFh expansion cards include:

  - QLogic Ethernet and 4 GB Fibre Channel Expansion Card (#8252)
  - 4X InfiniBand DDR Expansion Card (#8258)
  - Cisco 4X InfiniBand DDR Expansion Card (#8297)

– Voltaire 4X InfiniBand DDR Expansion Card (#8298)

3.13.7 Software requirements

The JS22 server supports the AIX and Linux operating systems, giving clients the flexibility of using applications written for either. IBM has qualified two popular Linux distributions for use with the JS22 server:

- Red Hat Enterprise Linux for POWER
- Novell SUSE Linux Enterprise Server for POWER

One or more of the following operating systems is required for an operational JS22:

- AIX V5.3 Technology Level 5300-06 with Service Pack 4, or later
- AIX V5.3 Technology Level 5300-07, or later
- AIX V6.1 Technology Level 6100-00, or later
- Novell SUSE Linux Enterprise Server 10 for POWER SP1 Update 1, or later
- Red Hat Enterprise Linux for POWER, V4.6, or later
- Red Hat Enterprise Linux for POWER, V5.1, or later

iSCSI boot is available through an iSCSI software initiator running on the embedded Ethernet controller or the iSCSI expansion card (#8238). The AIX operating system supports iSCSI boot through an iSCSI software initiator on any supported Ethernet controller including Ethernet expansion cards.

3.13.8 Integrated systems management processor

The enhanced BMC for the JS22 blade server is a flexible service processor that provides support for the following functions:

- Alert Standard Format (ASF) and RS-485 interfaces to the management modules
- Intelligent Platform Management Interface (IPMI)
- Power control and advanced power management
- Reliability, availability, and serviceability (RAS) features
- Serial over LAN (SOL)
- Continuous health monitoring and control
- Configurable notification and alerts
- Event logs that are time stamped and saved in nonvolatile memory and that can be attached to e-mail alerts
- Point-to-point protocol (PPP) support
- Remote power control
- Remote firmware update and access to critical blade server settings

3.13.9 Systems management and administrative tools

Integrated diagnostic and administrative tools like IBM Predictive Failure Analysis and light path diagnostics are designed to simplify administration to help lower costs and improve control of the IT environment. Remote management capabilities allow automating IT

networking tasks. Proven technology like VIOS allows the sharing of disk drives, communications, and Fibre Channel adapters.

Systems management tools for the JS22:

- AIX 5.3 or AIX 6: IBM Director V5.1 and Cluster Systems Management (CSM) for AIX, V1.7
- Novell SUSE Linux Enterprise Server 10 for POWER SP1 Update 1, or later: IBM Director V5.1 and CSM for Linux on POWER, V1.7
- Red Hat Enterprise Linux for POWER, V4.6, or later: IBM Director V5.1 and CSM for Linux on POWER, V1.7
- Red Hat Enterprise Linux for POWER, V5.1, or later: IBM Director V5.1 and CSM for Linux on POWER, V1.7

3.14 QS21 type 0792

The IBM BladeCenter QS21 type 0792 is based on the innovative Cell Broadband Engine (Cell/B.E.) processor, providing a new level of parallelism and performance to targeted workloads.

The Cell/B.E. processor chip utilizes a multicore architecture optimized for high-performance computing and media-rich applications. The QS21 Cell/B.E. technology-based system is designed for businesses that need the dense computing power and unique capabilities of the Cell/B.E. processor to tackle tasks involving graphic-intensive, numeric applications.

The QS21 is particularly well suited for image processing, video processing, and signal processing applications in Aerospace and Defense, Medical Imaging, EDA, Digital Video Surveillance (DVS), Seismic, Financial markets, Digital Media industries.

The system uses the Cell/B.E. processor to help accelerate key algorithms like 3D rendering, compression, and encryption, to help enable companies to create and run highly visual, immersive, real-time applications.

To read more about the application programming methods and tools that are available for the Cell/B.E. processor, see IBM Redbooks publication *Programming the Cell Broadband Engine: Examples and Best Practices*, SG24-7575, available from:

http://www.redbooks.ibm.com/abstracts/sg247575.html

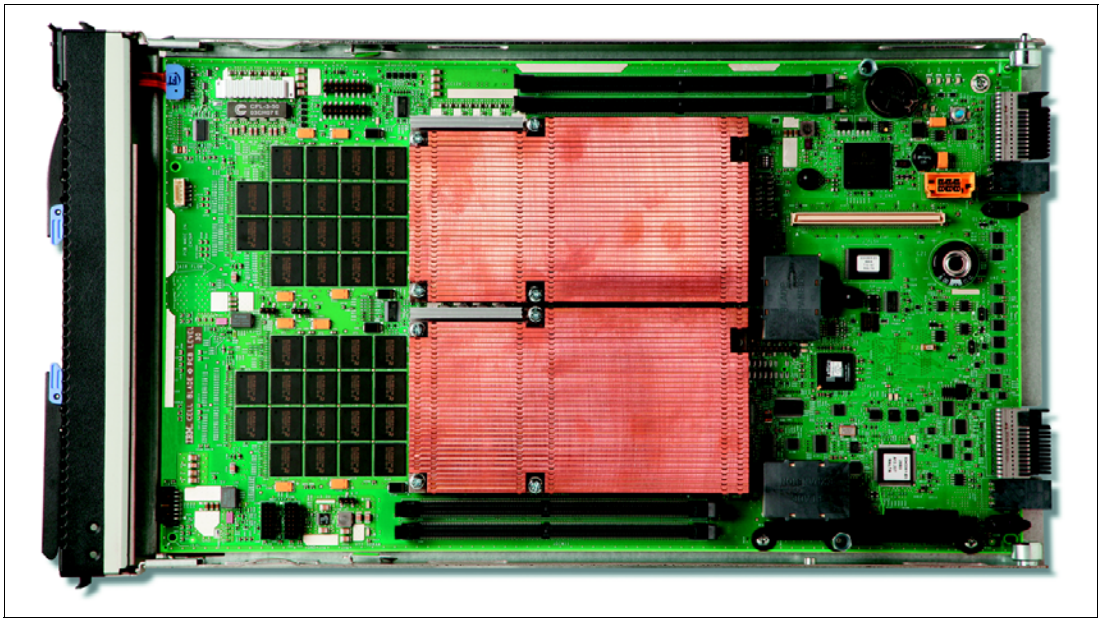

Figure 3-50 The QS21 Blade Server

Key features of the QS21 include the following:

- Two 3.2 GHz Cell/B.E. processors
- 2 GB XDR memory
- Integrated dual GB Ethernet
- IBM Enhanced I/O Bridge Chip
- Optional two or four 512 MB VLP DIMMs as I/O buffer
- Optional InfiniBand adapter
- Serial over LAN
- Advanced power management including over-subscription and thermal throttling
- Single-wide blade for BladeCenter H chassis
- Support by IBM SDK for Multicore Acceleration v3.0

See Table 3-84 for an overview of QS21 type 0792 models.

Table 3-84 Models of the QS21 type 0792 at a glance

Model	CPU Speed/L2 Cache	CPU Qty (Std/Max)	Memory (Std/Max)	Standard SAS disk
0792-32Y	2-way CBE 3.2 GHz/512 KB	2/2	2 GB/2 GB	None

Note: QS21 is currently supported in the BladeCenter S, H, and HT chassis only.

3.14.1 Processor

The BladeCenter QS21 server use Cell/B.E. processors. All QS21 models have two processors installed which are directly mounted to the planar board.

3.14.2 Memory

The memory subsystem consists of 18 XDR memory modules per processor, creating 1 GB of ECC memory per processor. The XDR modules are soldered on the planar for fixed memory configurations.

Additionally, the QS21 supports two or four optional VLP DDR2 memory DIMMs to act as an I/O buffer. There are two I/O buffers, one for each Cell/B.E. companion chip. Each buffer is implemented as two DIMM sockets. Only 512 MB DIMMs are supported and they must be installed so that the two buffers have the same amount of memory. Table 3-85 lists the supported memory options. If you only have one memory option, install one DIMM in one buffer and one DIMM in the other.

Note: The DIMMs are used as memory for the I/O buffers only. You cannot increase the size of system memory, which is fixed at 1 GB for each Cell/B.E processor.

Table 3-85 QS21 Memory Module options

Part number	Size	Memory speed
39M5861	1 GB (2x 512 MB)	667 MHz (PC2-5300) DDR2

3.14.3 Onboard network controllers

The QS21 blade server has a dual channel Gigabit Ethernet Broadcom 5704S chip. These two Ethernet interfaces are routed to the bays 1 and 2 through midplane providing a foundation for fault-tolerant network setups.

3.14.4 Integrated systems management processor

The QS21 type 0792 blade server has an integrated service processor that provides:

- BMC, IPMI compliant code stack
- Serial over LAN (SOL)
- Wake on LAN (WOL)
- Power management functions
- Light path diagnostics support (Lightbox LEDs)
- Automatic server restart (ASR)
- Predictive Failure Analysis for CEs on the Cell BE companion chip attached I/O Buffer DDR2 DIMMs
- Error logging
- Environmental monitoring (voltages and temperature on the system board)

Note that this onboard BMC cannot be accessed externally through a dedicated port or be daisy-chained to other servers or blades. The BMC is forwarding information to the Management Module (MM) within the BladeCenter chassis. The MM is a single point of control for hardware management of all server and switch components contained within the chassis.

3.14.5 Local storage options

Currently, BladeCenter QS21 have no integrated SAS storage controller, and other techniques should be used for OS boot. For example, network boot (most commonly used) or boot from local removable media are supported.

The QS21 type 0792 supports connecting to the IBM BladeCenter Boot Disk System using the SAS Expansion Card.

3.14.6 I/O expansion options

The blade server QS21 type 0792 has two onboard connectors. One of these connectors is a 133 MHz PCI-X connection, and other connector is a PCI Express connector.

Currently the QS21 type 0792 only supports the expansion options as listed in Table 3-86.

Table 3-86 Blade server QS21 type 0792 I/O expansion card options

Part number	FC	Description	Form factor	Page
39Y9190	2979	SAS Expansion for IBM BladeCenter	CFFv	252
32R1760	1466	Cisco Systems InfiniBand 4X HCA Expansion Card	HSFF	251
43W4423	2991	4X InfiniBand DDR Expansion Card	CFFh	251

3.14.7 Reliability, availability, and serviceability features

Component-level features include:

- Transparent CPU Hardware error recovery
- ECC for XDR and DDR2 memory
- PFA for CEs on the Cell BE companion chip attached I/O Buffer DDR2 DIMMs
- Memory Scrubbing on XDR system memory
- XDR System memory failure isolation for memory errors to a single memory interface (Cell BE-0 or Cell BE-1 attached)
- DDR2 I/O Buffer memory failure isolation to a single DIMM
- PCI Bus Parity

Blade level features include:

- Degraded boot for both XDR and DDR2 memory errors. XDR memory errors might create holes in the memory map.
- Automatic server recovery and restart:
  - Automatic reboot after boot hangs (with switch of the boot flash to the PERM side)
  - Automatic reboot after checkstop (without switch of the boot flash side)
- Environmental monitors and alerts
- System VPD and VPD on all major electronic components
- System Management Services (SMS) menu support
- Checkstop detection with data logging and automated reboot
- Light path diagnostics provides true hardware-based monitoring and alerting.

3.15 QS22 type 0793

The IBM BladeCenter QS22 type 0793 is based on the innovative multi-core IBM PowerXCell 8i processor, a new generation processor based on the Cell Broadband Engine (Cell/B.E.)Architecture. Offering enhanced double-precision floating-point processing power,

the QS22 can yield application results faster and with more fidelity. This can enable organizations to get information faster to facilitate important business decisions.

The Cell/B.E. processor chip utilizes a multicore architecture optimized for high-performance computing and media-rich applications. The new QS22 Cell/B.E. technology-based system is designed for businesses that need the dense computing power and unique capabilities of the Cell/B.E. processor to tackle tasks involving graphic-intensive, numeric applications.

The QS22 Blade Server, shown in Figure 3-51, is particularly well suited for image processing, signal processing, and graphics rendering applications in aerospace/defense, financial markets, medical imaging, EDA, digital video solutions, petroleum exploration and other industries.

The system uses the Cell/B.E. processor to help accelerate key algorithms like 3D rendering, compression, and encryption, to help enable companies to create and run highly visual, immersive, real-time applications.

To read more about the application programming methods and tools that are available for the Cell/B.E. processor, see IBM Redbooks publication *Programming the Cell Broadband Engine: Examples and Best Practices*, SG24-7575, available from:

http://www.redbooks.ibm.com/abstracts/sg247575.html

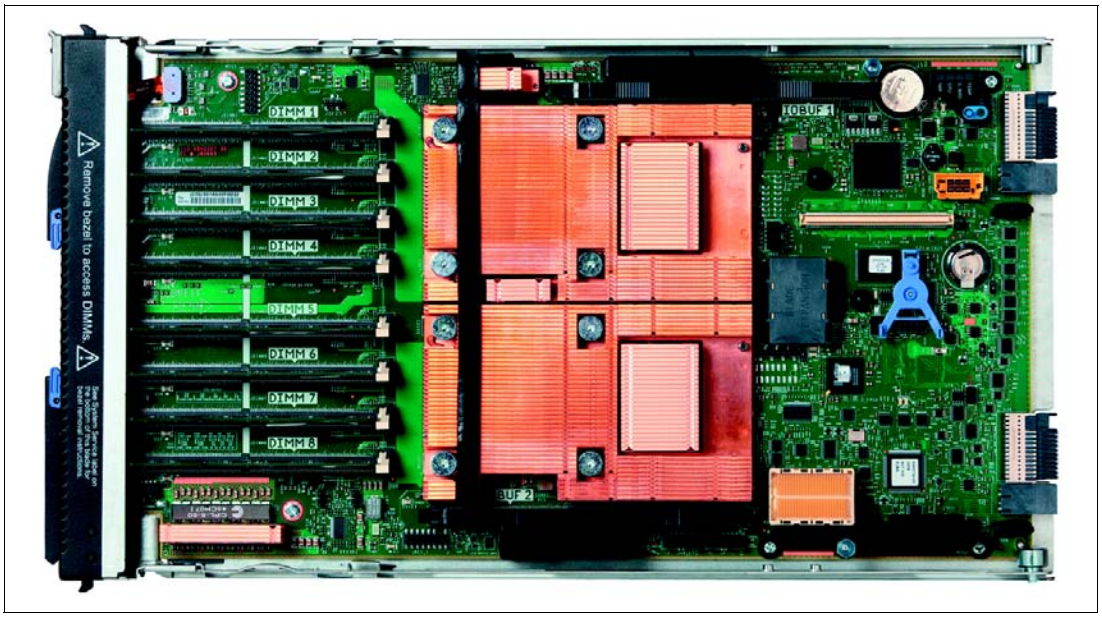

Figure 3-51 The QS22 Blade Server

Key features of the QS22 include the following:

- Two 3.2 GHz IBM PowerXCell 8i Processors
- Up to 32 GB PC2-6400 800 MHz DDR2 of memory (16 GB per processor)
- 460 single precision (SP) GFLOPS per blade (peak)
- 217 double precision (DP) GFLOPS per blade (peak)
- Integrated dual 1 Gb Ethernet BCM5704S
- IBM Enhanced I/O Bridge Chip
- Support for IBM System Storage Boot Disk System (172622B)
- Two DIMM slots for optional two 1 GB VLP DIMMs as I/O buffer
- Serial over LAN
- Advanced power management including over-subscription and thermal throttling
- Single-wide blade

- Support by IBM SDK for Multicore Acceleration v3.0 and v3.1 (v3.1 is recommended version)

Table 3-87 lists QS22 type 0793 models.

Table 3-87 Models of the QS22 type 0793

Model	CPU Speed/L2 Cache	CPU Qty (Std/Max)	Memory (Std/Max)	Standard SAS disk
0793-38Y	Two 3.2 GHz/512 KB	2/2	8 GB/32 GB	None
0793-40Y	Two 3.2 GHz/512 KB	2/2	16 GB/32 GB	None
0793-41Y	Two 3.2 GHz/512 KB	2/2	16 GB/32 GB	None
0793-42Y	Two 3.2 GHz/512 KB	2/2	32 GB/32 GB	None

Note: QS22 is currently supported in the BladeCenter S (with SAS Connectivity Module only), BladeCenter H, and BladeCenter HT chassis only.

For more information about QS22-based solutions, refer to the Cell Broadband Engine Resource Center on IBM DeveloperWorks:

http://www.ibm.com/developerworks/power/cell/

3.15.1 Processor

The BladeCenter QS22 is based on the 64-bit IBM PowerXCell 8i processor operating at 3.2 GHz. Two processors per blade are directly mounted to the blade planar board to provide multiprocessing capability. Each processor includes 32/32 KB L1 (data/instruction) and 512 KB L2 cache.

3.15.2 Memory

The memory subsystem consists of 8 DIMM slots, enabling configurations from 4 GB (if 4 slots are populated with 1 GB DIMMs) to 32 GB (if all 8 slots are populated with 4 GB DIMMs) of ECC memory. The QS22 is available in four base model configurations:

- QS22 type 0793-38Y with 8 GB Memory (8x 1GB DIMMs)
- QS22 type 0793-40Y with 16 GB Memory (4x 4GB DIMMs)
- QS22 type 0793-41Y with 16 GB Memory (8x 2GB DIMMs)
- QS22 type 0793-42Y with 32 GB Memory (8x 4GB DIMMs)

IBM BladeCenter QS22 also supports an option to add special purpose I/O buffer memory. If desired, two 1 GB VLP DIMMs can be added. This optional memory is sold in kits of two VLP DIMMs, and can be purchased from Options by IBM under Part Number 46C0501.

The function of I/O buffer memory is unlike typical processor memory in that applications must be written or modified to utilize these memory DIMMs, which are attached to the IBM I/O bridge chips on the blade. Table 3-88 on page 351 lists the QS22 Memory Module options.

Table 3-88 QS22 Memory Module options

Part number	Size	Memory speed
46C0501	2 GB (2x1 GB) VLP	PC2-6400 DDR2
46C0509	4 GB (2x2 GB) VLP	PC2-6400 DDR2
46C0510	8 GB (2x4 GB) VLP	PC2-6400 DDR2

Note: The 2 GB DIMM option (part number 46C0501) can be used as memory for the I/O buffers, as well as for system memory.

3.15.3 Onboard network controllers

The QS22 blade server has a dual channel Gigabit Ethernet Broadcom 5704S chip. These two Ethernet interfaces are routed to bays 1 and 2 through the midplane, providing a foundation for fault-tolerant network setups.

3.15.4 Integrated systems management processor

The QS22 type 0793 blade server has an integrated service processor that provides:

- Monitoring of system and battery voltage, system temperature, fans, power supplies, processor and DIMM status
- Fan speed control
- Product ID and Family ID detection
- Highly secure remote power on/off
- System reset control
- NMI/SMI detection and generation
- System diagnostic LED control (power, HDD, activity, alerts, heartbeat)
- IPMI over LAN
- Serial Over LAN
- Proxy server support
- LAN messaging and alerting
- Text console redirection over LAN
- VLAN support
- Enhanced authentication and encryption algorithms (RMCP+, SHA-1, AES)
- Local update of BMC firmware
- Firmware firewall
- Support for IPMI v1.5 compliant management software (for example, xCAT)
- Other mandatory and optional IPMI BMC functions

Note that this onboard BMC cannot be accessed externally through a dedicated port or be daisy-chained to other servers or blades. The BMC forwards information to the Management Module (MM) within the BladeCenter chassis. The MM is a single point of control for hardware management of all server and switch components contained within the chassis.

3.15.5 Local storage options

The QS22 supports either IBM 8GB Modular Flash Drive (part number 43W3934) or IBM 8 GB Modular Solid State Disk (part number 60H4324) for local storage and to boot supported operating systems. These drive options plug into a socket on the QS22 blade. Table 3-90 lists the QS22 type 0793 internal storage options.

Table 3-89 Blade server QS22 type 0793 storage options

Part number	Description
43W3934	IBM 8 GB Modular Flash Drive
60H4324	IBM 8 GB Modular Solid State Disk

3.15.6 Boot options

The BladeCenter QS22 has several scenarios to boot the server to an operating system.

One solution is to implement a boot server. This option involves installing the RHEL5.2 operating system to the local disk of JS21, JS22, or other POWER-compatible system. After the installation is complete the image of the installation is copied to the remote boot server, which is then used to boot the QS22 blades over the network.

An alternative boot solution for QS22 is to use the BladeCenter Boot Disk System, a 2U, 12-drive enclosure that supports hot-swap SAS drives. This storage product can be used to consolidate the boot functionality for up to 28 QS22 blades to a single, highly available enclosure. To use the boot disk system, an optional SAS expansion adapter card (39Y9190) is required for each QS22 blade. The QS22 blades with optional SAS adapters connect to the boot disk system through the SAS Connectivity Module for BladeCenter (39Y9195). SAS adapters and SAS connectivity modules are purchased separately in addition to the BladeCenter Boot Disk System.

A third boot solution is to use the optional IBM 8GB Modular Flash Drive (part number 43W3934) or IBM 8 GB Modular Solid State Disk (part number 60H4324). These drive options plug into a socket on the QS22 blade. When using these options, the Linux operating system can be installed and QS22 can boot directly from the drive. Table 3-90 lists the QS22 type 0793 storage options.

Table 3-90 Blade server QS22 type 0793 boot options

Part number	Description
43W3934	IBM 8 GB Modular Flash Drive
60H4324	IBM 8 GB Modular Solid State Disk
172622B	IBM BladeCenter Boot Disk System

For information about installing Linux on a BladeCenter QS22, see the following:

http://publib.boulder.ibm.com/infocenter/systems/index.jsp?topic=/liaaj/linuxonbladecenterqs22.htm

3.15.7 I/O expansion options

The blade server QS22 type 0793 has two onboard connectors. One of these connectors is a 133 MHz PCI-X connection, and other is a PCI Express connector.

Currently, the QS22 type 0793 only supports the expansion options listed in Table 3-91.

Table 3-91 Blade server QS22 type 0793 I/O expansion options

Part number	Description	Form factor	Page
39Y9190	SAS Expansion for IBM BladeCenter	CFFv	252
43W4423	4X InfiniBand DDR Expansion Card	CFFh	251

PCI Express I/O Expansion Unit (PEU3e)

The QS22 supports the attachment of one PCI Express I/O Expansion Unit, part number 43W4391. The PEU3e allows standard PCI Express-type adapters to be attached to the server.

For more information about the PEU3e, see 3.18.5, “PCI Express I/O Expansion Unit (PEU3e)” on page 368

The QS22 does not support the attachment of the PCI I/O Expansion Unit 2 (or any previous PCI I/O expansion unit).

3.16 PN41 type 3020

The IBM BladeCenter PN41 blade server (3020-1RX) performs Deep Packet Inspection (DPI) and analysis of network traffic, which enables new classes of applications. DPI looks beyond addresses and ports into packets to help deliver visibility, service control, and security of network traffic. Additionally, PN41 offers development environment using Eclipse IDE, which enables rapid application development.

The BladeCenter PN41 is a highly programmable high performance blade that helps to improve detection and management of network traffic, to enable advanced visibility, service management, and security of network traffic, and to provide improved protection of mission-critical network infrastructure, which are key concerns in today's telecommunications and national government environments.

This blade combines the powerful capabilities of BladeCenter with the performance of CloudShield, an industry-leading network service control and security provider. This combination creates an optimal platform for building and operating next-generation network service management solutions.

Note: IBM provides only the IBM BladeCenter PN41 hardware. Customers will receive a voucher for the licensed internal code, which they will receive from CloudShield under the CloudShield end-user license.

PN41 is only supported within BladeCenter H (8852) and BladeCenter HT (8740/8750) chassis. Additionally, each DPI blade requires its own Application Server Module (ASM). The ASM runs as a virtual machine in a VMware ESX 3.0.1 or later environment on any compatible blade server in the same BladeCenter.

IBM BladeCenter PN41 blade is shown on Figure 3-52 on page 354.

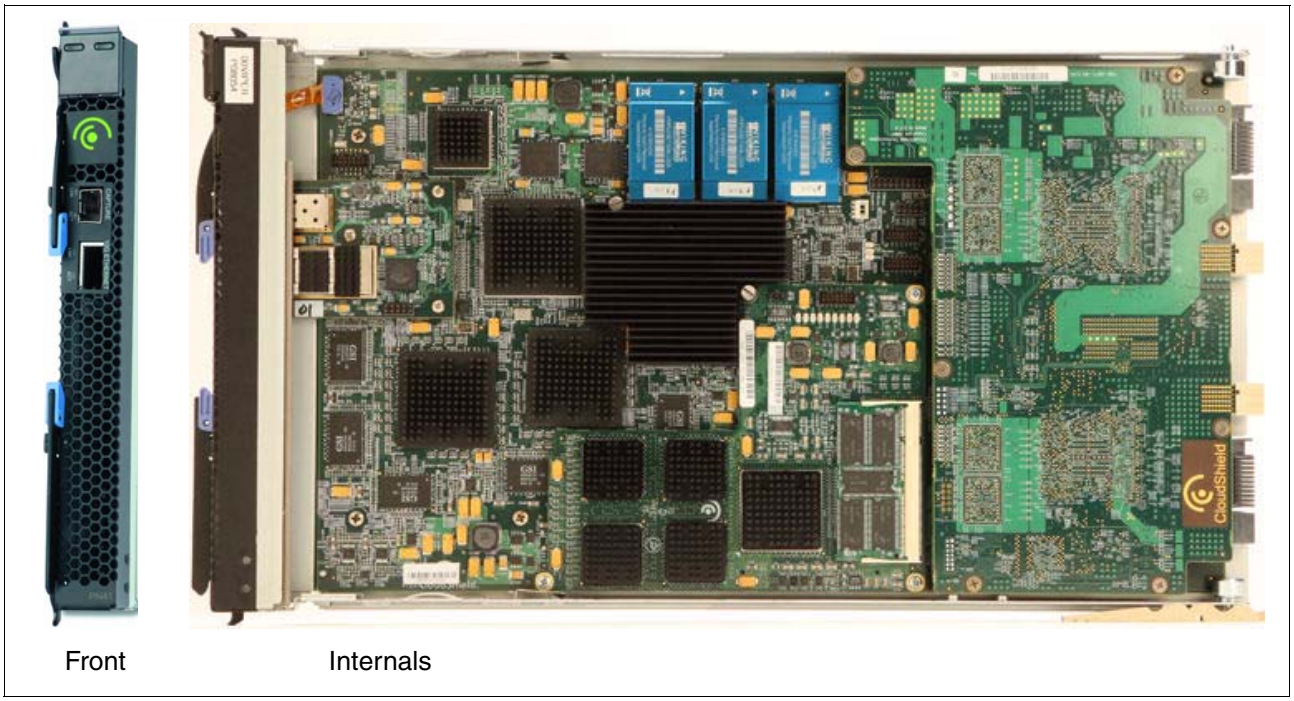

Figure 3-52 The PN41 DPI blade

Key features of the PN41 include the following:

- Contains Intel IXP2805 network processor for use in handling packets.
- Contains a Regular Expression® engine (REGEX engine) to speed up searches.
- Contains multiple types of memory such as TCAM, QDR, and RDRAM. Memory is dedicated to packet processing and is not expandable.
- Contains no drives or flash drives.
- Four internal 1 Gb Ethernet ports used for management interface.
- Four internal 10 Gb Ethernet ports used for data flow analysis.
- One external XFP-based 10 Gb Ethernet port (on the blade's front panel) used for data flow analysis
- One external SFP-based 1 Gb Ethernet port used for data capture.
- Baseboard Management Controller with IPMI support
- RS-485 interface for communication with the management module
- Full Layer 7 processing and control
- Full payload screening (DPI)
- Up to 20 Gbps of DPI throughput per blade
- Selective traffic capture. rewrite, and redirect
- Single-wide blade

For more information about PN41-based solutions, refer to the following publications:

- *IBM BladeCenter PN41 DPI Blade Getting Started Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5077816

- *IBM BladeCenter PN41 DPI Blade Installation and User Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5077817

3.16.1 Processor

The BladeCenter PN41 contains Intel IXP2805 network processor for use in handling packets. This is a dedicated processor and is not available for general applications.

3.16.2 Memory

The BladeCenter PN41 contains multiple types of memory such as TCAM, QDR, and RDRAM. Memory is dedicated to packet processing and is not expandable.

3.16.3 Onboard network controllers

The PN41 blade server has a total of ten network ports:

- Four internal 1 Gb Ethernet ports used for management interface
- Four internal 10 Gb Ethernet ports used for data flow analysis
- One external XFP-based 10 Gb Ethernet port (on the blade's front panel) used for data flow analysis
- One external SFP-based 1 Gb Ethernet port (on the blade's front panel) used for data capture

XFP and SFP modules are not included and must be purchased separately. The following XFPs can be used with PN41:

- IBM option part number 32R1877 10GbE 850 nm Fiber XFP Transceiver (SR) or equivalent
- IBM option part number 32R1878 10GbE 1310 nm Fiber XFP Transceiver (LR) or equivalent

Note: If SFPs are required, they may be purchased from third-party suppliers.

3.16.4 Integrated systems management processor

The PN41 type 3020 blade server has an integrated service processor that provides:

- IPMI support
- Power control and advanced power management
- Automatic server restart
- Built-in monitoring for temperature and voltage
- System error logging

Note that this onboard BMC cannot be accessed externally through a dedicated port or be daisy-chained to other servers or blades. The BMC forwards information to the Management Module (MM) within the BladeCenter chassis. The MM is a single point of control for hardware management of all server and switch components contained within the chassis.

3.16.5 Local storage options

There are no local storage options available for BladeCenter PN41.

3.16.6 I/O expansion options

There are no I/O expansion options available for BladeCenter PN41.

3.17 Concurrent KVM Feature Card (cKVM)

The Concurrent KVM Feature Card, part number 26K5939, can be installed in blade servers that support it (see Table 1-3 on page 5). It provides the ability for multiple users to remotely manage blade servers simultaneously through the management module.

Without the cKVM, for multiple users to access blades within the chassis, administrators have to use serial over LAN, Director Remote Control, or tools like RDC, VNC, or PC Anywhere, all of which require appropriate software on the blade server. The cKVM card provides this function via hardware and does not require additional software on the blades. You can also use the remote control functions even during boot or during an operating system installation.

The cKVM card installs under the I/O daughter cards and does not restrict the physical installation of other expansion cards. The cKVM card is shown in Figure 3-53.

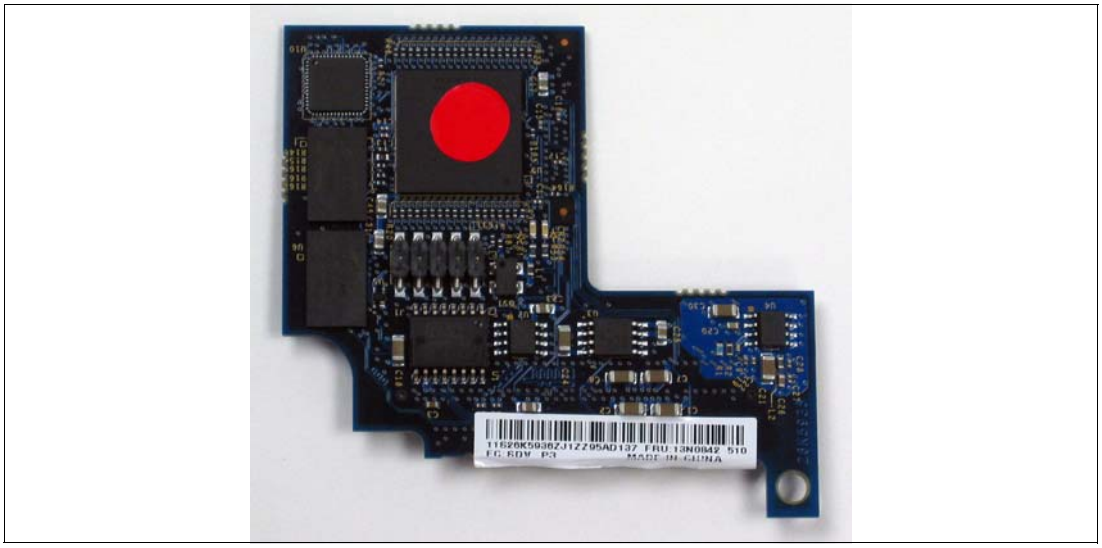

Figure 3-53 Concurrent KVM Feature Card

The cKVM card is installed in a dedicated socket, as shown in Figure 3-54.

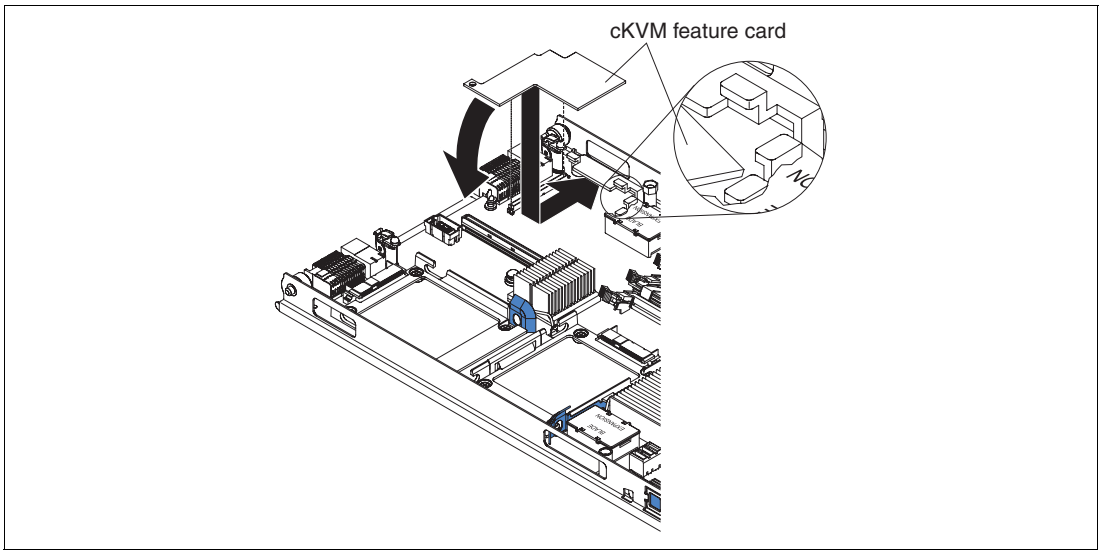

Figure 3-54 Installing the cKVM card in a blade server

The cKVM provides two functions:

- Concurrent keyboard/video/mouse (KVM): For blades in the chassis with the cKVM card installed, this feature allows up to 20 browser-based users to each control the console of one or more blades simultaneously.

  Up to four remote control sessions per blade can be active using cKVM, although the maximum number of remote control sessions per blade can be limited to one if required.

  The video connection is digital and provides a performance improvement over the analog video connectivity through the management module without the cKVM card.

- Concurrent remote media (CD-ROM, diskette drive): For blades in the chassis with the cKVM card installed, this feature allows multiple remote users to simultaneously connect a CD-ROM or diskette drive (inserted in their mobile computer, for example) and have it connected to a blade server. Every blade with cKVM installed can have its own mounted remote disk simultaneously (or concurrently).

  This concurrent remote media feature is independent of the media tray in the chassis itself. The local can belong to only one blade at time.

  Note: Concurrent remote media refers to multiple users connecting remote media (CD-ROM, diskette drive) to multiple blade servers, all concurrently. It does *not* refer to simultaneous access to the media tray in the BladeCenter chassis.

  There had been some restrictions associated with installing operating systems to blades via remote media mounted through a cKVM card. This is described in the following:

  `http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5071326`
  `http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5071121`

  Most of the restrictions have reportedly been resolved with the latest release of the cKVM card firmware.

The configuration rules are listed here:

- Not all blades in the chassis need to have the cKVM card installed. However, those that do not have the card will be remotely controllable (KVM and remote disk) only one at time.

- The Management Module does not support concurrent KVM or concurrent media, only the Advanced Management Module. The cKVM card works in any chassis that has an Advanced Management Module installed.
- The maximum number of supported concurrent KVM sessions per chassis is 20.
- The maximum number of concurrent KVM sessions to the one blade is 4.
- At least one Ethernet switch module installed in bay 1 or 2 is required to use cKVM functions. Pass-thru modules cannot be used. See RETAIN tip H19786:

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5070957
- Not all blade servers support the cKVM card. See Table 1-3 on page 5 for compatibility.

3.18 Expansion blades

There are four expansion blades currently available that add functions to existing blade servers. These expansion blades attach to the server blade and double the width of the overall server. In some configurations, two expansion blades can be attached to a single blade server.

Table 3-92 lists the available expansion blades and the servers that support them.

Table 3-92 Currently available expansion blades and supported server blades

Expansion unit	HS12 — 8028	HS12 — 8014	HS20 — 8843	HS21 — 8853	HS21 XM — 7995	LS21 — 7971	LS22 — 7901	LS41— 7972	LS42 — 7902	HC10 — 7996	JS12 — 7998-60x	JS20 — 8842	JS21 — 8844	JS22 — 7998-61x	QS21 — 0792	QS22 — 0793	PN41 — 3020	Page
Memory and I/O Expansion Unit (MIO), 42C1600	N	N	N	Y	N	N	N	N	N	N	N	N	N	N	N	N	N	361
Storage and I/O (SIO) Expansion Unit, 39R7563	Y	Y	N	Y	Y	Y	N	Y	N	N	N	N	N	N	N	N	N	363
PCI Expansion Unit II (PEU2), 25K8373	N	N	Y	Y	Y	N[a]	Y	N[a]	Y	N	N	N	N	N	N	N	N	367
PCI Express I/O Expansion Unit (PEU3e), 43W4391	Y	N	N	Y	Y	Y	Y	Y	Y	N	N	N	N	N	Y	Y	N	368

a. The PEU2 is only supported on the LS21 and LS41 with custom testing using the SPORE process.

As you can see in Table 3-92, some blades such as the HS21 type 8853 and HS21 XM type 7995 support more than one expansion blade. There are configuration requirements related to connecting multiple expansion blades:

- For HS21 type 8853, you can only attach two expansion blades to the server blade. Attaching three is not supported.
- The following configurations are supported. Note that the sequence is important (for example, the MIO must be directly attached to the HS21).
  - HS21 - MIO - PEU2
  - HS21 - MIO - PEU3e
  - HS21 - MIO - SIO[1]

[1] In this configuration, the PCI-X connectors on the SIO will not function.

- HS21 - SIO - PEU2
- HS21 - SIO - PEU3e[2]

The following are *not* supported combinations:

- HS21 - PEU2 - MIO (not supported)
- HS21 - PEU3e - MIO (not supported)
- HS21 - SIO - MIO (not supported)
- HS21 - PEU2 - SIO (not supported)
- HS21 - PEU3e - SIO (not supported)

Although the HS21 supports all three expansion blades, only two can be attached to a server at one time. If used, the MIO must be *closest* to the blade server.

- For HS21 XM type 7995, the SIO and PEU2/PEU3e are supported in the following connection sequence:
  - HS21 XM - SIO - PEU2
  - HS21 XM - SIO - PEU3e

  The following are *not* supported combinations:
  - HS21 XM - PEU2 - SIO is *not* supported.
  - HS21 XM - PEU3e - SIO is *not* supported.

  Note: When the HS21 XM - SIO - PEU3e sequence is used, only the first PCIe slot in the PEU3e will function.

- For LS21, the SIO and PEU3e are supported in the following connection sequence:
  - LS21 - SIO - PEU3e

  The following is an *unsupported* combination:
  - LS21 - PEU3e - SIO is *not* supported.
- For LS41, the SIO and PEU3e are supported in the following connection sequence:
  - LS41 - MPE - SIO - PEU3e

  The following is an *unsupported* combination:
  - LS41 - MPE - PEU3e - SIO is *not* supported.
- The HS20 8843 supports both the PEU2 (listed in Table 3-92 on page 358) and the BSE2 (listed in Table 3-93 on page 361), but these two expansion units are *not* supported together.
- You can only attach one blade of each type (for example, only one SIO expansion blade).

[2] In this configuration, only the first PCIe slot in the PEU3e will function.

Figure 3-55 shows example configurations.

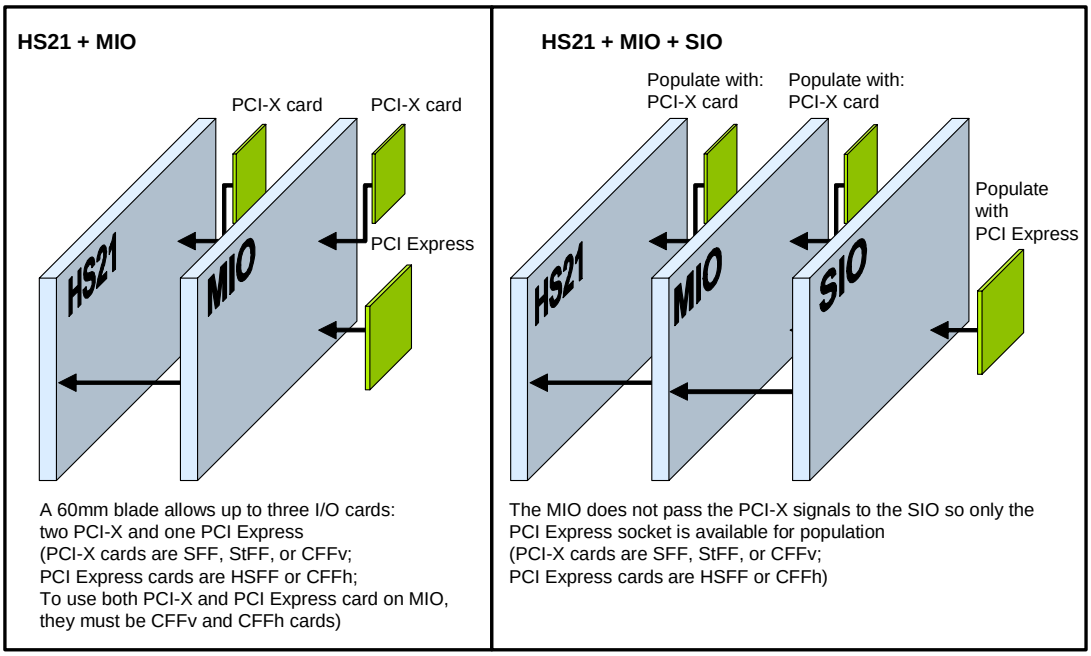

Figure 3-55 Examples of supported configurations

Figure 3-56 illustrates other supported configurations.

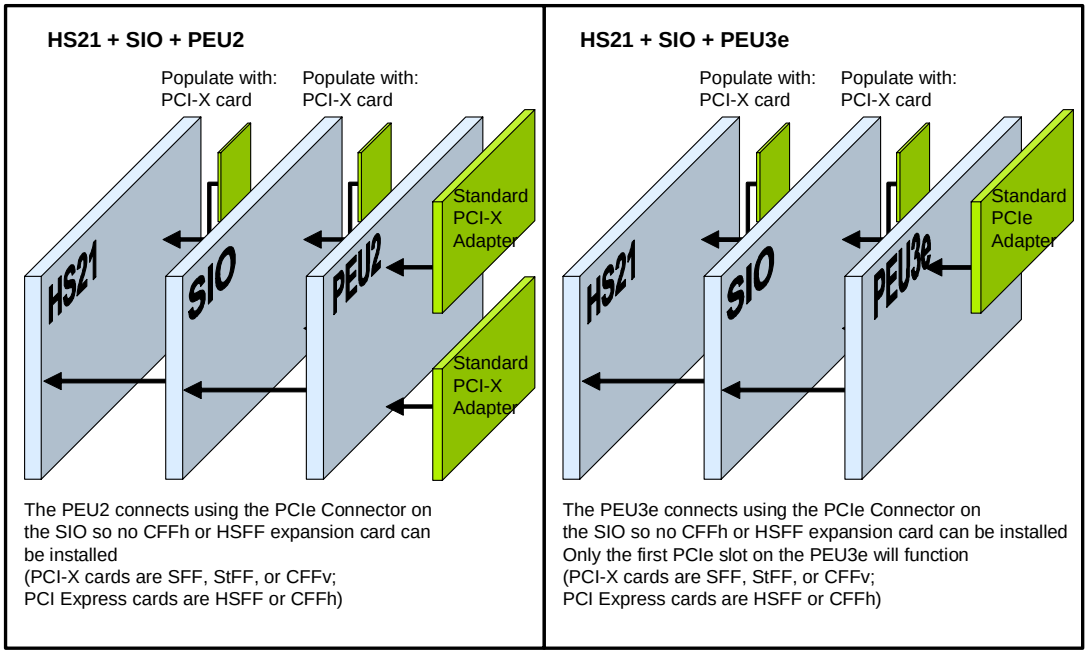

Figure 3-56 More examples of supported configurations

Table 3-93 lists expansion units that are now withdrawn from marketing.

Table 3-93 Withdrawn expansion blades and supported server blades

Expansion unit	HS12 — 8028	HS12 — 8014	HS20 — 8843	HS21 — 8853	HS21 XM — 7995	LS21 — 7971	LS22 — 7901	LS41— 7972	LS42 — 7902	JS12 — 7998-60x	JS20 — 8842	JS21 — 8844	JS22 — 7998-61x	QS21 — 0792	QS22 — 0793	Page
BladeCenter SCSI Storage Expansion Unit II, 39R8625	N	N	Y	N	N	N	N	N	N	N	N	N	N	N	N	366
BladeCenter SCSI Storage Expansion Unit, 26K4817	N	N	Y	N	N	N	N	N	N	N	N	N	N	N	N	None
PCI Expansion Unit, 90P3721	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	None

3.18.1 BladeCenter Memory and I/O Expansion Unit

The Memory and I/O Expansion Unit (MIO), part number 42C1600, allows users to expand the number of DIMMs, I/O slots, and NICs on a blade. It is perfect for customers looking to do more with their blade than traditionally possible. The MIO includes two non-TOE enabled Broadcom NICs, and space for either a StFF, SFF, or CFFv I/O expansion card (in addition to the one on the base server) and/or a High-Speed I/O expansion card (HSFF or CFFh).

Note: The MIO is currently only supported on the HS21 type 8853 blade server.

This expansion unit attaches to the blade server and takes up an extra blade slot. If the MIO is used in conjunction with other expansion units, the MIO needs to installed closest to the blade server. See 3.18, “Expansion blades” on page 358 for details about supported combinations

Figure 3-57 Memory and I/O (MIO) Expansion Unit attached to the HS21

Figure 3-58 shows the system board of the MIO expansion unit.

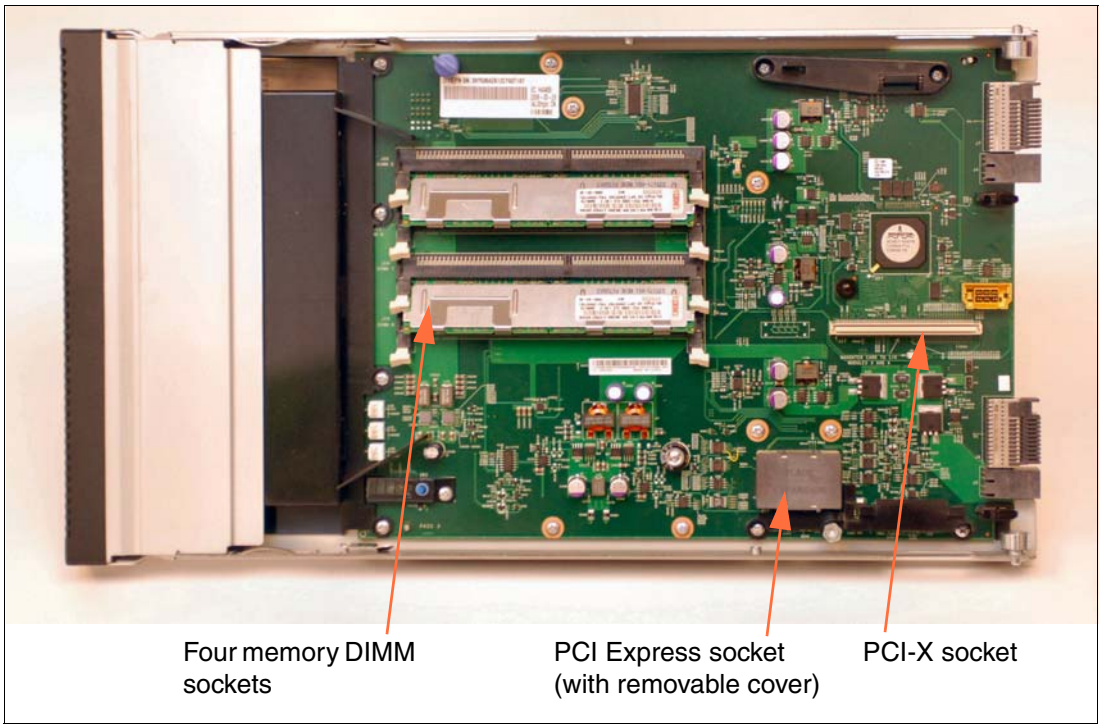

Figure 3-58 Memory and I/O Expansion Unit

The MIO has the following features:

- Four DIMM slots
- Two Gigabit Ethernet ports via a dual channel Broadcom (BCM 5714S) Gigabit Ethernet controller built into the motherboard (non-TOE enabled)
- One PCI-X slot for a StFF, SFF, or CFFv expansion card
- One PCI Express connector for a HSFF or CFFh expansion card

Note: The PCI Express connector is only usable when the server and MIO are installed in a BladeCenter H, BladeCenter HT, or BladeCenter S chassis.

Connection rules are as follows:

- You can install any StFF, SFF, or CFFv expansion card in the PCI-X connector, but it limits what cards can be used in the PCI Express connector.
- You can install an HSFF expansion card in the PCI Express connector, but if you do, you cannot install any card in the PCI-X connector (as the cards will occupy the same physical space).
- You can install a CFFh expansion card in the PCI Express connector, by if you do, you can only install a CFFv expansion card in the PCI-X connector.
- If you also plan to also attach a PEU2 or SIO expansion unit to the MIO, then it connects using the PCI Express connector, so you cannot also install a CFFh or HSFF expansion card in the MIO.
- If you attach a SIO expansion unit to the MIO, then the two PCI-X slots on the SIO will not function. Only the PCIe connector on the SIO will be usable.

The onboard Ethernet controllers and the expansion cards installed must match the switch modules installed in the chassis switch module bays.

- The two onboard Ethernet controllers require an Ethernet switch module or Pass Thru Module in bays 1 and 2
- The PCI-X connector requires a switch module in bays 3 and 4
- The PCI Express connector requires a switch module in bays 7–10 (BladeCenter H or HT only)

3.18.2 Storage and I/O Expansion Unit

The BladeCenter Storage and I/O Expansion Unit (SIO), part number 39R7563, adds additional storage and networking capabilities to a supported blade server. Table 3-92 on page 358 shows the blade servers that support the attachment of the SIO.

The SIO has the following features:

- Three hot swap 2.5” SAS HDD bays
- Dedicated SAS RAID controller
  - Standard configuration includes the ServeRAID 8k-l
  - Optional upgrade to the ServeRAID 8k
- Two PCI-X I/O expansion connectors for StFF, SFF or CFFv expansion cards.
- One PCI Express connector for an HSFF or CFFh expansion card

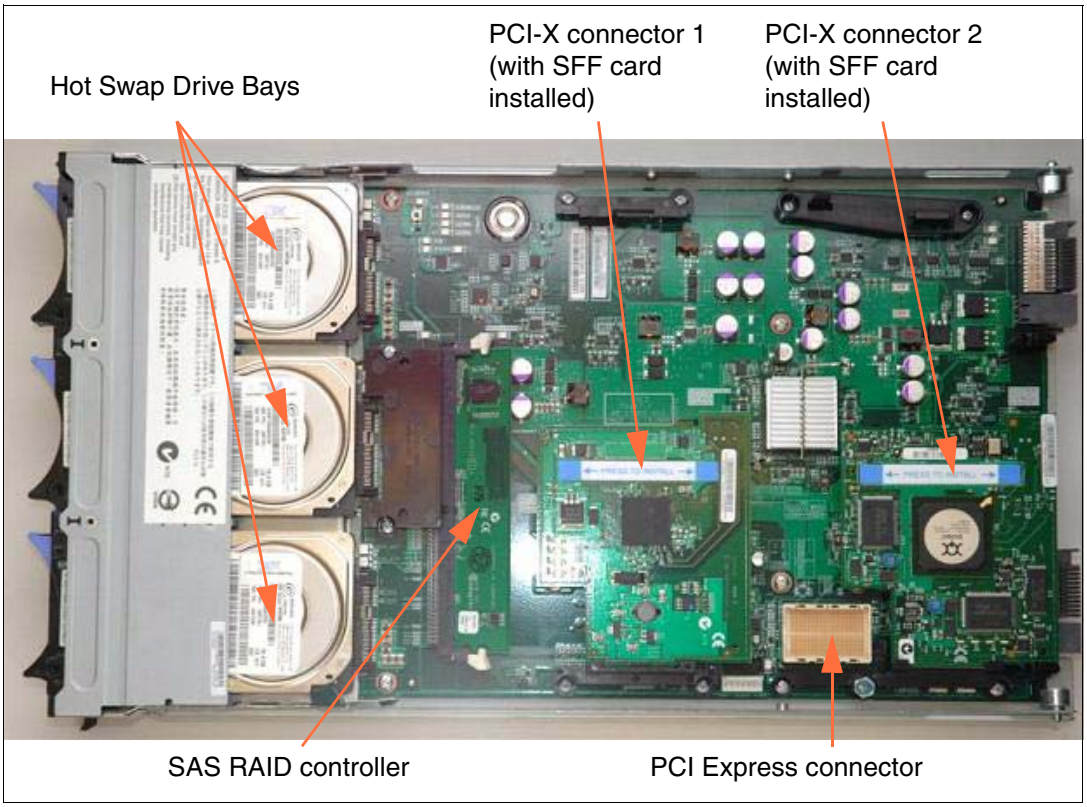

Figure 3-59 Storage and I/O (SIO) Expansion Unit

The BladeCenter Storage and I/O Expansion Unit (SIO) is sometimes also referred to as the BladeCenter Storage Expansion Unit 3 (BSE3).

The SIO unit attaches to the blade server unit via a dedicated connector on the system board. The top cover of the blade server needs to be removed and the SIO unit attaches to it, forming a double-width blade server unit. This will limit the maximum number of blade servers that can fit into a chassis.

Figure 3-60 shows an HS21 8853 blade server with BladeCenter Storage and I/O (SIO) Expansion unit.

Figure 3-60 HS21 8853 with the Storage and I/O (SIO) Expansion Unit attached

Storage configuration

See Table 3-94 for BladeCenter Storage and I/O (SIO) Expansion Unit Options.

Table 3-94 BladeCenter Storage and I/O (SIO) Expansion Unit Options

Part number	Description
39R7389	73 GB SFF 10K HS drive
43X0853	73 GB SAS 15K SFF HS drive
43X0832	146 GB SAS 10K SFF HS drive
25R8064	IBM ServeRAID-8k SAS Controller (replaces the ServeRAID 8k-l)

Note: The tray on the supported hot-swap hard disk drives is different from the tray used in System x servers and the HS12.

The standard RAID controller in the SIO is an ServeRAID 8k-l which is a small, DIMM-like adapter that plugs into a dedicated connector on the SIO motherboard. It has the following characteristics:

- Support for RAID levels 0, 1, and 10
- No cache
- No battery backup

The SIO also supports the optional ServeRAID 8k, part number 25R8064. The ServeRAID 8k RAID controller has the following features:

- Support for RAID levels 0, 1, 1E, 5, 6 and 10
- 256 MB DDR2 533 MHz memory unbuffered DIMM cache
- Battery backup

Depending on the blade server used, you might or might not be able to form an array spanning the SAS drives internal to the blade server and those in the SIO (this is referred to as *drive annexation*).

- Blade Server HS21, LS21, and LS41

  When the SIO is used with one of these servers, the ServeRAID 8k-l in the SIO (or ServeRAID 8k, if installed) can form arrays using combinations of SAS drives in both the SIO and in the attached blade server. This means that you can form one RAID array of four or five drives: three in the SIO and one or two in the blade server (depending on the server model).

  As mentioned, the ability to give control of the internal disk drives to the RAID controller on the SIO is known as drive annexation. Drive annexation needs to be set in the BIOS of the blade server. All disk drives need to be SAS for drive annexation.

- Blade Server HS12 and HS21 XM

  With the HS12 and HS21 XM, the ServeRAID controller in the SIO does not communicate with the SAS drives in the blade server. This means that you can only form RAID arrays using the three disk drives in the SIO with the ServeRAID 8k-l or ServeRAID 8k. The disk drives on the attached blade server can still be installed, but they will be controlled separately by the blade server controller.

For more specifics about these RAID devices, and for a comparison of features of members of the ServeRAID family, see:

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

Networking configuration

Each blade can support up to three I/O expansion cards, one in each of three connectors, as shown in Figure 3-61.

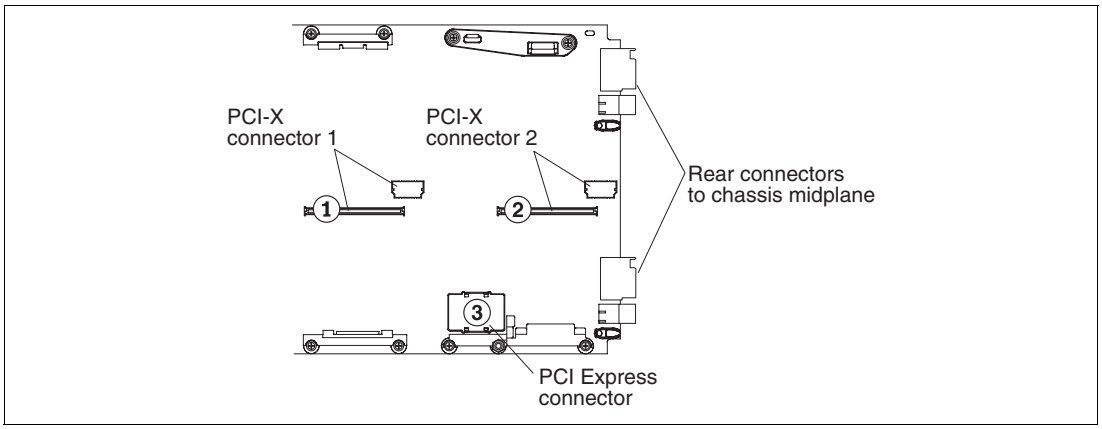

Figure 3-61 PCI-X and PCI Express connectors on the SIO

The connectors are as follows:

- Connector 1, a PCI-X connector suitable for SFF, StFF, or CFFv expansions cards
- Connector 2, a PCI-X connector suitable for SFF, StFF, or CFFv expansions cards

- Connector 3, a PCI Express connector suitable for HSFF, or CFFh expansions cards

Note: Connector 3 is only usable when the server and SIO are installed in a BladeCenter H or BladeCenter HT chassis.

The connection rules are explained here:

- You can install any StFF, SFF, or CFFv expansion card in connector 1. It must be an Ethernet expansion card because this connector is routed to module bays 1 and 2 in the BladeCenter chassis.
- You can install any StFF, SFF, or CFFv expansion card in connector 2, but these limit what cards can be used in connector 3. This connector is routed to module bays 3 and 4 in the BladeCenter chassis.
- You can install an HSFF expansion card in connector 3, but if you do, you cannot install any card in connector 2 (as they occupy the same physical space).
- You can install a CFFh expansion card in connector 3, but if you do, you can only install a CFFv expansion card in connector 2.
- If you also plan to also attach a PEU2 expansion unit to the SIO, then it connects using the PCI Express connector 3, so you cannot also install a CFFh or HSFF expansion card.
- If you install the SIO on a HS21 type 8853 blade server with an MIO also installed, then the two PCI-X connectors on the SIO will not function. The PCIe connector on the SIO will function.

If there is an HSFF or CFFh expansion card installed in the blade server, it will need to be removed to install the SIO. It can then be re-installed in the SIO Expansion Unit.

The expansion cards installed must match the switch modules installed in the chassis switch module bays.

- Connector 1 requires an Ethernet switch module in bays 1 and 2.
- Connector 2 requires a switch module in bays 3 and 4.
- Connector 3 requires a switch module in bays 7–10 (BladeCenter H or HT only).

3.18.3 BladeCenter SCSI Storage Expansion Unit II

The BladeCenter SCSI Storage Expansion Unit II (BSE2 or BSE II), part number 39R8625 (also named the IBM BladeCenter SFF SCSI Storage Expansion Unit), provides the following:

- Two hot swap 3.5 in. SCSI HDD bays for local SCSI hard disk drives for operating system installation.
- Two additional I/O expansion connectors:
  - The first I/O expansion slot connects through to module bays 1 and 2
  - The second I/O expansion slot connects through to module bays 3 and 4

Only the HS20 type 8843 supports this expansion unit as shown in Table 3-93 on page 361.

See Table 3-95 for supported hard drives.

Table 3-95 BladeCenter Storage Expansion Unit II (BSE-II) Options

Part number	Description
40K1026	36 GB (Gen 3) Hot-Swap 3.5" 15K RPM Ultra320 SCSI HDD
40K1023	73 GB Hot-Swap 3.5" 10K RPM Ultra320 SCSI HDD

Part number	Description
40K1027	73 GB (Gen 3) Hot-Swap 3.5" 15K RPM Ultra320 SCSI HDD
40K1024	146 GB Hot-Swap 3.5" 10K RPM Ultra320 SCSI HDD
40K1028	146 GB (Gen 3) Hot-Swap 3.5" 15K RPM Ultra320 SCSI HDD
40K1025	300 GB Hot-Swap 3.5" 10K RPM Ultra320 SCSI HDD

If you require hot swap SCSI hard disk drives with 15K RPM or 300 GB capacity, refer to the ServerProven Web site. There you find supported hard disk drives under the BladeCenter Expansion Unit which is listed under the blade server type you intend to use:

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

The BSE-II Unit can both accommodate the new small form factor (SFF) I/O expansion cards and the older standard form factor (StFF) I/O expansion cards.

The BSE-II Unit attaches to the blade server unit via a dedicated connector on the system board. The top cover of the blade server needs to be removed and the BSE-II unit attaches to it, forming a double-width blade server unit.

It is possible to use two internal SCSI SFF hard disk drives and a BSE-II Unit with two hot swappable hard disk drives at the same time, providing more local storage capacity. The BSE-II unit uses the onboard LSI 53C1020 SCSI controller of the 8843 blade server. However, it is only possible to create one RAID array with all four hard disk drives, for example, creating a RAID-1E. If you have 2 disks, it is a RAID-1 mirror. If you have more than 2 disks, they can be used as additional single disks, a hot spare, or included in the array. So, for 3 drives, you can have a 3 drive RAID-1E, for 4 drives, a 4 drive RAID-1E, or a 3 drive RAID-1E with a hot spare drive.

Note: The BSE-II Unit is different from the BSE-I Unit because the interface and connector used between the unit and the blade system board are different in specification. Also, while the BSE-I unit had an additional SCSI controller, the BSE-II unit uses the onboard SCSI controller of the blade server. In addition, the physical location of the connector is different.

3.18.4 BladeCenter PCI I/O Expansion Unit 2

The PCI Expansion Unit II (PEU2), 25K8373, attaches to supported blade servers via a connector on the server motherboard. The PEU2 has two standard full-length PCI-X slots which operate at a maximum of 100 MHz if two cards are installed, or 133 MHz if only one card is installed.

Only one PEU2 can be attached to the blade server. See 3.18, “Expansion blades” on page 358, for a discussion of all expansion units and the supported combinations for each blade server type. Table 3-92 on page 358 shows the blade servers that support the attachment of the PEU2.

For the list of PCI-X adapters supported in the two full-length slots of the PEU2, see ServerProven at:

http://www.ibm.com/servers/eserver/serverproven/compat/us/peu/25K8373.html

Figure 3-62 PCI I/O Expansion Unit II (PEU2)

There are no PCI-X or PCI Express sockets on the PEU2 suitable for BladeCenter expansion cards such as the StFF, CFFh, or CFFv expansion cards.

Figure 3-63 shows the internal circuitry of the PEU2.

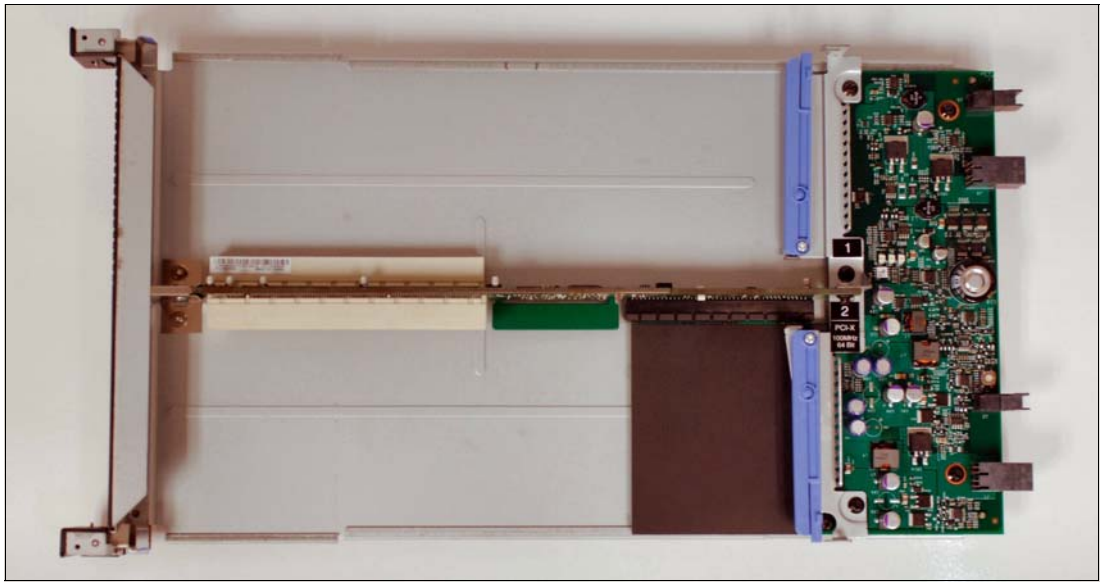

Figure 3-63 PCI I/O Expansion Unit II internals

3.18.5 PCI Express I/O Expansion Unit (PEU3e)

The PCI Express I/O Expansion Unit (PEU3e), part number 43W4391, attaches to supported blade servers via a connector on the server motherboard. The PEU3e has two standard full-length PCI Express slots that support the attachment of 2 full sized, 25W PCI Express adapters. These two slots are PCI Express x16 slots mechanically, however they can only operate at speeds up to x8. The PEU3e provides two power connectors that can be used to provide 12V, via a cable, to PCI Express adapters that have external power connectors.

Note: Using PCI Express adapters that require more than 25W requires special testing and support to be completed. The cable or cables used to connect to the 12 volt connectors on the PEU3e are only available if this testing is completed successfully. Contact your IBM representative to engage the SPORE process.

Only one PEU3e can be attached to the blade server. When a PEU3e is added to a blade server along with the Storage I/O Expansion (SIO), only slot 1 of the PEU3e will be functional. See 3.18, “Expansion blades” on page 358, for a discussion of all expansion units and the supported combinations for each blade server type. Table 3-92 on page 358 lists the blade servers that support the attachment of the PEU3e.

For the list of PCI Express adapters that are supported in the two full-length slots of the PEU3e, see ServerProven at:

http://www.ibm.com/servers/eserver/serverproven/compat/us/peu/43W4391.html

Figure 3-64 PCI Express I/O Expansion Unit (PEU3e)

There are no PCI-X or PCI Express sockets on the PEU3e that are suitable for BladeCenter expansion cards such as the StFF, CFFh, or CFFv expansion cards. Therefore, if the PEU3e is attached to a blade server, then it is not possible to also install a CFFh I/O adapter.

Figure 3-65 on page 370 shows the internal circuitry of the PEU3e.

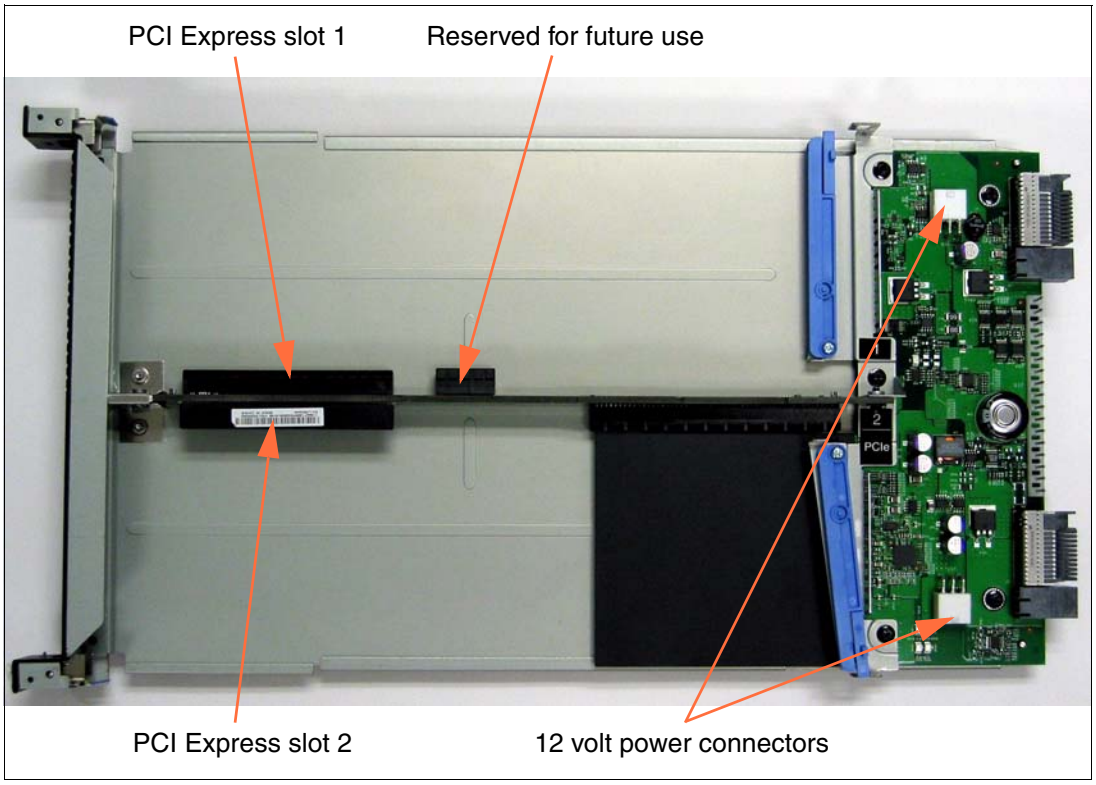

Figure 3-65 PCI Express I/O Expansion Unit internals

4

Network integration

This chapter describes different aspects of planning and implementing an IBM BladeCenter-based network infrastructure. You need to take several considerations into account to achieve a successful implementation. These considerations include systems management, performance, high availability and redundancy features, VLAN implementation, interoperability, and others.

This chapter includes the following topics:

4.1 Ethernet switch module selection

There are a number of I/O modules that can be used to provide network connectivity. They include Ethernet switch modules that provide integrated switching capabilities for BladeCenter and pass-thru modules that make internal blade ports available to the outside. Plan to use the Ethernet switch modules whenever possible, because they often provide the required functionality and simplified cabling. However, under some circumstances (for instance, specific security policy or certain network requirements) it is not possible to use integrated switching capability so pass-thru modules are needed.

Make sure that the external interface ports of the switches selected are compatible with physical cabling used or planned to be used in the client's data center (for information about supported cables, refer to 2.2, "Ethernet switch modules" on page 76).

Also make sure that features and functions required in the network are supported by proposed switches modules (refer to Table 2-10 on page 77 and switch module-related information in 2.2, "Ethernet switch modules" on page 76).

If you plan to use pass-thru modules, be aware that the Copper Pass-thru Module and Optical Pass-thru Module support 1 Gb Ethernet connections only (there is no support for 10 Mbps or 100 Mbps Ethernet), so the external switches must have the required number of 1 Gbps Ethernet ports.

Currently, you can use a wide range of network bandwidth capabilities in the BladeCenter chassis, starting from basic end-to-end Gigabit Ethernet connectivity, through 10 Gb uplinks to the external switches, and finishing with end-to-end 10 Gb Ethernet connectivity (10 Gb ports on the blades itself).

Table 4-1 lists common selection considerations that might be useful to you when selecting an appropriate switch module. Refer to Table 2-10 on page 77 for an expanded feature comparison.

Table 4-1 Switch module selection criteria

Requirement	Suitable switch module
Gigabit Ethernet Uplinks	► "IBM Server Connectivity Module" on page 113 ► "Cisco Systems Intelligent Gigabit Ethernet Switch Module" on page 84 ► "Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module" on page 79 ► "Cisco Catalyst Switch Module 3110G" on page 89 ► "Cisco Catalyst Switch Module 3012" on page 106 ► "Nortel Networks L2/3 Copper GbE Switch Module" on page 115 ► "Nortel Networks L2/3 Fiber GbE Switch Module" on page 120 ► "Nortel 1/10 Gb Uplink Ethernet Switch Module" on page 130 ► "Nortel Networks Layer 2-7 Gigabit Ethernet Switch Module" on page 123
10 Gb Ethernet Uplinks	► "Cisco Catalyst Switch Module 3110X" on page 97 ► "Nortel 10 Gb Uplink Ethernet Switch Module" on page 126 ► "Nortel 1/10 Gb Uplink Ethernet Switch Module" on page 130
Basic Layer 2 Gigabit Ethernet switching connectivity (VLAN, port aggregation)	► "IBM Server Connectivity Module" on page 113

Requirement	Suitable switch module
Advanced Layer 2 Gigabit Ethernet switching: standards-based features (STP, QoS)	► “Cisco Systems Intelligent Gigabit Ethernet Switch Module” on page 84 ► “Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module” on page 79 ► “Cisco Catalyst Switch Module 3110G” on page 89 ► “Cisco Catalyst Switch Module 3110X” on page 97 ► “Cisco Catalyst Switch Module 3012” on page 106 ► “Nortel Networks L2/3 Copper GbE Switch Module” on page 115 ► “Nortel Networks L2/3 Fiber GbE Switch Module” on page 120 ► “Nortel 10 Gb Uplink Ethernet Switch Module” on page 126 ► “Nortel 1/10 Gb Uplink Ethernet Switch Module” on page 130
Advanced Layer 2 Gigabit Ethernet switching: proprietary features (CDP, VTP)	► “Cisco Systems Intelligent Gigabit Ethernet Switch Module” on page 84 ► “Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module” on page 79 ► “Cisco Catalyst Switch Module 3110G” on page 89 ► “Cisco Catalyst Switch Module 3110X” on page 97 ► “Cisco Catalyst Switch Module 3012” on page 106
Layer 3 Gigabit Ethernet switching (IP routing, filtering)	► “Cisco Catalyst Switch Module 3110G” on page 89 ► “Cisco Catalyst Switch Module 3110X” on page 97 ► “Cisco Catalyst Switch Module 3012” on page 106[a] ► “Nortel Networks L2/3 Copper GbE Switch Module” on page 115 ► “Nortel Networks L2/3 Fiber GbE Switch Module” on page 120 ► “Nortel 10 Gb Uplink Ethernet Switch Module” on page 126 ► “Nortel 1/10 Gb Uplink Ethernet Switch Module” on page 130
Layer 3 10 Gb Ethernet switching on all ports (ip routing, filtering)	► “Nortel 10 Gb Ethernet High Speed Switch Module” on page 134
Layer 3 IPv6 routing	► “Cisco Catalyst Switch Module 3110G” on page 89 ► “Cisco Catalyst Switch Module 3110X” on page 97 ► “Nortel 1/10 Gb Uplink Ethernet Switch Module” on page 130[b]
Layer 4-7 Gigabit Ethernet switching (content-based switching, traffic inspection, server load balancing)	► “Nortel Networks Layer 2-7 Gigabit Ethernet Switch Module” on page 123
InifiniBand-to-Gigabit Ethernet connectivity	► “QLogic InfiniBand Ethernet Bridge Module” on page 141

a. This supports IP forwarding based on static routes and routes learned via RIP only.
b. This feature is planned for release in 2008.

4.2 VLAN

VLANs are commonly used in the Layer 2 network. VLAN considerations include the number and types of VLANs supported, management VLANs, VLAN tagging protocols supported, and specific VLAN configuration protocols implemented (such as Cisco VLAN Trunking Protocol).

All switch modules for IBM BladeCenter (including Cisco) support the 802.1Q protocol for VLAN tagging. Additionally, the ISL protocol, which is the Cisco proprietary protocol for VLAN tagging, is supported on Cisco Catalyst 3110X, 3110G, and 3012.

Management VLANs are discussed in more detail in 4.5, “Systems management” on page 380.

Another use of 802.1Q VLAN tagging is to divide one physical Ethernet interface on a blade into several logical interfaces that belong to different VLANs. In other words, a blade server can send and receive tagged traffic from different VLANs on the same physical interface. This can be done with Broadcom Advanced Server Program (BASP) software (the same as used for NIC teaming, as described in 4.4.3, "NIC teaming" on page 379). Each logical interface will appear as a separate network adapter in the operating system with its own set of characteristics, such as IP addresses, protocols, services, and so on.

Having several logical interfaces can be useful in cases when an application requires more than two separate interfaces and you do not want to dedicate a whole physical interface to it for some reason (not enough interfaces or low traffic, as examples). It may also help to implement strict security policies for separating network traffic using VLANs while having access to server resources from different VLANs, without needing to implement Layer 3 routing in the network.

To be sure that the application deployed supports logical interfaces, check the application documentation for possible restrictions applied to the NIC teaming configurations, especially in the case of a clustering solutions implementation.

Refer to Table 2-10 on page 77 and to switch module-related information in 2.2, "Ethernet switch modules" on page 76 for information about the switch module's supported features and protocols.

4.3 High availability and redundancy

Clients might have requirements for continuous access to their network services and applications. Providing high availability for client network resources is a complex task that involves fitting multiple "pieces" together on a hardware and a software level. One HA component is to provide network infrastructure availability.

Network infrastructure availability can be achieved by implementing certain techniques and technologies. Most of them are widely used standards, but some of them are specific to BladeCenter. This section discusses the most common technologies that can be implemented in a BladeCenter environment to provide high availability for network infrastructure.

In general, a typical LAN infrastructure consists of server NICs, client NICs, and network devices such as Ethernet switches and cables connecting them. The potential failures in a network include port failures (both on switches and servers), cable failures, and network device failures.

To provide high availability and redundancy:

- Avoid or minimize single points of failure; that is, provide redundancy for network equipment and communication links. The BladeCenter chassis has built-in redundancy:
  - Two Ethernet ports on each blade server
  - Two ports on I/O expansion cards on each blade server
  - Two separate communication paths to I/O modules through dual midplane connections
  - Two I/O bays per dual port for device redundancy

  For examples illustrating this redundancy, refer to Figure 2-1 on page 60, Figure 2-2 on page 63, and Figure 2-3 on page 66.

- Implement technologies that will provide automatic failover in case of any failure. This can be done by using certain feature protocols that are supported by network equipment, together with server-side software.

  Consider implementing the following technologies which could help you to achieve a higher level of availability in an IBM BladeCenter network solution (depending on your network architecture):

  - Spanning Tree Protocol
  - Trunk Failover with NIC Teaming
  - Virtual Router Redundancy Protocol
  - Routing Protocol (such as RIP or OSPF)

4.3.1 Redundant network topologies

IBM BladeCenter can be connected to the enterprise network in several ways (see Figure 4-1).

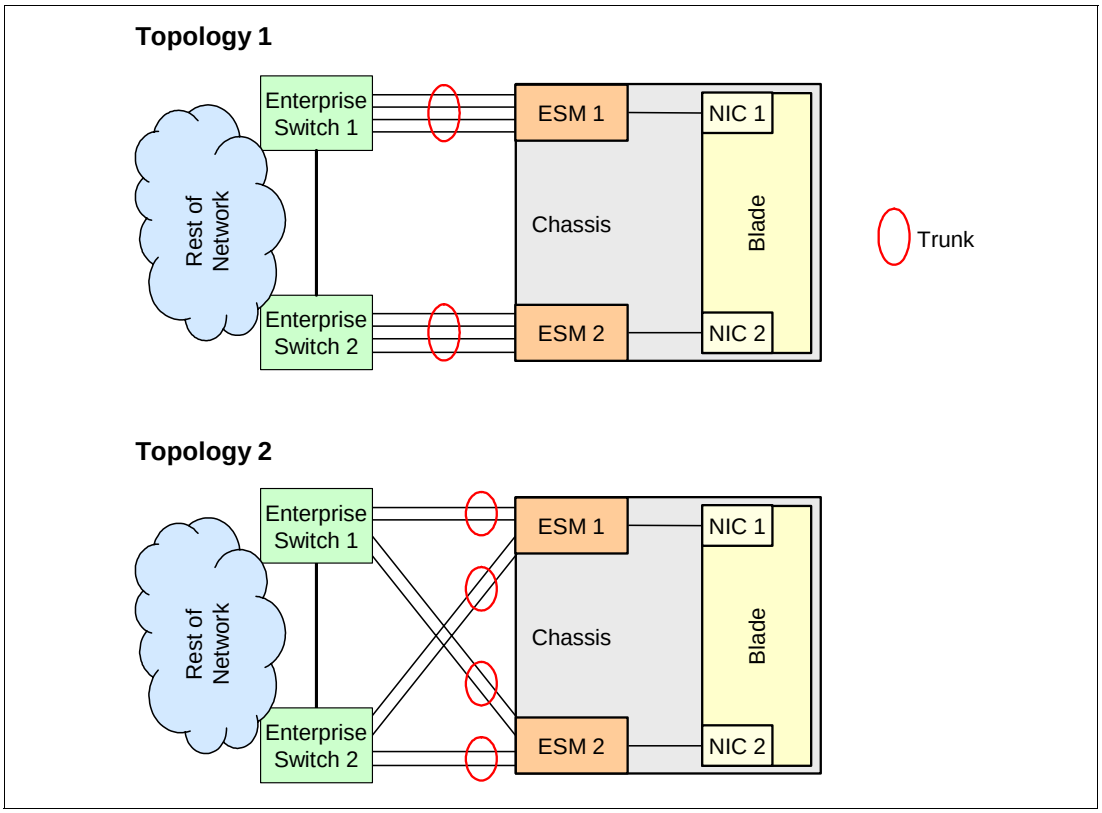

Figure 4-1 IBM BladeCenter redundant LAN integration topologies

Topology 1 in Figure 4-1 has each ESM in BladeCenter directly connected to the one of enterprise switches via aggregation links using all four or six external ports on the ESM. This topology is the simplest way to integrate BladeCenter into an existing network, or to build a new one. The IBM Server Connectivity Module or other ESMs can be used for this configuration.

Topology 2 in Figure 4-1 has each ESM with two direct connections to two enterprise switches. That topology is more advanced, and it has higher level of redundancy, but some specific protocols such as Spanning Tree have to be implemented. This excludes the IBM Server Connectivity Module from the list of switches that can be used in a such way.

The next section describes how these different technologies can be used in these topologies.

4.3.2 Spanning Tree Protocol

Spanning Tree Protocol (STP) is a 802.1D standard protocol used in Layer 2 redundant network topologies (Topology 2 in Figure 4-1 on page 375) to enable automatic network reconfiguration in case of failure. For example, enterprise switches 1 and 2, together with ESM 1, create a loop in a Layer 2 network. We must use STP in that case as a loop prevention mechanism (because a Layer 2 network cannot operate in a loop).

Assume that the link between enterprise switch 2 and ESM 1 is disabled by STP to break a loop, so traffic is going through the link between enterprise switch 1 and ESM 1. In case of link failure, STP will reconfigure the network and activate the previously disabled link. The process of reconfiguration can take tenths of a second, and service will be unavailable during this period of time.

Spanning Tree is required on IBM BladeCenter ESMs only when they are a part of a redundant Layer 2 network topology (see Topology 2 in Figure 4-1 on page 375). In other cases, STP is not required.

Whenever possible, plan to use trunks (aggregated links) with VLAN tagging for interswitch connections. This can help achieve higher performance by increasing interswitch bandwidth, as well as promote higher availability by providing redundancy for links in the aggregation bundle. For more information about link aggregation, see 4.4.1, “Link aggregation” on page 378.

STP modifications, such as Port Fast Forwarding or Uplink Fast, may help to improve STP convergence time as well as the performance of the network infrastructure. Additionally, several instances of STP may run on the same switch simultaneously, on a per-VLAN basis (that is, each VLAN has its own copy of STP to load balance traffic across uplinks more efficiently).

For example, assume a switch has two uplinks in a redundant loop topology and several VLANs are implemented. If single STP is used, then one of these uplinks will be disabled and the other will carry traffic from all VLANs. However, if two STP instances are running, then one link will be disabled for one set of VLANs while carrying traffic from another set of VLANs, and vice versa. In other words, both links will be active, thus enabling more efficient use of available bandwidth.

Refer to Table 2-10 on page 77 for a list of switches supporting STP. Refer to 2.2, “Ethernet switch modules” on page 76, for additional information.

4.3.3 Trunk Failover with NIC Teaming

Depending on the configuration, each blade can have one IP address per each Ethernet port, or it can have one virtual NIC consisting of two or more physical interfaces with one IP address. This is known as an NIC teaming technology. From the BladeCenter perspective, NIC Teaming is useful when you plan to implement high availability configurations with automatic failover in case of internal or upstream link failures.

We can use only two ports on blades per virtual NIC for high availability configurations. One port is active (transferring data), and another is standby (waiting for failure of the active port). One port (for instance, the active port) should be connected to ESM in I/O bay 1, and another port (for instance, the standby port) should be connected to ESM in I/O bay 2. If you plan to use an Ethernet expansion card for high availability configurations, then the same rules apply.

Active and standby ports should be connected to ESMs on different midplanes (that is, active ports to bay 1 or 3, and standby ports to bay 2 or 4).

In the case of internal port or link failure of the active NIC, the teaming driver switches the port roles. The standby port becomes active and the active port becomes standby. This is done very quickly, within a few seconds. After restoration of the failed link, the teaming driver can perform a failback or can simply do nothing, depending on the configuration.

Look at Topology 1 in Figure 4-1 on page 375. Assume that NIC Teaming is on and blade NIC port connected to ESM 1 is active and the other is standby. If something goes wrong with the internal link to ESM 1, then the teaming driver will perform a failover. But what will happen if external connections are lost (the connection from ESM 1 to Enterprise Switch 1)? The answer is that nothing happens because the internal link is still on and the teaming driver does not perform a failover. So network service becomes unavailable.

To address this issue, the Trunk Failover technique is used. Trunk Failover can disable all internal ports on ESM in the case of an upstream links failure. A disabled port means no link, so the NIC teaming driver performs a failover. This is a special feature supported only on IBM BladeCenter ESMs. Thus, if Trunk Failover is enabled and you lose connectivity with Enterprise Switch 1, then the NIC teaming driver will perform a failover and service will be available through Enterprise Switch 2 and ESM 2.

Trunk Failover should be used with NIC active/standby teaming. Before planning to use NIC teaming, verify whether it is supported by the operating system and applications deployed.

All current blade servers use integrated Broadcom Gigabit Ethernet controllers. NIC Teaming can be implemented by using certain software drivers from Broadcom (for instance, the Broadcom Advanced Server Program (BASP) is used with IBM blade servers and Windows and Novell operating systems) or by using specific drivers in the operating system (such as bonding in Linux).

Therefore, if BASP or bonding is used for NIC teaming and you are implementing Trunk Failover, then only two NICs (one active, one standby) can be grouped into one logical interface (BASP supports only one standby NIC per virtual NIC, bonding mode 1 used for active/backup supports only two interfaces). Consult the corresponding documentation for additional information.

When you are planning to use NIC teaming for a NetXen 10 Gb Ethernet High Speed Expansion Card, you can use bonding for Linux operating systems, and either BASP or NetXen teaming drivers for Windows.

Refer to Table 2-10 on page 77 for a list of switches supporting Trunk Failover.

Note: We do not recommend using automatic failback for NIC teaming to avoid possible issues when you replace failed ESM. A newly installed ESM has no configuration data, and it can cause service disruption.

Refer to 2.2, “Ethernet switch modules” on page 76, for additional information.

4.3.4 Virtual Router Redundancy Protocol

If you are integrating IBM BladeCenter into a Layer 3 network with different subnets, routing and routing protocols, some Layer 3 techniques can be used to provide service availability to clients. Traditionally, in multi-subnet IP networks, client workstations and servers use IP default gateways to be able to communicate with each other. In a redundant network, in case

of router failure (which is a default gateway), certain protocols should be used to keep network connectivity between servers and clients. One of them is Virtual Router Redundancy Protocol (VRRP).

VRRP uses the virtual router concept, meaning the logical device that can be hosted by different physical network devices. This protocol ensures that in case of physical network device failure (on which the virtual router was hosted), the virtual router will be up and running on another Layer 3 device.

Some ESMs for IBM BladeCenter can be used for VRRP. Refer to Table 2-10 on page 77 for a list of switches supporting VRRP. Also refer to 2.2, “Ethernet switch modules” on page 76, for additional information.

Cisco Catalyst 3012, 3110G, and 3110X use the Cisco proprietary router redundancy protocol called Hot Standby Routing Protocol (HSRP), which provides almost the same functions as VRRP.

4.3.5 Routing protocols

In a complex IP network with multiple subnets and multiple redundant links, the routing protocols are commonly used to provide integrity of routing information and automatic network reconfiguration capability in case of any failures. Typical standard routing protocols existing in enterprise networks include Routing Information Protocol (RIP) and Open Shortest Path First (OSPF). Additionally, ISPs and other network service providers might use Border Gateway Protocol (BGP).

Refer to Table 2-10 on page 77 for a list of switches supporting routing protocols, and refer to 2.2, “Ethernet switch modules” on page 76 for additional information.

4.4 Performance

Another major topic to be considered during BladeCenter network planning is network performance. Planning network performance is a very complicated task, so the following sections provide guidance regarding the performance features of IBM BladeCenter-based network infrastructures. The commonly used features include link aggregation, jumbo frames, NIC Teaming, and network or server load balancing.

4.4.1 Link aggregation

Link aggregation is a simple way to acquire more network bandwidth between switches. Link aggregation (also commonly referred as a trunk) is a technique that combines several physical links into one logical link to get more bandwidth for network applications. A trunk also provides some level of redundancy for its physical links. That is, if one of the physical links in the trunk fails, then traffic will be distributed between the remaining functional links.

Basically, there are two ways of establishing a trunk: static and dynamic. Static trunks can be mostly used without any limitations. As for dynamic trunks, there are differences between them. There are several protocols for dynamic trunking. First is a standard IEEE 802.1ad protocol also known as LACP. This protocol is supported by the following ESMs:

- IBM Server Connectivity Module
- Nortel Networks L2/3 Copper GbE Switch Module for BladeCenter
- Nortel Networks L2/3 Fiber GbE Switch Module for BladeCenter
- Nortel Networks Layer 2-7 GbE Switch Module

- Nortel 10 Gb Uplink Ethernet Switch Module
- Nortel 1/10 Gb Uplink Ethernet Switch Module
- Nortel 10 Gb Ethernet High Speed Switch Module
- Cisco Systems Intelligent Gigabit Ethernet Switch Module
- Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module
- Cisco Catalyst Switch Module 3110G
- Cisco Catalyst Switch Module 3110X
- Cisco Catalyst Switch Module 3012

Additionally, Cisco Systems uses its own protocol called Port Aggregation Protocol (PAgP). It is supported by the Cisco ESMs:

- Cisco Systems Intelligent Gigabit Ethernet Switch Module
- Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module
- Cisco Catalyst Switch Module 3110G
- Cisco Catalyst Switch Module 3110X
- Cisco Catalyst Switch Module 3012

Refer to 2.2, “Ethernet switch modules” on page 76, for additional information.

4.4.2 Jumbo frames

Jumbo frames or extended frames are used to speed up server network performance. Unlike a traditional Ethernet frame size of up to 1.5 KB, the Ethernet jumbo frames can be up to 12 KB in size. This can speed up server network processing, and can provide better utilization of network. For instance, jumbo frames can be helpful in iSCSI implementations.

Jumbo frames must be supported by all network devices in the communication path. For instance, if you plan to implement iSCSI storage with jumbo frames, then all components including server NICs, network switches, and storage system NICs must support jumbo frames.

Currently, all IBM BladeCenter ESMs support jumbo frames. Refer to Table 2-10 on page 77 for a list of switches supporting jumbo frames. Refer to 2.2, “Ethernet switch modules” on page 76 for additional information.

4.4.3 NIC teaming

NIC teaming can be used for high availability purposes as described in 4.3.3, “Trunk Failover with NIC Teaming” on page 376. However, in some cases it is necessary to get more network bandwidth for certain servers. NIC teaming can help you to achieve that.

Static trunking or 802.1ad modes of NIC teaming can be used for interfaces connected to the same ESM (for instance, blade’s integrated Ethernet port and a port on an expansion card installed into slot 1 of the storage expansion unit). In other cases when NICs are connected to different ESMs, the BASP Smart Load Balancing or Linux bonding modes 0 or 2 should be used.

Refer to BASP documentation for additional information.

4.4.4 Server Load Balancing

In a scale-out environment, the performance of network applications can be increased by implementing load balancing clusters. There are different methods to do that including IP load balancing (for example, Microsoft Network Load Balancing), application load balancing using specific software features (for instance, WebSphere® Load Balancer), or application load

balancing using hardware features integrated into network devices (for example, Server Load Balancing with Nortel L2/7 GbE Switch for IBM BladeCenter).

Consider using load balancing features of Nortel L2/7 ESM if you plan to implement server farms for terminal services, Web services, or other application services needed to be load balanced. Besides performance, Server Load Balancing also provides high availability by redistributing client's requests to the operational servers in case of any server or application failure.

Server Load Balancing uses virtual server concept similar to virtual router described above, and together with VRRP can provide even higher level of availability for network applications. Also VRRP and Server Load Balancing can be used for inter-chassis redundancy and even disaster recovery solutions.

Server Load Balancing feature is supported by Nortel L2/7 GbE Switch Module for IBM BladeCenter.

Refer to 2.2, "Ethernet switch modules" on page 76 for additional information.

4.5 Systems management

An important part of IBM BladeCenter solution is systems management. This section covers several aspects of management network planning.

4.5.1 Management interfaces

IBM BladeCenter has the following management network interfaces (see Figure 4-2 on page 381):

- Advanced Management Module, external 100 Mbps Ethernet
- Advanced Management Module, internal 100 MBps Ethernet interfaces for I/O module management
- Advanced Management Module, internal RS-485 interface for blade server management

Although Figure 4-2 on page 381 shows management interfaces for the IBM BladeCenter E chassis, the management interfaces for other chassis are essentially the same except for the number of connections: BladeCenter T has eight blade bays, BladeCenter H has an additional six I/O module bays (that is, an additional six internal Ethernet management interfaces per management module), and BladeCenter HT has an additional four I/O bays.

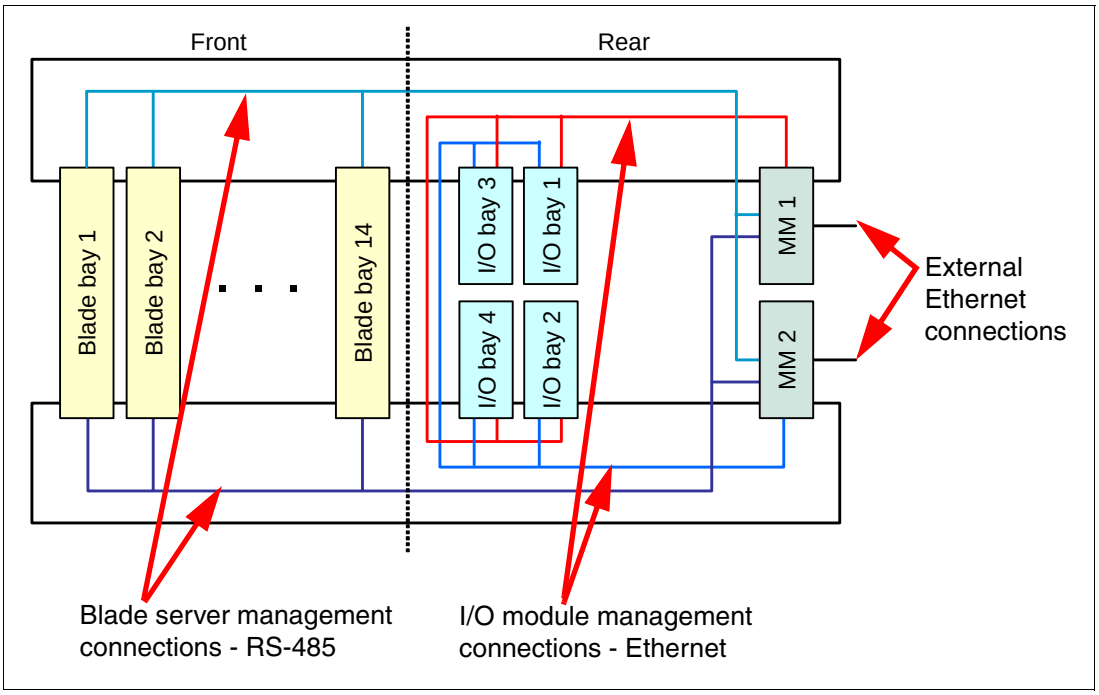

Figure 4-2 Management paths in IBM BladeCenter E chassis

For redundancy purposes, each component has dual management paths. For instance, each I/O module has two Ethernet ports for management, but only one of them is active. Another interface is activated when the primary management module fails or when manual switchover of management modules is initiated.

Devices requiring management in the IBM BladeCenter chassis include I/O modules and blade servers:

- Pass-thru modules cannot be managed and they have no IP addresses.
- Fibre Channel switch modules and InfiniBand switch modules can be managed only through the management module.
- Ethernet switch modules can be managed through the management module or through external Ethernet ports (except for the IBM Server Connectivity Module, which is managed only through internal management interfaces).

If an I/O module is managed through an internal management interface, then it always uses the management module management path. If you plan to use external interfaces of ESMs for management purposes, then certain considerations have to be taken into account as described in 4.5.2, “Management network topology” on page 382.

Each blade server has two RS-485 management paths (one for each management module). Only one path is active at the same time, as with I/O modules. This path is used for functions such as blade power control, VPD, error logs, and remote control.

Serial Over LAN (SOL) for blade server management uses standard Ethernet ports on the blade itself, but SOL traffic is routed to the management module. You can implement SOL only through the management module, and an SOL-capable Ethernet switch is required (pass-thru modules do not support SOL). See Table 2-10 on page 77 for a list of SOL-capable switches.

Whenever possible, plan to perform all management tasks through management module management path. Furthermore, management tasks for some I/O modules such as FCSM or IBM SCM can only be done through management module interface.

4.5.2 Management network topology

The recommended management network topology is shown on Figure 4-3.

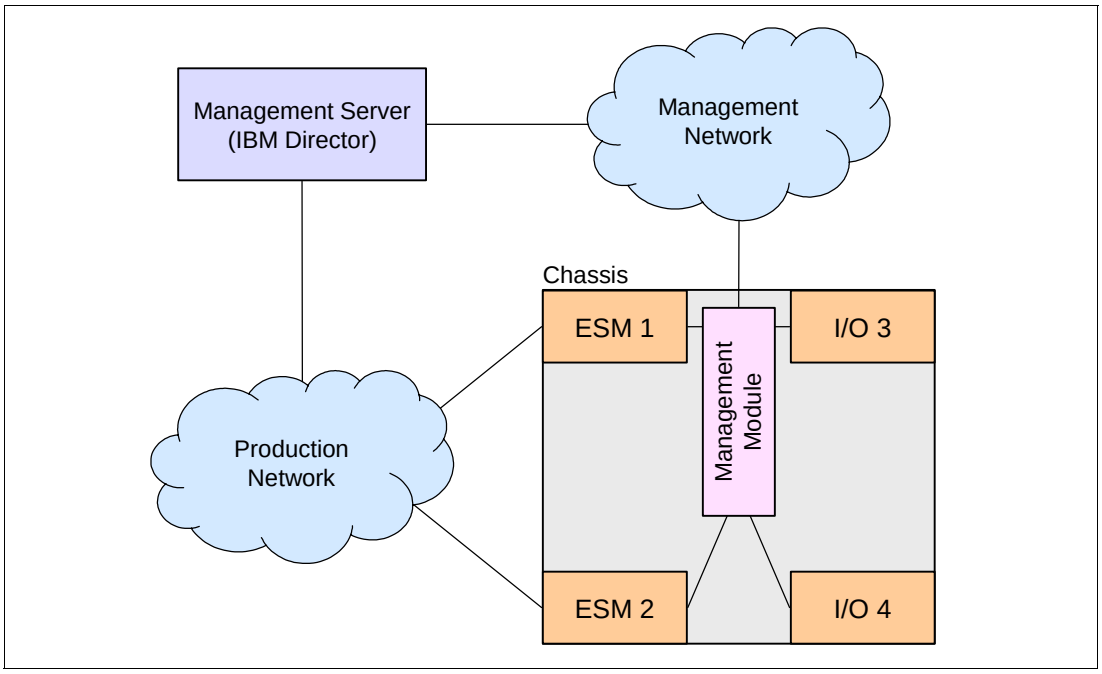

Figure 4-3 Recommended management network topology

Plan to use a dedicated management network for almost all management tasks required to perform for BladeCenter servers and modules. This dedicated network should include management interfaces of all management modules and internal management interfaces of all I/O modules in all chassis deployed. This dedicated management network should have a separate IP subnet (or even IP network) which differs from the production network.

However, if you plan to use a management server such as IBM Director (and it is highly recommended that you do) then server deployment and support tasks (such as installation of the operating system or distribution software after installation or inventory) are performed through the production network (until you designate one of NICs on blades for management purposes, for instance, by assigning them to a management VLAN). NICs dedicated for management are rarely used in a BladeCenter environment; for most environments, it is suitable to use in-band network for operating system deployment.

Sometimes, a management VLAN is implemented across the enterprise network for network management. In this case, the right way to integrate BladeCenter into it is to connect the management module external interface to a port of an external switch assigned to management VLAN. Then assign IP addresses from this management subnet to the management module interface and all internal management interfaces of I/O modules.

Note: In any case, all internal management interfaces of I/O modules and the management interface of the management module must be part of the same IP subnet.

If you plan to use external Ethernet ports on ESMs for in-band management (or dedicate one of them for out-of-band management for some reason), consider the following. Internal management ports on all Nortel switches are always in VLAN 4095 and this cannot be changed. IP addresses assigned by the management module are also always in VLAN 4095. So, if you want to manage a switch via an external Ethernet interface, then you must use a different VLAN on the switch and assign an IP address to that VLAN from a subnet other than the management subnet.

The Nortel 10 Gb Ethernet High Speed switch module is the only external Gigabit Ethernet port dedicated for out-of-band management. It belongs to VLAN 4094, and this cannot be changed.

Note: If you are using an external 1 Gb Ethernet port on a Nortel 10 Gb Ethernet high-speed switch for management, then this port must be on a different subnet than the management module's management interface and an internal management interface of this switch.

Cisco Systems ESMs use a slightly different approach. They use a management VLAN for internal management interfaces. This VLAN can be changed, but internal interfaces will always belong to the new management VLAN. So, if you want to manage a switch through external Ethernet interfaces, then you must use a different VLAN and assign an IP address to that VLAN from a subnet other than the management subnet.

Under some circumstances, such as when the server administration team and the network administration team are two separate teams, it is required that the server team should not be able to control network switch settings. IBM implemented the Protected Mode feature in the Advanced Management Module to deal with these situations. Protected Mode, when enabled, prevents the management module from performing switch administration tasks such as change management IP addresses, reset switch configuration, or answer to ARP requests on behalf of the switch IP address. The switch's firmware must support Protected Mode as well.

Protected Mode is supported on the following ESMs:

- Nortel Networks L2/3 Copper GbE Switch Module for BladeCenter
- Nortel Networks L2/3 Fiber GbE Switch Module for BladeCenter
- Nortel 10 Gb Uplink Ethernet Switch Module
- Nortel 1/10 Gb Uplink Ethernet Switch Module
- Cisco Systems Intelligent Gigabit Ethernet Switch Module
- Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module
- Cisco Catalyst Switch Module 3110G
- Cisco Catalyst Switch Module 3110X
- Cisco Catalyst Switch Module 3012

With the introduction of the Protected Mode feature for the Cisco IGESM, it is now possible to manage the IGESM over its own uplinks, while still having the Management Module and IGESM use the same VLAN and IP subnet. However, the Cisco switch will display an error message that the management module and external management IP addresses are on the same subnet, in this case. This same VLAN and IP subnet could also be shared by servers if so desired. However, best practices still recommend that management and data travel on separate VLANs.

For more details about understanding and implementing the Protect Mode feature, see *Configuring protected mode with the IBM BladeCenter Advanced Management Module and Cisco Intelligent Ethernet Switch Module*, which is available from:

`http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101034`

Note: Even if you are not planning to manage Ethernet switches through the management module and will use Protected Mode and management over external ports, consider using an IP address for switch management from a subnet that differs from the management module's management subnet.

We highly recommend that you use a dedicated management server (such as IBM Director) for centralized systems management. That server should have at least two network interfaces: one for the management network and one for the production network (unless all management tasks are performed through a dedicated management network).

If you plan to use advanced deployment tools such as RDM, then some requirements have to be met especially in the case of a routed network. For instance, if the management server and deployed servers are on different subnets, then routers between them must support DHCP relay. If you use multicast traffic for deployment, then routers must support IGMP. Refer to the documentation of your systems management software for additional information.

4.5.3 Management protocols and tools

When planning to implement a BladeCenter solution, consider using the following management tools for deployment, monitoring, maintenance, support, and optimization tasks:

- For the management of individual servers and modules, the simplest (from the implementation point of view, though not from the usage point of view) way is to use standard telnet or Web-based management for I/O modules.
- For the management of a group of I/O modules, use the available advanced management tools, such as:
  - BladeHarmony Manager from Blade Network Technologies for centralized management of Nortel ESMs using SNMP
  - Cisco Network Assistant to graphically manage groups of Cisco ESMs
  - SAN Surfer or QuickTools for managing QLogic FC switch modules
- BladeCenter Open Fabric Manager provides I/O virtualization capabilities with flexible predefined blade I/O configurations for quick deployment and failover.
- Plan to implement IBM Director for BladeCenter management. IBM Director, together with Remote Deployment Manager, can provide a smart and powerful way of managing your entire infrastructure including servers, switches, and even certain storage systems from IBM.
- IBM BladeCenter management can also be integrated into enterprise management systems such as Tivoli®.

4.6 Virtual Blade Switch technology

Virtual Blade Switch (VBS) technology was introduced by Cisco Systems with the Cisco Catalyst 3110 family of switches for IBM BladeCenter. This technology allows you to combine several physical network switches into one logical entity by using high-speed dedicated stacking ports on these switches. This logical switch appears as a single network device to the blade servers and external network devices. A sample topology is shown in Figure 4-4 on page 385.

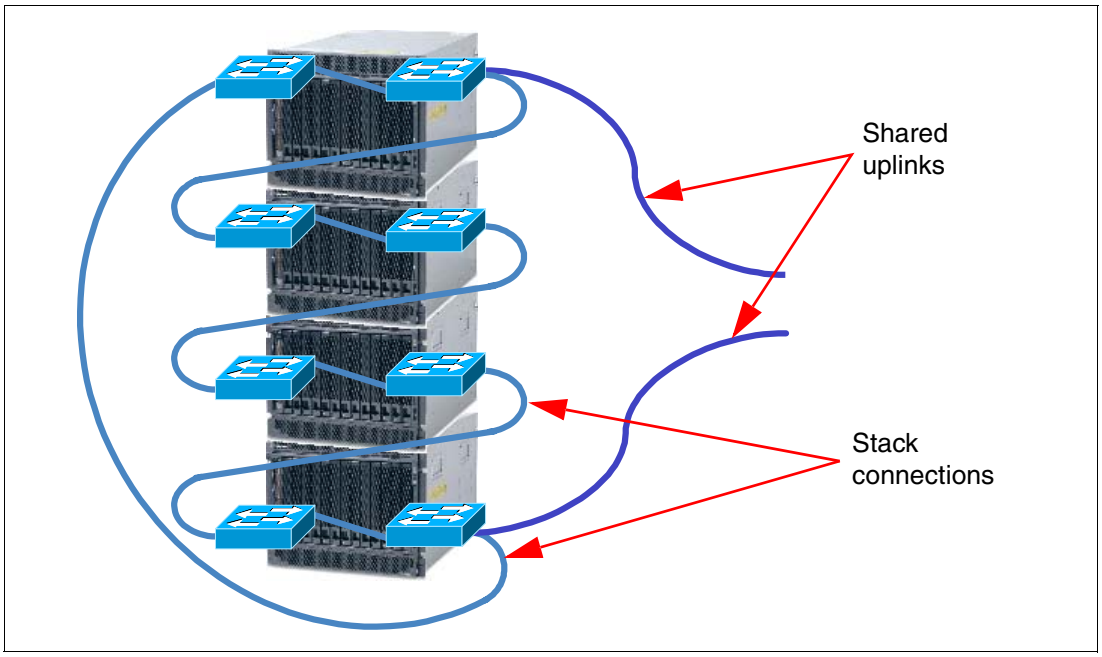

Figure 4-4 Virtual Blade Switch sample topology

A VBS stack has the following major characteristics:

- Up to nine physical switches can be combined in one stack, one master switch, and up to eight member switches.
- Two dedicated connectors on a switch are used to create VBS, and switches are connected in a ring. Stack connections form two counter-rotating unidirectional links, with up to 32 Gbps of raw capacity per stack member (16 Gbps per ring link).
- Single point of management
  - All switches in a stack are managed through a single IP address on the master switch.
  - Single configuration file per stack.
  - Single software upgrade for entire stack.
- Single Spanning Tree Protocol (STP) instance for Layer 2 networks.
- Single router for Layer 3 networks.
- Consolidation of uplink ports (or uplink sharing).

You need to take several considerations into account when planning VBS deployment:

- VBS stacking is well-suited to the concept of rack-level management, so we recommend that you build one VBS (or more VBSes, if the network paths from different server NICs must be physically separated) from the switches in the same rack whenever possible.
- In case of a master switch failure, any other member switch is eligible to become master to provide availability and redundancy.
- In case of a switch failure, the stack connections are looped back to keep the ring operational.
- True active-active NIC teaming (link aggregation of Ethernet NICs) is possible using supported link aggregation protocols like EtherChannel or 802.3ad LACP. For example, for network bandwidth-intensive applications it is possible to use up to four Ethernet ports

per blade combined into a single aggregated bundle (integrated Ethernet and Ethernet ports on expansion cards).

Physically, each port is connected to a different physical switch. However, because of VBS, aggregation is done on a stack level and the blade appears to be connected to the single switch.

- Uplink sharing can help simplify design by reducing the number of external links going out of the rack; that is, not all switches in the stack must have uplinks. For example, instead of using two 1 Gb ports per switch as a redundant uplink for eight switches for a total of 16 cables going out of rack, you may choose to use just two 10 Gb links going out of rack if applicable.
- If a fully redundant topology is required, that is, if each blade server must have two separate paths to the external infrastructure, then you can use two VBS stacks. One stack combines switches from the upper Ethernet switch bays of a chassis, and the other stack combines switches from the lower Ethernet switch bays. In this case, the entire rack will be represented as two separate network switches, both from the blade server side and from the external infrastructure side.
- Networking technologies and protocols, including VLANs, STP and its modifications (such as PVST+, RSTP, MSTP), link aggregation, link state tracking, and routing, are supported on the VBS level as well.

Note: When the switch stack is formed, the Advanced Management Module cannot manage any member of the stack (including the master switch) over IP. Management over external ports or a serial console cable is required.

4.7 BladeCenter Open Fabric Manager

IBM BladeCenter Open Fabric Manager (BOFM) is a tool for simplifying deployment, failover, and repurposing of blade servers by controlling assignment of physical Ethernet MAC and Fibre Channel WWN addresses for each blade's integrated Ethernet interfaces and I/O expansion cards.

In other words, BOFM provides the ability to create predefined configurations for the external infrastructure connections including network switches, SAN switches, storage systems and so on. With BOFM, you can configure MAC-based VLANs and security policies, and WWN-based zones on switches and storage partitions on storage systems by using your own pool of MAC addresses and WWNs created in advance. You can then assign your predefined addresses to a blade's I/O ports with BOFM instead of using burned-in hardware addresses that were assigned during manufacture.

This BOFM functionality, together with boot-from-SAN techniques (it is highly recommended that you plan BOFM together with boot-from-SAN to make the most of the entire solution), enables you to deploy new blades easily by assigning predefined MAC addresses and WWNs to them and, thus, connect them to the existing infrastructure in a minutes. You can also repurpose blades quickly by assigning addresses from different services or even provide simplified failover by assigning addresses from a failed blade to the spare blade.

The BladeCenter Open Fabric Manager offering consists of two components:

- BladeCenter Open Fabric Manager
- BladeCenter Open Fabric Manager - Advanced Upgrade

We describe these offerings in the following sections.

4.7.1 BladeCenter Open Fabric Manager

BladeCenter Open Fabric Manager, part number 44W3981, feature code 4941, is a software upgrade to the Advanced Management Module. It interacts with BMC microcode in the blades and the firmware in the I/O expansion cards to replace the burned-in addresses with ones specified at the blade-bay level by the administrator. The BladeCenter Open Fabric Manager user interface is browser-based.

When a blade server is inserted into a specific bay, it inherits hardware addresses for I/O ports that were assigned to this bay by administrator. This enables simplified deployment of blades because you can preconfigure your infrastructure before the physical installation of the blade server into the chassis. In case of blade failure, all you need to do is to replace the failed blade with a new one.

BOFM allows up to eight ports per blade to be configured, and it supports up to 100 chassis (1400 blades) in a single domain. These blades can be managed from single management module.

Note: Every management module must have the BOFM software update installed.

Solution components that must support BOFM include blade servers, expansion cards, and management modules in the chassis. The BOFM-enabled firmware upgrade must be available for blades (BIOS and BMC), expansion cards, and advanced management module.

BOFM is compatible with most I/O modules because it is not dependent on them. It also uses address reassignment at the blade level rather than at the I/O module level.

With the AMM-based version of BOFM, you can provide quick deployment and redeployment of blades, and manual slot-based failover or manual blade replacement. To provide automated failover capabilities, the BOFM Advanced Upgrade is required.

Solution components that must be supported by BOFM include blade servers, expansion cards, and management modules in the chassis.

Table 4-2 lists supported chassis, blade servers, expansion cards, and I/O options.

Table 4-2 BOFM supported options

Part number	Feature Code	Description	BOFM support
Chassis			
8677	Not applicable	BladeCenter E	Supported
8852	Not applicable	BladeCenter H	Supported
8720, 8730	Not applicable	BladeCenter T	Supported
8740, 8750	Not applicable	BladeCenter HT	Supported
8886	Not applicable	BladeCenter S	Supported
Blade servers			
8014	Not applicable	HS12	Supported
8028	Not applicable	HS12	Supported
8678	Not applicable	HS20	No
8832	Not applicable	HS20	No

Part number	Feature Code	Description	BOFM support
7981	Not applicable	HS20	No
8843	Not applicable	HS20	No
8839	Not applicable	HS40	No
8853	Not applicable	HS21	Supported
7995	Not applicable	HS21 XM	Supported
8850	Not applicable	LS20	No
7971	Not applicable	LS21	Supported
7901	Not applicable	LS22	Supported
7972	Not applicable	LS41	Supported
7902	Not applicable	LS42	Supported
7996	Not applicable	HC10	No
0200	Not applicable	QS20	No
0792	Not applicable	QS21	No
7998-60X	Not applicable	JS12	Supported
8842	Not applicable	JS20	No
7988	Not applicable	JS21	No
7998	Not applicable	JS22	Supported
0792	Not applicable	QS21	No
0793	Not applicable	QS22	Supported
3020	Not applicable	PN41	No
Expansion cards			
None	None	Integrated Gigabit Ethernet	Supported
13N2203	1519	BladeCenter Fibre Channel Card	No
26K4841	1549	BladeCenter 2 Gb Fibre Channel Card	No
73P9030	1547	BladeCenter Gigabit Ethernet Card	No
39R8624	1548	BladeCenter Gigabit SFF Ethernet Card	Supported
32R1896	1492	Cisco Systems InfiniBand 1X HCA Card	No
32R1760	1466	Cisco Systems InfiniBand 4X HCA Card	No
39Y9186	2925	Emulex 4 Gb SFF Fibre Channel Card	Supported
43W6859	2994	Emulex 4 Gb CFFv Fibre Channel Card	Supported
73P6000	1546	Myrinet Cluster Expansion Card	No
32R1923	1458	QLogic iSCSI Expansion Card	No
26R0890	1577	QLogic 4 Gb SFF Fibre Channel Card	Supported
26R0884	1597	QLogic 4 Gb StFF Fibre Channel Card	Supported

Part number	Feature Code	Description	BOFM support
39Y9271	2967	NetXen 10 Gb Ethernet Card	Supported
39Y9310	2969	Ethernet Expansion Card	Supported
39Y9306	2968	QLogic Ethernet and 4 Gb FC Exp. Card	Supported
41Y8527	2970	QLogic 4 Gb Fibre Channel Exp. Card	Supported
39Y9190	2979	SAS Expansion Card	No
43W3974	1591	SAS Connectivity Card (CFFv)	No
46C7167	5490	ServeRAID-MR10ie (CIOv)	No
43W4423	2991	4X InfiniBand DDR Expansion Card	No
43W4421	2992	Cisco 4X InfiniBand DDR Expansion Card	No
43W4420	2993	Voltaire 4X InfiniBand DDR Expansion Card	No
44W4479	5476	2/4 Port Ethernet Exp. Card (CFFh)	No
44W4465	5479	Broadcom 10 Gb 4-port Eth. Exp. Card	No
44W4465	5489	Broadcom 10 Gb 2-port Eth. Exp. Card	No
44X1940	5485	QLogic Eth. & 8 Gb FC Exp. Card (CFFh)	No
Ethernet I/O Modules			
32R1888	1498	Cisco Systems Fiber Intelligent Gigabit ESM	Supported
32R1892	1497	Cisco Systems Intelligent Gigabit ESM	Supported
41Y8523	2989	Cisco Catalyst 3110G ESM	Supported
41Y8522	2988	Cisco Catalyst 3110X ESM	Supported
43W4395	5450	Cisco Catalyst 3012 ESM	Supported
39Y9324	1484	IBM Server Connectivity Module	Supported
32R1860	1495	Nortel L2/3 Copper GbE Switch Module	Supported
32R1861	1496	Nortel L2/3 Fiber GbE Switch Module	Supported
32R1859	1494	Nortel Networks Layer 2-7 Gb ESM	Supported
32R1783	1493	Nortel Layer 2/3 10 Gigabit Uplink ESM	Supported
44W4404	1590	Nortel Layer 2/3 1/10 Gigabit Uplink ESM	Supported
39Y9267	2952	Nortel 10 Gb High Speed Switch Module	Supported
InfiniBand I/O Modules			
32R1900	1533	Cisco Topspin InfiniBand Switch Module	No
32R1756	1574	Cisco Systems 4X InfiniBand Switch Module	No
39Y9207	2941	QLogic InfiniBand Ethernet Bridge Module	No
39Y9211	2942	QLogic InfiniBand FC Bridge Module	No
Fibre Channel I/O Modules			
32R1812	1569	Brocade 4 Gb 20-port SAN Switch Module	Supported

Part number	Feature Code	Description	BOFM support
32R1813	1571	Brocade 4 Gb 10-port SAN Switch Module	Supported
39Y9280	2983	Cisco 4 Gb 20 port FC Switch Module	Supported
39Y9284	2984	Cisco 4 Gb 10 port FC Switch Module	Supported
26R0881	1560	QLogic 4 Gb 20-port FC Switch Module	Supported
32R1904	1575	QLogic 4 Gb 10-port FC Switch Module	Supported
43W6725	2987	QLogic 20-port 4 Gb SAN Switch Module	Supported
43W6724	2986	QLogic 10-port 4 Gb SAN Switch Module	Supported
43W6723	2985	QLogic 4 Gb Intelligent Pass-thru Module	Supported
32R1833	1562	McDATA 4 Gb 20-port FC Switch Module	No
32R1905	1576	McDATA 4 Gb 10-port FC Switch Module	No
44X1905	5478	QLogic 20-Port 8 Gb SAN Switch Module	Supported
44X1907	5482	QLogic 8 Gb Intelligent Pass-thru Module	Supported
Pass-thru Modules			
39Y9316	1556	IBM BladeCenter Optical Pass-thru Module	Supported
39Y9320	2900	IBM BladeCenter Copper Pass-thru Module	Supported
44W4483	5452	Intelligent Copper Pass-thru Module	Supported
43W6723	2985	QLogic 4 Gb Intelligent Pass-thru Module	Supported
44X1907	5482	QLogic 8 Gb Intelligent Pass-thru Module	Supported
43W4419	2990	4X InfiniBand Pass-thru Module	No
SAS I/O Modules			
39Y9195	2980	SAS Connectivity Module	No
43W3584	3734	SAS RAID Controller Module	No

4.7.2 BladeCenter Open Fabric Manager - Advanced Upgrade

BladeCenter Open Fabric Manager - Advanced Upgrade is available in two versions:

- Standalone utility, part number 46C3551
- IBM Director extension, part number 46C3552

This upgrade adds capabilities for managing address assignments through the IBM Director GUI and policy-based automated failover to the standard BOFM. With this upgrade, you can define failover pools of spare blades across chassis and failover rules with Event Action Plans.

Unlike standard BOFM, which uses the management module's Web interface for configuration tasks, the BOFM Advanced Upgrade uses GUI (either IBM Director or standalone utility).

One BOFM license is required per chassis. BOFM Advanced Upgrade (46C3551 or 46C3552) requires a standard BOFM license (44W3981).

Notes:

- The BladeCenter Open Fabric Manager - Advanced Upgrade standalone utility is not compatible with the Active Energy Manager standalone version. IBM Director and its corresponding extensions must be used in this case.
- Currently, BladeCenter Open Fabric Manager does not support address assignments and failover for iSCSI configurations.

4.7.3 Hot spare blades

With the introduction of BladeCenter Open Fabric Manager (BOFM), the process of automating blade failover becomes much easier, because BOFM takes all responsibility for blade address management including LAN MAC addresses and FC WWNs assignments.

Additionally, with BOFM Advanced Upgrade, the BOFM-related tasks can be integrated into the IBM Director environment. This allows you to use a single management GUI, as well as to perform certain BOFM tasks based on IBM Director event triggers. Using BOFM Advanced Upgrade together with Boot from SAN techniques provides an easy and robust way to implement hot spare blades by using Standby Blade Pool.

The idea of the Standby Blade Pool is to designate one or several blades across the BOFM domain as spare blades, and to restore application functionality on one of the spare blades in case of failure of one of active blades, including VLAN configuration.

The process of blade failover consists of several steps (assuming that failover was initiated already by an event trigger or manually):

- BOFM performs a series of checks to fail over successfully.
- BOFM reads addresses from the failed blade through AMM and the port-based VLAN configuration from Ethernet switches.
- BOFM applies addresses to the one of the eligible spare blades and the VLAN configuration to the Ethernet switches in the chassis where the spare blade is installed.
- BOFM powers down the failed blade and powers up the standby blade.

The process of failover takes a certain amount of time, depending on the existing networking environment.

The first step in the failover process performs a series of checks including machine type match, machine model match, VLAN migrations, and others. The perform or skip check step is controlled by the properties file located in C:\Program Files\IBM\Director\data\failover.properties by default.

You can enable or disable a specific check step by editing this file. The content of the file is listed in Example 4-1.

Example 4-1 Content of file failover.properties

```
# 1 = skip check for matching standby blade machine type (default = 0)
ignore_machine=0
# 1 = skip check for matching standby blade model type (default = 1)
ignore_model=1
# 1 = skip check for standby blade initially power state to be off (default = 0)
ignore_power=0
# 1 = skip check for source blade powering off (default = 1)
```

```
ignore_power_off=1
# 1 = skip migration of VLANs from source blade's switches to standby blade's switches (default = 0)
ignore_vlan=0
```

To implement spare blades with BOFM Advanced Upgrade successfully, you must take these considerations into account:

- Blades must be booted from SAN (see 5.8, “Boot from SAN” on page 409 for more information).
- The hardware used supports BOFM (see Table 4-2 on page 387).
- IBM Director is used for systems management.
- Initially, failover can only be done to blades in the standby pool that have exactly the same Machine Type Model (MTM). However, this can be overridden by changing the settings in the failover configuration file.

  In practice, to fail over successfully on different types of server, the software image (operating system and applications) running on the failed blade will be able to boot on different hardware (that is, it will recognize all hardware devices built into the system, and so on).
- If you use the Microsoft Windows Server operating system, we recommend that you use volume licensing to avoid activation issues when the existing Windows Server operating system image will be booted on a different blade server.
- ESMs used in the failover scenario must be supported by IBM Director to be able to transfer VLAN configuration. Currently, IBM Director supports Cisco IGESMs, Nortel L2/3 ESMs, and Nortel L2/7 ESMs as its standard functions.

  Additionally, Nortel 10 Gb Uplink ESMs, Nortel 10 Gb High-speed ESMs, and IBM Server Connectivity Module are supported with a free IBM Director plug-in, the Nortel 10Gb ESM Plugin for IBM Director v5.20.1.
- Up to 1400 blade servers (up to 100 chassis), and about 11000 addresses are supported in BOFM deployments.
- MAC-based VLANs may significantly reduce failover time.

Refer to *IBM BladeCenter Open Fabric Manager Installation and User Guide* for more information. It is available from:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5073103

4.8 InfiniBand and Ethernet LANs

If you plan to implement a high-performance enterprise grid computing solution such as Oracle® RAC on the BladeCenter platform, it might require a high-bandwidth low-latency 4X InfiniBand fabric for internode communications to provide a scalable high-performance environment. At the same time, to reduce the costs of implementation and management, you might choose to use the same InfiniBand fabric for access to the network resources on your traditional LAN or SAN.

In this case, the QLogic InfiniBand Ethernet Bridge Module or QLogic InfiniBand Fibre Channel Bridge Module can help integrate InfiniBand with your other networks. This section gives you a high-level overview of the InfiniBand-to-Ethernet bridging solution for planning purposes. An InfiniBand-to-Fibre Channel bridging solution is discussed in 5.11, “InfiniBand and Fibre Channel SANs” on page 432.

One useful feature of InfiniBand is that it can simultaneously run multiple upper-layer protocols (ULPs) like MPI, IP over InfiniBand (IPoOB) or Virtual NIC. This allows you to use a single InfiniBand fabric for different types of communications, including integration into an existing Ethernet network with the QLogic InfiniBand Ethernet Bridge Module. In this case, each blade that requires access to network resources (or clients that require network access to these servers) can be configured with one or more Virtual NICs, and these Virtual NICs provide the required communications.

The high-speed InfiniBand fabric itself consists of:

- One or more high speed-capable chassis (IBM BladeCenter H or BladeCenter HT)
- Blade servers with 4X InfiniBand expansion cards (one card per blade)
- Cisco Systems 4X InfiniBand High Speed Switch Module (up to two modules per chassis can be used concurrently, and they must be installed into I/O bays 7 and 9 only)
- One or more QLogic InfiniBand Ethernet Bridge Modules (depending on your requirements.) However, up to two bridge modules per chassis can be used concurrently, and they must be installed into I/O bays 3 and 5.

To be able to communicate with your resources with InfiniBand using ULPs, a host-side software called QuickSilver Fabric Access Software is required. This software (as well as the user guides for the software and bridge modules) can be found on the following resource:

http://support.qlogic.com/support/oem_detail_all.asp?oemid=377

The QuickSilver Fabric Access Software is supported with both Linux and Windows operating systems, but Linux supports more advanced functionality such as Port Recovery, which provides failover capabilities to the Virtual NICs with BladeCenter's redundant infrastructure.

The ULP called Virtual NIC is used for integration with Ethernet LANs. It allows you to create virtual network interfaces on a server that can be used in the same way as physical NICs. This protocol, in conjunction with the Ethernet bridge module, supports the following major features and functions that can be implemented in your network:

- VLANs and 802.1Q VLAN tagging are used to separate network traffic in the same manner as done with traditional Ethernet networks, including 802.1Q frame tagging.
- 802.3ad Link Aggregation provides increased bandwidth and redundancy by combining up to three external physical Ethernet links on the bridge module into one logical bundle.
- Port Sharing provides the ability to use up to 24 active Virtual NICs per one Ethernet bridge module (up to six external Ethernet ports) simultaneously.
- Port Recovery (Linux only) can be used to provide failover capability to the Virtual NICs, including failover to another port of InfiniBand HCA expansion card, failover to another Ethernet port of the same bridge module, or failover to another bridge module (in the same chassis or in a different chassis).

5

Storage integration

IBM BladeCenter offers several possibilities of integration into storage infrastructure, such as SAS, Fibre Channel or iSCSI. This chapter covers major considerations to take into account during IBM BladeCenter storage infrastructure planning. These considerations include storage system interoperability, switch module selection and interoperability rules, performance, high availability and redundancy, backup, boot from SAN, and others.

This chapter covers the following topics:

5.1 External storage for IBM BladeCenter

There are currently several options to use to attach external storage systems to IBM BladeCenter, including:

- Directly-attached SAS storage
- Storage area networks (SANs) based on iSCSI
- SANs based on Fibre Channel technologies
- InfiniBand-based storage with native InfiniBand interfaces, or indirectly with InifiniBand-to-Fibre Channel bridges

Traditionally, Fibre Channel-based SANs are the most common and advanced design of external storage infrastructure. They provide high levels of performance, availability and redundancy, and scalability. However, the cost of implementing FC SANs will be higher in comparison with SAS or iSCSI. The major components of almost every FC SAN include server's HBAs, FC switches, FC storage servers, FC tape devices, and optical cables for connecting these devices to each other.

iSCSI-based SANs are becoming more and more popular today because they provide all the benefits of centralized shared storage in terms of storage consolidation and adequate levels of performance, but use traditional IP-based Ethernet networks instead of expensive optical cabling. iSCSI SANs consist of server hardware iSCSI adapters or software iSCSI initiators, traditional network components such as switches, routers, and so forth, and storage servers with an iSCSI interface such as IBM System Storage DS3300.

SAS is a relatively new technology that is used primarily for server direct-attached storage. However, it can also be used for external shared storage attachments as in the IBM BladeCenter Boot Disk System solution or IBM System Storage DS3200.

InfiniBand is a high-bandwidth, low-latency universal fabric for connecting different types of devices including storage systems. Storage systems can have native InfiniBand interfaces to connect to InfiniBand fabric directly, or FC interfaces to connect to InfiniBand fabric through InfiniBand-to-FC bridge modules.

iSCSI, SAS, and FC-based SANs can be used for diskless blade solutions to provide greater levels of utilization, availability, and cost effectiveness.

5.2 Fibre Channel storage system interoperability

In general, if IBM BladeCenter is integrated into FC storage fabric then ensure that following requirements are met. Check the compatibility guides from your storage system vendor for confirmation.

- IBM BladeCenter server hardware and HBA are supported by the storage system
- The FC fabric used or proposed for use is supported by the storage system
- The operating systems deployed are supported both by IBM blades and the storage system
- Multipath drivers exist and are supported by the operating system and storage system (in case you plan for redundancy)

If any of these requirements are not met, then consider proposing another solution that is supported.

Almost every vendor of storage systems or storage fabrics has extensive compatibility matrixes, which include supported HBAs, SAN switches, and operating systems.

Refer to the following publications for additional information:

- *DS3000 Interoperability Matrix*
- *DS4000 Interoperability Matrix*
- *IBM BladeCenter 4Gb SAN Solution*, SG24-7313

5.3 FC switch selection and fabric interoperability rules

IBM BladeCenter provides integrated FC switching functionality by using several switch options (see 2.4, “Fibre Channel switch modules” on page 144):

- Brocade 10 and 20 port 4 Gb SAN Switch Modules (32R1813 and 32R1812)
- QLogic 20 port 4 Gb Fibre Channel Switch Modules (26R0881)
- QLogic 10 and 20 port 4 Gb SAN Switch Modules (43W6725 and 43W6724)
- McDATA 10 and 20 port 4 Gb Fibre Channel Switch Modules (32R1905 and 32R1833)
- Cisco 10 and 20 port 4 Gb Fibre Channel Switch Modules (39Y9280 and 39Y9284)
- QLogic 20-Port 8 Gb SAN Switch Module (44X1905)

An Optical Pass-thru Module (39Y9316) can be used in cases when IBM BladeCenter has to be connected directly to an external FC enterprise switch.

Note: The Optical Pass-thru Module (39Y9316) only supports 2 Gb Fibre Channel and not 4 Gb or 8 Gb Fibre Channel.

Brocade Access Gateway is a firmware upgrade for the existing Brocade 4 Gb FC switch, which enables certain NPIV processing on it to provide intelligent pass-thru capabilities for the BladeCenter chassis.

QLogic Intelligent Pass-thru Module (43W6723) and QLogic 8 Gb Intelligent Pass-thru Module (44X1907) are embedded solutions that provide seamless integration of BladeCenter into existing Fibre Channel fabric, avoiding any multivendor interoperability issues by using N_Port ID Virtualization (NPIV) technology. This module adds ports to the fabrics, and not domain IDs like switches, and it is absolutely transparent to the switches in the fabric. Following is a brief overview describing how the NPIV concept works for the Intelligent pass-thru Module and the Brocade Access Gateway.

In general, there are several basic types of ports used in Fibre Channel fabrics, including:

- N_Ports (Node ports) represent an end-point FC device (such as host, storage system, or tape drive) connected to the FC fabric.
- F_Ports (Fabric ports) are used to connect N_Ports to the FC switch (that is, the host HBA’s N_port is connected to the F_Port on the switch).
- E_Ports (Expansion ports) provide interswitch connections. That is, if you need to connect one switch to another, then E_ports are used (the E_port on one switch is connected to the E_Port on another switch).

When one switch is connected to another switch in the existing FC fabric, it uses the Domain ID to uniquely identify itself in the SAN (like a switch address). Because every switch in the fabric has the Domain ID and this ID is unique in the SAN, this can limit the number of switches and number of ports (for example, QLogic theoretically supports up to 239 switches, and McDATA supports up to 31 switches), thus limiting SAN scalability.

Another problem with E_Ports is an interoperabilty issue between switches from different vendors. In many cases only the so-called “interoperability mode” can be used in these fabrics, thus disabling most of the vendor’s advanced features.

On the other hand, each switch requires some management tasks to be performed on it, and an increased number of switches can increase the complexity of the management solution, especially in heterogeneous SANs consisting of multivendor fabrics. NPIV technology helps to address these issues.

Initially, NPIV technology was used in virtualization environments to share one HBA by multiple virtual machines and assign unique port IDs to each of them. This allows you to separate traffic between VMs and to deal with VMs in the same way as physical hosts, that is, by zoning fabric or partitioning storage.

For example, if NPIV is not used, then every virtual machine shares one HBA with one WWN. This means that you are not able to separate traffic between these machines and to isolate LUNs for these machines because all of them use the same ID. In contrast, when NPIV is used, every VM has its own port ID and these port IDs are treated as N_Ports by FC fabric, so you can perform storage partitioning or zoning based on the VM’s port ID. The switch that the virtualized HBAs are connected to must support NPIV as well; check the documentation that comes with the FC switch.

The QLogic Intelligent Pass-thru Module and the Brocade Access Gateway use the NPIV technique to present the blade’s port IDs as N_Ports to the external fabric switches, thus eliminating the need for E-Ports connections between the BladeCenter chassis and external switches. In this way, all 14 internal blades’ FC ports are multiplexed and distributed across up to six external FC links and presented to the external fabric as N_Ports.

This means that external switches connected to the IPM do not see IPM at all, they see only N_ports connected to the F_ports. So IPM can help to achieve a higher port count for better scalability without using Domain IDs, and avoid multivendor interoperability issues. However, these intelligent pass-thru modules cannot be directly attached to the storage system. They must be attached to an external NPIV-capable FC switch. Refer to the switch’s documentation about NPIV support.

The first and most important step in planning is to select the correct SAN module that can provide the required functionality together with seamless integration into the existing client’s storage infrastructure (see Table 5-1 on page 399). There are no strict rules to follow during integration planning. However, several considerations must be taken into account.

Almost all switches support FC-SW2 interoperability standards, which means that almost any switch can be integrated into existing fabric using interoperability mode. *Interoperability mode* is a special mode used for integration of different vendors’ FC fabrics into one, and only standards-based functionality is available in the interoperability mode, although advanced features of a storage fabric’s vendor might not be available. Check the compatibility matrixes for a list of supported and unsupported features in the interoperability mode. Table 5-1 on page 399 provides a high-level overview of standard and advanced functions available for particular BladeCenter SAN switches, and lists how these switches might be used for designing new storage networks or integrating with existing storage networks.

Note: Advanced (proprietary) FC connectivity features from different vendors might be incompatible with each other, although providing almost the same functionality. For example, both Brocade and Cisco support port aggregation. However, Brocade uses ISL trunking and Cisco uses PortChannels, and they are incompatible with each other.

Table 5-1 SAN Module feature comparison and interoperability

	Brocade SAN SM	QLogic FCSM	McDATA FCSM	Cisco FCSM	QLogic IPM	Brocade AG
Basic FC connectivity						
FC-SW2 interoperability	Yes	Yes	Yes	Yes	N/A	N/A
Zoning	Yes	Yes	Yes	Yes	N/A	N/A
Maximum number of Domain IDs	239	239	31	239	N/A	N/A
Advanced FC connectivity						
VSAN	No	No	No	Yes	N/A	N/A
Port Aggregation	Yes	No	No	Yes	N/A	N/A
Advanced fabric security	Yes	Yes	Yes	Yes	N/A	N/A
Interoperability (existing fabric)						
Brocade interoperability	Yes	No	No	No	Yes	Yes
QLogic interoperability	Yes	Yes	Yes	Yes	Yes	Yes
McDATA interoperability	No	No[a]	Yes	No	Yes	Yes
Cisco interoperability	No	No	No	Yes	Yes	Yes

a. Optional McDATA Mode Firmware Upgrade is available

Note the following points:

- Where Yes is listed in the Interoperability section of Table 5-1, it indicates that a feature is supported without any restrictions for existing fabric, but with restrictions for added fabric, and vice versa.
- Where No is listed, it does not necessarily mean that a feature is not supported. Instead, it means that severe restrictions apply to the existing fabric; that is, some functions of the existing fabric potentially have to be disabled (if used).

For example, if we integrate Brocade FCSM into QLogic fabric then we are not be able to use Brocade proprietary features such as ISL trunking; however, QLogic fabric does not lose functionality. Conversely, if we integrate QLogic fabric into existing Brocade fabric, then we should place all Brocade switches in interoperability mode thus losing Advanced Fabric Services functionality.

If you plan to integrate IBM BladeCenter into a Fibre Channel fabric that is not listed here, then QLogic could be a good choice (but remember, this is possible with interoperability mode only, so extended functionality will not be supported). A better way would be to use the QLogic Intelligent Pass-thru or Brocade Access Gateway. The use of an Optical Pass-thru Module could also be an option.

The rules for switch selection and interoperability can be summarized as follows:

- McDATA FCSM is basically used when IBM BladeCenter is integrated into existing McDATA fabric.
- QLogic FCSM is essentially used when IBM BladeCenter is integrated into existing QLogic fabric or when basic FC functionality is required; that is, with one IBM BladeCenter with a directly connected storage server.

- Brocade SAN SM is used when IBM BladeCenter is integrated into existing Brocade fabric or when advanced FC connectivity is required, such as several chassis connected to high performance storage systems.
- Cisco FCSM is used when IBM BladeCenter is integrated into an existing Cisco fabric.

If you plan to use advanced features such as ISL trunking, then you might need to acquire specific licenses for these features.

Note: We highly recommend using FC storage fabric from the same vendor to avoid possible operational, management, and troubleshooting issues.

If IBM BladeCenter is attached to a non-IBM storage system, then support for that configuration should be provided by the storage system's vendor. Even if non-IBM storage is listed on IBM ServerProven, it means that the configuration has been tested but IBM will not provide support for it, and support is provided by a third-party vendor. Refer to the vendor's compatibility information for supported configurations.

For IBM storage, refer to "interoperability matrix" documents at:

http://www.storage.ibm.com

For instance, the DS4000® interoperability matrix can be found at:

http://www.ibm.com/servers/storage/disk/ds4000/interop-matrix.html

In addition you can check IBM ServerProven for a list of supported configurations at:

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

Additional information can be found in following publications:

- *IBM BladeCenter 4Gb SAN Solution*, SG24-7313
- *Fibre Channel Switch Interoperability Guide - IBM BladeCenter*
- *Brocade SAN Switch Modules: Design, Deployment and Management (DDM) Guide - IBM BladeCenter*
- *DS3000 interoperability matrix*
- *DS4000 interoperability matrix*

5.4 High availability and redundancy

The IBM BladeCenter chassis has built-in redundancy. All Fibre Channel HBAs, SAS daughter cards, or iSCSI cards for blade servers are dual port, and all storage switches (FC, SAS, or Ethernet) can be installed in a pair into an IBM BladeCenter chassis to avoid single points of failure in the storage infrastructure.

All major vendors, including IBM, use dual controller storage systems to provide redundancy. A typical topology for integrating IBM BladeCenter into FC infrastructure is shown in Figure 5-1 on page 401.

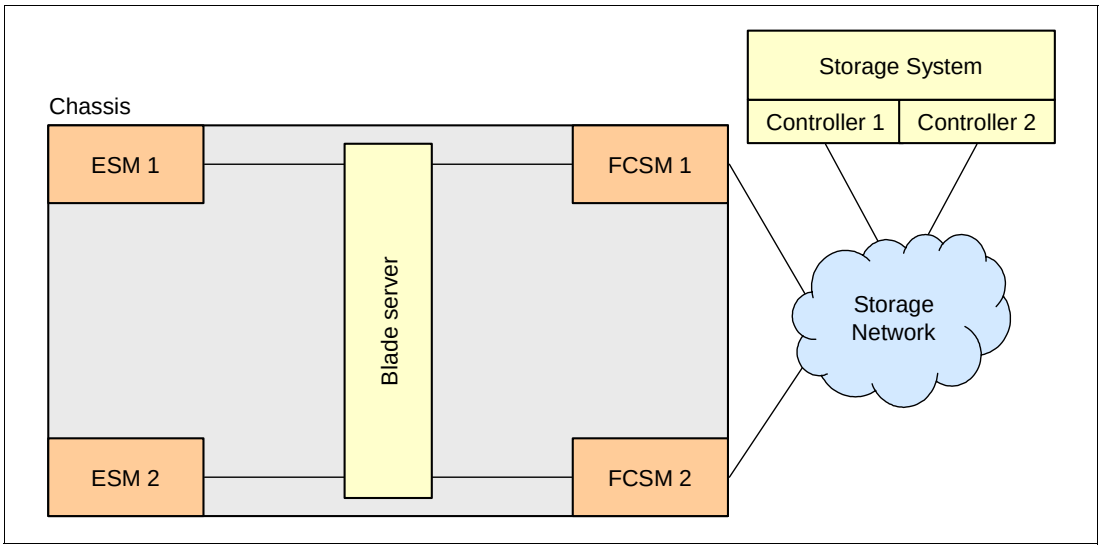

Figure 5-1 IBM BladeCenter SAN infrastructure topology

This topology includes a dual port FC expansion card installed into the blade's expansion slot, two FC switch modules installed into bays 3 and 4 of the IBM BladeCenter chassis, additional storage fabric infrastructure devices (if they exist), and storage systems or storage servers. The iSCSI SAN or SAS infrastructure uses a similar implementation concept.

In case of failure, the specific operating system driver provided by the storage system manufacturer is responsible for the automatic failover process (also known as multipathing capabilities). For instance, midrange IBM storage (the DS4000 family) uses RDAC for that purpose.

Therefore, if you plan to use redundancy and high availability for storage fabric, then ensure that failover drivers satisfy the following requirements:

- They are available from the vendor of the storage system.
- They come with the system or can be ordered separately (remember to order them in such cases).
- They support the blade's operating system.
- They support the redundant multipath fabric that you plan to implement (that is, they support the required number of redundant paths). For example, RDAC for IBM storage supports only four end-to-end redundant paths per storage system.

Refer to the storage system vendor's documentation for additional information.

5.5 Performance

Performance is an important consideration during storage infrastructure planning. Providing the required end-to-end performance for your SAN can be accomplished in several ways.

First, the storage system's failover driver can provide the functionality of load balancing across redundant paths in addition to high availability. IBM System Storage Multi-path Subsystem Device Driver (SDD) used with DS6000™ and DS8000® provides this

functionality. If you plan to use such drivers, then ensure that they satisfy the following requirements:

- They are available from the storage system vendor.
- They come with the system, or can be ordered separately.
- They support the blade's operating system.
- They support the multipath fabric that you plan to implement (that is, they support the required number of paths implemented).

Also, you can use static LUN distribution between two storage controllers in the storage system; that is, some LUNs will be served by controller 1 and others will be served by controller 2. A zoning technique can also be used together with static LUN distribution if you have redundant connections between FC switches and the storage system's controllers.

And finally, you might have to implement trunking or PortChannels between FC or Ethernet switches to increase network bandwidth, thus increasing performance. Trunks in the FC network use the same concept as in standard Ethernet networks; that is, several physical links between switches are grouped into one logical link with increased bandwidth. This is typically done when IBM BladeCenter is integrated into existing advanced FC infrastructures. However, keep in mind that not all FC switches support trunking (only Brocade can do that). Also be aware that this is an optional feature that requires the purchase of an additional license.

Refer to the storage system vendor's documentation and to the switch vendor's documentation for additional information.

5.6 Backup solutions

Backup is one of the most important parts of the IBM BladeCenter solution. First, you need to decide which kind of tape backup solution will be implemented. The valid IBM BladeCenter possibilities are listed here (see 5.6.1, "Dedicated server for centralized LAN backup" on page 403 for details):

- Centralized LAN backup with dedicated backup server (blade in chassis) with SCSI-attached or SAS-attached tape autoloader or tape library
- Centralized LAN backup with dedicated backup server (blade in chassis) with FC-attached tape autoloader or tape library
- Centralized LAN backup with dedicated backup server (non-blade) with SCSI-attached or SAS-attached tape drive, autoloader, or tape library
- Centralized LAN backup with dedicated backup server (non-blade) with FC-attached tape autoloader or tape library
- LAN-free backup with FC-attached tape autoloader or library; see 5.6.2, "LAN-free backup for blades" on page 405

If you plan to use blade server as a dedicated backup server or LAN-free backup for blades, then you must use only certified tape autoloaders and tape libraries (see IBM ServerProven for list of supported tape devices for specific blades). If you plan to use a dedicated backup server on a non-blade system, then you should use tape devices that are certified for that server. Also verify that the tape device and kind of backup you select are supported by the backup software you plan to use.

The following sections explain the valid IBM BladeCenter backup possibilities in more detail.

5.6.1 Dedicated server for centralized LAN backup

The simplest way to provide backup for blades is to use a blade or non-blade server with an SCSI-attached, SAS-attached, or FC-attached tape unit. In this case, all blade servers requiring backup will have backup agents, and backup traffic from these agents to the backup server will use standard LAN paths.

If you use an FC-attached tape drive, then it should be connected to FC fabric (or at least to HBA) that is dedicated for backup; it should *not* be connected to the FC fabric carrying disk traffic. If you cannot use dedicated switches, then use zoning techniques on FC switches to separate these two fabrics.

Note: We strongly recommend that you do not mix disk storage and tape storage on the same FC HBA. If you experience issues with your SAN and the root of these issues is found to be having the tape and disk on the same HBA, IBM will recommend separating these devices.

So, if you plan to use the blade server as a dedicated backup server with FC-attached tape, your choices are:

- Use either one port of the FC expansion card for tape and another for disk (but there is no redundancy in this case).
- Use a suitable expansion blade (such as the SIO, MIO, or MPE, depending on the blade servers you have) for additional FC ports to separate traffic and to provide redundancy.
- For the BladeCenter H and BladeCenter HT chassis, use an MSIM in conjunction with a suitable CFFh expansion card to increase the number of FC ports per blades, thereby gaining redundancy without losing blade density.

If you plan to use the blade server as a dedicated backup server with SCSI-attached tape, then you must use PCI Expansion Unit II (part number 25K8373), IBM Ultra320 SCSI Controller 2 (part number 39R8743), and tape connected to this controller. Refer to Table 3-92 on page 358 for a list of blades supporting the PCI Expansion Unit II (PEU2).

Possible topologies and traffic flows for LAN backups and FC-attached storage devices are shown on Figure 5-2 and Figure 5-3 on page 405.

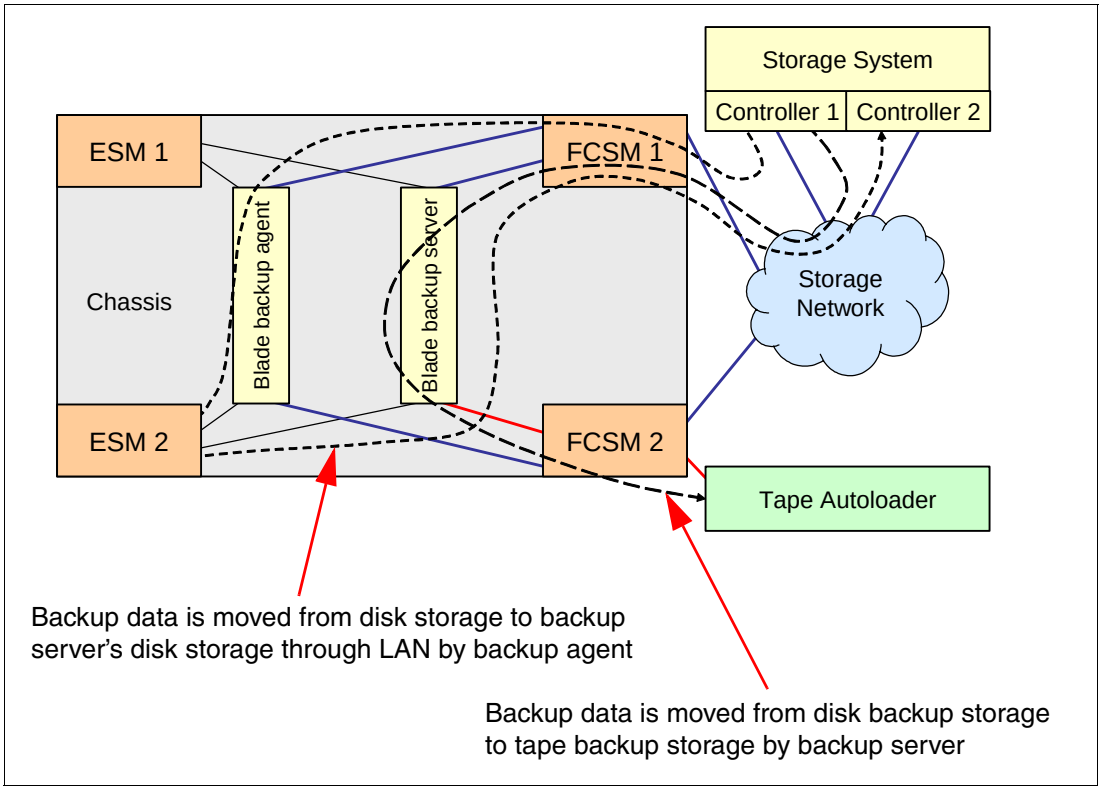

Figure 5-2 LAN backup topology and traffic flow - type 1

The topology shown in Figure 5-2 has the following characteristics. Each blade server participating in backup, except the backup server itself, has dual redundant connections to the disk storage system. The backup server has only one disk storage connection. The other port of the FC HBA is dedicated for tape storage. A backup agent is installed onto each blade requiring backup.

The backup traffic flow is as follows: the backup agent will transfer backup data from the disk storage to the backup server through LAN. The backup server will store this data on its disk storage, for instance, on the same storage system. Then the backup server will transfer data from its storage directly to the tape device. Zoning is implemented on a FCSM 2 to separate disk and tape data flows. Zoning is almost like VLANs in networks.

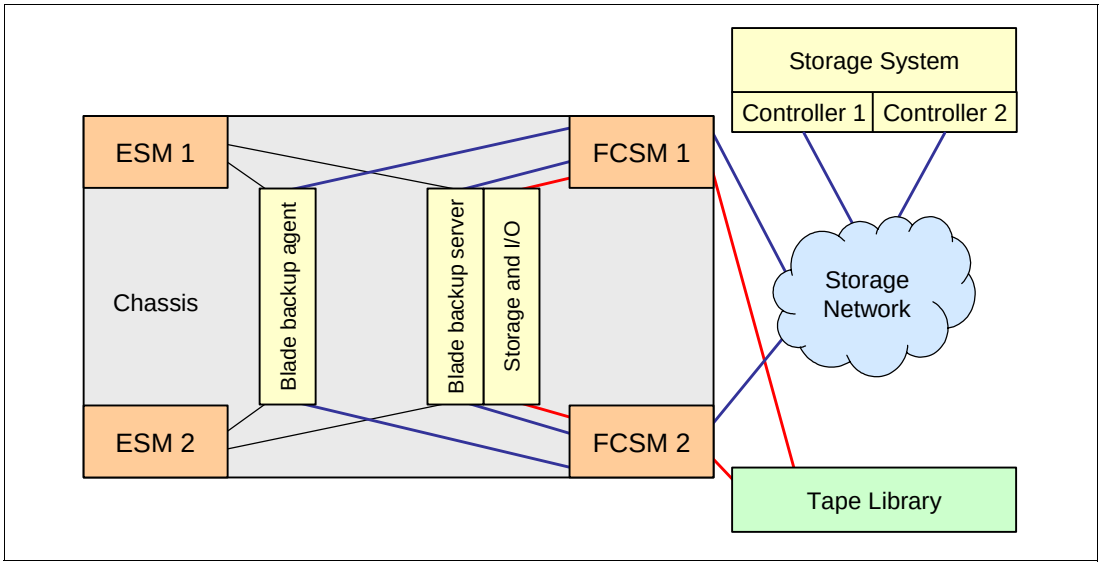

Figure 5-3 LAN backup topology - type 2

LAN backup topology shown on Figure 5-3 is almost the same as in Figure 5-2 on page 404 except that the backup server has a Storage and I/O (SIO) Expansion Unit, thus providing redundant connections both to the disk storage and the tape library. Having a tape library with at least two tape drives is highly recommended here because it can provide redundancy with two or more FC connections.

5.6.2 LAN-free backup for blades

LAN-free backup means that the SAN fabric is used for the backup data flow instead of LAN, and LAN is used only for passing control information between the backup server and agents. LAN-free backup can save network bandwidth for network applications, thus providing better network performance. The backup agent transfers backup data from the disk storage directly to the tape storage during LAN-free backup. This is illustrated in Figure 5-4 and Figure 5-5 on page 406.

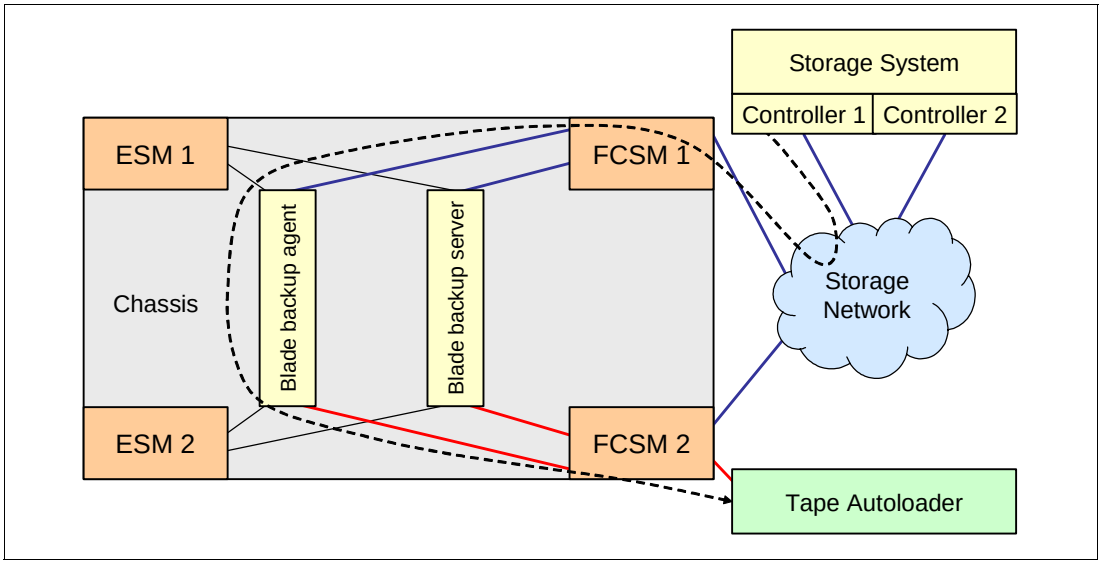

Figure 5-4 LAN-free backup without disk storage redundancy

Figure 5-4 on page 405 shows the simplest topology for LAN-free backup. With this topology, the backup server controls the backup process and the backup agent performs backup data movement from the disk storage directly to the tape storage. In this case there is no redundancy provided both for disk storage and tape storage. Zones are not required because the second FCSM is exclusively used for the backup fabric.

If you need redundancy you can implement a more complex topology, as shown in Figure 5-5.

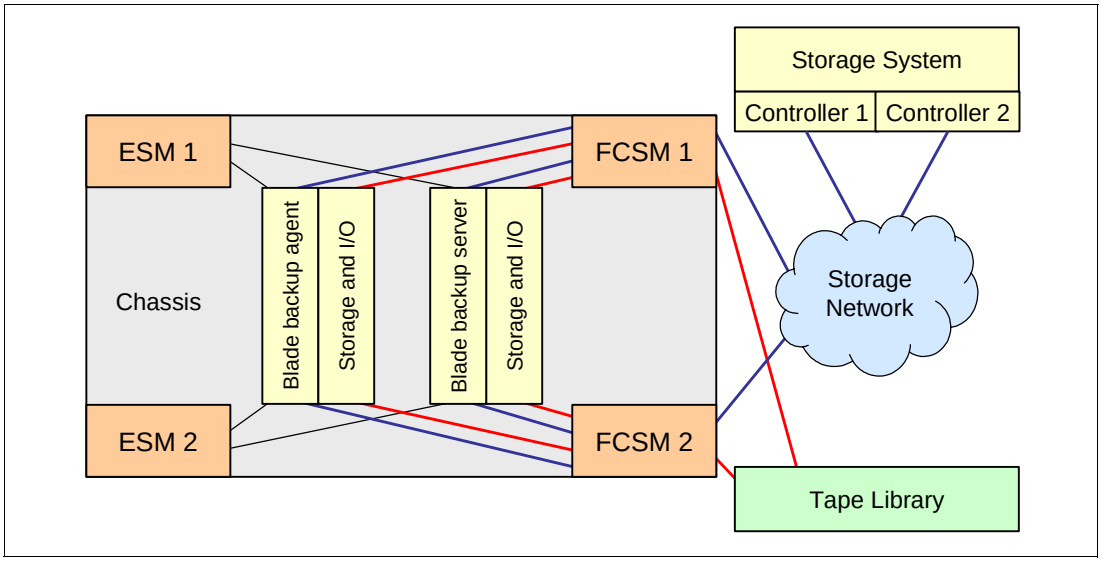

Figure 5-5 LAN-free backup with disk storage redundancy

This topology provides full redundancy both for disk storage and tape storage subsystems. But in this case, you have to use a Storage and I/O Expansion Unit (or another expansion unit providing additional slots for FC expansion cards) for each blade requiring backup and for the backup server, as well. The backup traffic flow is the same as in the previous topology. FC switch zoning is required to separate disk and tape data flows.

Backup software vendors can use other (or additional) topologies and protocols for backup operations. Consult the backup software vendor's documentation for a list of supported topologies and features, as well as additional information.

5.7 iSCSI solutions

iSCSI uses a traditional Ethernet network for block I/O between storage system and servers. That is, servers are connected to the LAN and storage systems are connected to the LAN, and they use SCSI over IP protocol (iSCSI) to communicate with each other. Because iSCSI uses a standard TCP/IP stack, you can use iSCSI connections across LAN, MAN, or even WAN connections.

IBM BladeCenter iSCSI solution components include the iSCSI initiator on the blade itself, traditional Ethernet switches, an iSCSI target such as IBM System Storage N series iSCSI models, an optional DHCP server, and a management station with iSCSI Configuration Manager.

In general iSCSI on blade servers can be implemented in two ways, by using a software iSCSI initiator or by using a hardware iSCSI initiator, as explained here:

- A software iSCSI initiator is specialized software that uses a server's CPU for iSCSI protocol processing. A software iSCSI initiator can exist on two different levels that complement each other:
  - As a driver in the operating system
  - As microcode that is built in to the onboard Ethernet controller on the blade

  Both implementations provide iSCSI capabilities for standard Ethernet NICs. However, an operating system driver can only be used after the locally installed operating system is booted and running. In contrast, the NIC built-in microcode is used for boot-from-SAN implementations, but cannot be used for storage access when the operating system is booted and running.

  Software initiators can be obtained from the operating system vendor (for example, Microsoft offers Software iSCSI initiator for download). Or they can be obtained as a part of an NIC firmware upgrade, as in blade servers.

  Table 5-2 lists blades that support a built-in software iSCSI initiator.

Table 5-2 IBM blade boot support with built-in software iSCSI initiator

Blade server (model)	iSCSI blade boot support[a]
HC10 (7996)	No
HS20 (8678)	No
HS20 (8832)	No
HS20 (8843)	Supported
HS20 (7981)	No
HS21 (8853)	Supported
HS21 XM (7995)	Supported
HS40 (8839)	No
LS20 (8850)	Supported
LS21 (7971)	Supported
LS22 (7901)	No
LS41 (7972)	Supported
LS42 (7902)	No
JS20 (8842)	No
JS21 (8844)	No
JS22 (7998)	No
QS21 (0792)	No
PN41 (3020)	No

a. BIOS, BMC, and NIC firmware upgrades might be required to support iSCSI boot with onboard Ethernet controllers.

- A hardware iSCSI initiator is a separate expansion card that provides processing of iSCSI protocol in dedicated hardware. IBM currently offers one hardware iSCSI initiator, the

QLogic iSCSI Expansion Card, part number (32R1923). It is possible to implement booting from iSCSI when using a hardware iSCSI initiator also.

The iSCSI BladeBoot solution from IBM is based on software iSCSI initiators integrated into a blade's NICs. It enables boot from SAN capabilities from external iSCSI storage systems as a standard blade feature without purchasing iSCSI-enabled expansion cards for blades.

Currently the iSCSI BladeBoot solution from IBM is supported with IBM System Storage N series iSCSI models and EMC AX150i. It includes the following operating systems:

- Microsoft Windows Server 2003
- SUSE Linux Enterprise 10
- Red Hat Enterprise Linux 5 (limited support - see Red Hat 5 Release Notes for restrictions)

Note: If you are planning to use Windows Server for iSCSI boot, then plan to have local disk storage for at least one blade in order to be able to install Windows onto a blade server, because Windows does not support installation to an iSCSI device directly. Then use your preferred disk cloning utility to create a disk image and put it onto external iSCSI storage.

For the latest compatibility information, consult the storage vendor's compatibility guides.

For more information about IBM BladeCenter iSCSI BladeBoot, go to:

`http://www.ibm.com/systems/bladecenter/storage/iscsi.html`

Note: We strongly recommend that you use a separate network segment for iSCSI traffic. That is, NICs, switches (or VLANs), and storage system ports participating in iSCSI communications should be isolated from other traffic.

If you plan for redundancy, you must use multipath drivers. Generally they are provided by the operating system vendor for iSCSI implementations, even if you plan to use hardware initiators.

It is possible to implement HA clustering solutions using iSCSI, but certain restrictions might apply (check the storage system vendor's compatibility guides). For instance, Microsoft Cluster Services are not supported with multipathing for IBM iSCSI solutions (DS300 and N3700).

When planning your iSCSI solution, consider the following:

- IBM blade hardware, the initiators, and the operating system should be supported by an iSCSI storage system; check the compatibility guides from the storage vendor for more information.
- Multipath drivers exist and are supported by the operating system and the storage system (when redundancy is planned); check the compatibility guides from the operating system vendor and storage vendor.

You can find additional information in the following publications:

- *IBM BladeCenter iSCSI SAN Solution*, REDP-4153
- *IBM System Storage N series Interoperability Matrix*
- *Microsoft Support for iSCSI* (from Microsoft)

5.8 Boot from SAN

Boot from SAN is a technique used when the blade server has no local disk drives and uses an external FC or iSCSI storage system LUN to boot the operating system. This means both the operating system and data reside in the SAN. This technique is commonly used to provide higher availability and better utilization of the blade's system storage (where the operating system resides). Hot Spare blades also can be implemented in a convenient way using boot from SAN.

You must ensure that the SAN boot configuration can be implemented as explained here:

- In general terms, all IBM blades are capable of using boot from SAN.
- Requirements listed in 5.2, "Fibre Channel storage system interoperability" on page 396 or 5.7, "iSCSI solutions" on page 406 should be met.
- The FC HBAs or iSCSI initiators used must support boot from SAN.
- The multipath drivers must support SAN boot.
- The operating systems must support boot from SAN; check the operating system documentation.
- The storage systems must support boot from SAN; check the compatibility guides from the storage system vendor.

The IBM System Storage Family compatibility guides contain FC SAN boot compatibility information for the DS3400 and DS4000 family (except the DS4200) with Microsoft Windows and SLES and RHEL Linux operating systems. For operating systems that are not listed in the compatibility guides, an RPQ is required in order for a boot from FC SAN configuration to be supported. RPQs are conducted through IBM System Storage.

You can also check the documentation for the operation system used for boot from SAN support. Additionally, several guides from IBM and other companies describe boot from SAN configurations for Linux and Novell NetWare operating systems.

Refer to the following links for additional SAN boot-related information:

- *IBM BladeCenter 4Gb SAN Solution*, SG24-7313
- *DS4000 interoperability matrix*
- *IBM BladeCenter Deployment Guide, Part 3 - Blades*
- *Boot from SAN: DS300/400/4x00 Configuration Guide* (from QLogic)
- *Boot from SAN in Windows Server 2003 and Windows 2000 Server* (from Microsoft)
- *ESX Server 3.x Storage/SAN Compatibility Guide* (from VMware)
- *Virtual Infrastructure 3: SAN Configuration Guide* (from VMware)

Note: There are special considerations to keep in mind if you plan to implement Microsoft Cluster Services with boot from SAN, which can be done with either boot disk and shared cluster disk using the same HBA port, or boot disk and shared cluster disk using separate HBAs. For additional information, refer to the Microsoft document *Boot from SAN in Windows Server 2003 and Windows 2000 Server*:

http://www.microsoft.com/windowsserversystem/wss2003/techinfo/plandeploy/bootfromsaninwindows.mspx

5.9 External SAS-based storage for IBM BladeCenter

SAS connectivity was added to the BladeCenter solutions to provide an affordable way of connecting external storage and tape devices to the IBM BladeCenter chassis. This section focuses on IBM System Storage DS3200-based solutions, as well as external SAS tape attachments:

- 5.9.1, “IBM BladeCenter Boot Disk System” on page 410
- 5.9.2, “IBM System Storage DS3200” on page 413
- 5.9.3, “SAS tape storage for IBM BladeCenter” on page 416
- 5.9.4, “EXP3000 attachment” on page 420

5.9.1 IBM BladeCenter Boot Disk System

The IBM BladeCenter Boot Disk System (1726-22B) is a part of the IBM System Storage DS3000 family. This solution enables clients to use all the advantages of diskless blades with a shared external disk system in a cost-effective manner both in terms of acquisition and cost of ownership.

The Boot Disk System is designed specifically to be a booting appliance; that is, to provide blade server boot capability from an external storage system. If you plan to work with application data stored on the external storage system as well (for example, a database application), then you might consider using an iSCSI or Fibre Channel solution instead of SAS for performance reasons. Alternatively, with BladeCenter H and MSIM, it is possible to use dedicated SAS fabric for booting blades, and add iSCSI or FC fabric for dedicated data storage infrastructure.

The IBM BladeCenter Boot Disk System is shown in Figure 5-6.

Figure 5-6 IBM BladeCenter Boot Disk System, 1726-22B

Note: All functional characteristics and supported configurations described here are based on the latest version of Boot Disk System controller firmware available at the time of writing (version 06.30.01.00 dated October 9, 2007).

The IBM BladeCenter Boot Disk System has the following characteristics:

- 2U rack-mountable enclosure.
- Two RAID controllers.
- RAID levels 0, 1, 10, 3, and 5 are supported.
- Up to two SAS host interfaces per controller.
- One Ethernet port per controller for out-of-band management.
- Up to 12 hot swap 3.5” HDDs with 3 Gbps SAS interface.
- Can be attached to up to 28 blade servers, 28 activated partitions ship standard.

- Two hot swap redundant power supplies with two fans each.
- DS3000 Storage Manager application support.
- Optional FlashCopy® and VolumeCopy options are available.

Note: The BladeCenter Boot Disk System does not support EXP3000 expansion units. Only SAS disk drives are supported.

The dual-controller Boot Disk System comes with one SAS host connector on each controller, enabling up to 14 blade servers in one BladeCenter chassis to be redundantly connected. To redundantly connect a second BladeCenter chassis with up to 14 blade servers to the boot disk system, you can install an optional SAS host port adapter (DS3200 SAS 2-port Daughter Card, part number 39R6509, feature code 4836) in each controller, which provides an additional SAS host connector.

Note: Only one of the two SAS host connectors on the optional SAS host port adapter can be connected to the IBM BladeCenter SAS Connectivity Module. The other SAS host port on the SAS host port adapter must remain unused.

At the time of writing, the IBM BladeCenter Boot Disk System supports attachment to the following chassis:

- IBM BladeCenter E
- IBM BladeCenter H

The following operating systems are supported for host servers that have mapped LUNs on the boot disk system:

- Microsoft
  - Windows Server 2003/2003 R2 x86 Web, Standard, Enterprise, and Datacenter
  - Windows Server 2003/2003 R2 x64 Standard, Enterprise, and Datacenter editions
  - Windows Storage Server 2003/2003 R2 Standard edition
  - Windows Small Business Server 2003/2003 R2 Premium edition
- Red Hat
  - Red Hat Enterprise Linux 4 for x86
  - Red Hat Enterprise Linux 4 for AMD64/EM64T
  - Red Hat Enterprise Linux 5 for System i/p (for JS21, QS21, and QS22)
- SUSE Linux
  - SUSE Linux Enterprise Server 9 for x86
  - SUSE Linux Enterprise Server 9 for AMD64/EM64T
  - SUSE Linux Enterprise Server 10 for x86
  - SUSE Linux Enterprise Server 10 for AMD64/EM64T
- IBM
  - AIX 5.3 (for JS22 only)
  - AIX 6.1 (for JS22 only)

Blade servers supported by the Boot Disk System include:

- HS20 (8843)
- HS21 (8853)
- HS21 XM (7995)
- LS20 (8850)
- LS21 (7971)
- LS41 (7972)

- JS21 (8844)
- JS22 (7998)
- QS21 (0792)
- QS22 (0793)

Note: Not all blades may support the operating systems listed here. Check the operating system compatibility information for specific blade servers.

For the latest information about host operating system support, refer to the latest DS3000 Storage Manager software README file and to the boot disk system interoperability matrix. The interoperability matrix is part of the *BladeCenter Interoperability Guide* which can be downloaded from:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5073016

IBM BladeCenter SAS boot solution components include:

- An IBM BladeCenter E or IBM BladeCenter H chassis; a maximum of two chassis is supported by the Boot Disk System.
- Supported blade servers with SAS expansion cards (part number 39Y9190), with one installed in each blade. A maximum of 28 blades can be attached to the Boot Disk System (14 per each chassis).
- SAS Connectivity Modules (part number 39Y9195), with two per chassis. A maximum of four SAS Connectivity Modules can be connected to the Boot Disk System.
- 1m or 3m SAS cables (part number 39R6529 or 39R6531), two per each chassis. A maximum of four SAS cables can be used.
- An IBM BladeCenter Boot Disk System (1726-22B) with optional features if required.

Figure 5-7 shows the connection topology for the Boot Disk System and two BladeCenter chassis.

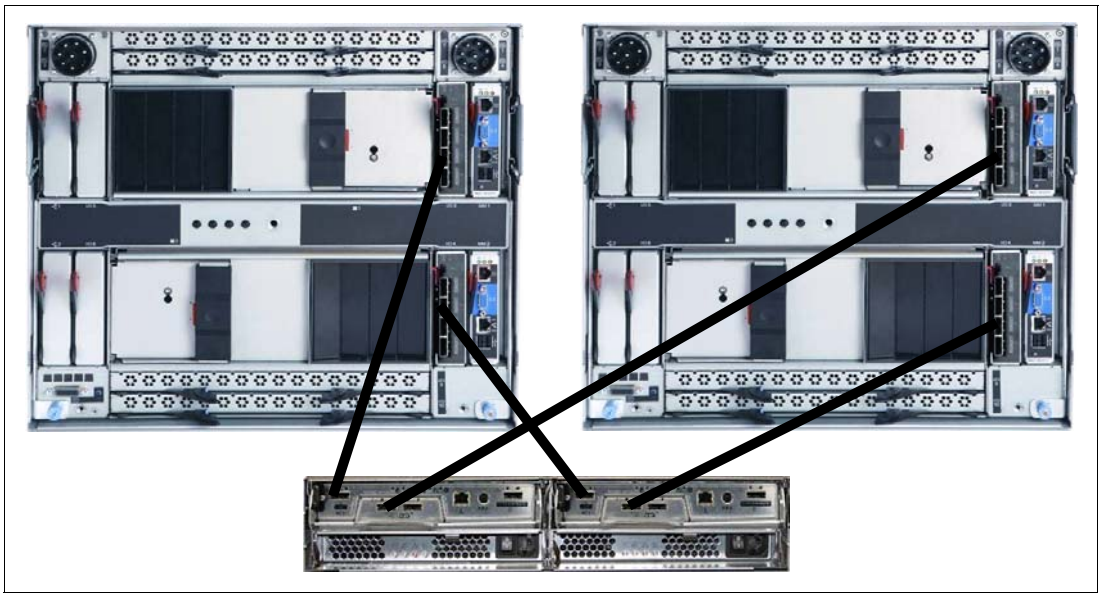

Figure 5-7 Sample BladeCenter SAS boot solution connection diagram for two BladeCenter chassis

Refer to *IBM BladeCenter Boot Disk System Installation and User Guide* for additional information:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5072756

For information about implementation, see the IBM Advanced Technical Support whitepaper *Implementing IBM BladeCenter Boot Disk System Solution* available from:

http://www.ibm.com/support/docview.wss?uid=tss1wp101198

5.9.2 IBM System Storage DS3200

The IBM System Storage DS3200 (1726-22X) is a part of the IBM System Storage DS3000 family. This solution enables clients to use all the advantages of shared data storage in a cost-effective manner, both in terms of acquisition and cost of ownership.

The IBM System Storage DS3200 is shown in Figure 5-6 on page 410.

Figure 5-8 IBM System Storage DS3200

Note: All functional characteristics and supported configurations covered here are based on the latest version of DS3000 controller firmware available at the time of writing (version 07.35.41.00 dated November 2, 2008).

The DS3200 has the following characteristics:

- 2U rack-mountable enclosure.
- One or two RAID controllers.
- RAID levels 0, 1, 10, 3, 5, and 6 are supported.
  - A maximum of 48 drives is supported in a single RAID 0, 1, and 10 arrays.
  - A maximum of 30 drives is supported in a single RAID array for other RAID levels (3, 5, and 6).
  - Logical volume size can be greater than 2 TB.
- Up to two SAS host interfaces per controller.
- One Ethernet port per controller for out-of-band management.
- Up to 48 hot swap 3.5” HDDs with SAS or SATA interface with up to three EXP3000 expansion units. Intermix of SAS and SATA drives is supported; however, both SAS and SATA drives cannot be used in the same array.
- Can be attached to up to 14 blade servers (one chassis). Partition expansion license is required to support more than four host server blades.
- Two hot swap redundant power supplies with two fans each.

- DS3000 Storage Manager application support.
- Optional FlashCopy and VolumeCopy options are available.

Note: Currently, DS3200 is supported for both data storage and boot from SAN when attached to the supported BladeCenter chassis.

At the time of writing, the IBM DS3200 supports attachment to the following chassis:

- IBM BladeCenter E
- IBM BladeCenter H
- IBM BladeCenter T
- IBM BladeCenter HT

The following operating systems are supported for blade servers that have LUNs on DS3200:

- Microsoft
  - Windows Server 2008 x86 Standard, Enterprise, and DataCenter editions
  - Windows Server 2008 x64 Standard, Enterprise, and DataCenter editions
  - Windows Server 2003 Service Pack 2/2003 R2 x86 Web, Standard, Enterprise, and DataCenter editions
  - Windows Server 2003 Service Pack 2/2003 R2 x64 Standard, Enterprise, and DataCenter editions
  - Windows Small Business Server 2003 Service Pack 2/2003 R2 Premium Edition
  - Windows Storage Server 2003 Service Pack 2/2003 R2 Standard Edition
- Red Hat
  - Red Hat Enterprise Linux 4 Update 7 for x86
  - Red Hat Enterprise Linux 4 Update 7 for AMD64/EM64T
  - Red Hat Enterprise Linux 5 Update 2 for x86
  - Red Hat Enterprise Linux 5 Update 2 for AMD64/EM64T
  - Red Hat Enterprise Linux 5 Update 2 for System i and System p (JSxx and QSxx blades only)
- SUSE Linux
  - SUSE Linux Enterprise Server 9 Service Pack 4 for x86
  - SUSE Linux Enterprise Server 9 Service Pack 4 for AMD64/EM64T
  - SUSE Linux Enterprise Server 10 Service Pack 2 for x86
  - SUSE Linux Enterprise Server 10 Service Pack 2 for AMD64/EM64T
  - SUSE Linux Enterprise Server 10 Service Pack 2 for POWER (JSxx blades only)
- IBM
  - AIX 5.3
  - AIX 6.1
- Sun
  - Solaris 10 Update 5 on x86

Blade servers supported:

- HS12 (8014 and 8028)
- HS20 (7981)
- HS20 (8843)
- HS21 (8853)
- HS21 XM (7995)
- LS20 (8850)
- LS21 (7971)

- LS22 (7901)
- LS41 (7972)
- LS42 (7902)
- JS12 (7998-60X)
- JS21 (8844)
- JS22 (7998)
- QS21 (0792)
- QS22 (0793)

Note: Not all blades may support the operating systems listed here. Check the operating system compatibility information for specific blade servers.

For the latest information about host operating-system support, refer to the latest DS3000 Storage Manager software README file and also refer to the DS3000 interoperability matrix. The interoperability matrix can be downloaded from:

http://www-03.ibm.com/systems/storage/disk/ds3000/pdf/interop.pdf

DS3200 SAS storage solution components include:

- A supported BladeCenter chassis; a maximum of one chassis is supported by DS3200.
- Supported blade servers with SAS Expansion Card (part number 39Y9190) or SAS Connectivity Card (43W3974), one per each blade. A maximum of 14 blades can be attached to the single DS3200 (14 in one chassis).
- SAS Connectivity Modules (part number 39Y9195); two per chassis is required.
- 1m or 3m SAS cables (part number 39R6529 or 39R6531); two per chassis is required.
- An IBM System Storage DS3200 Dual Controller model (1726-22X) with optional features if required.

Note: SAS Connectivity Card (43W3974) is supported with HS12 8028 models only. HS12 8014 and HS21 XM are not supported with this card; the SAS Expansion Card (39Y9190) must be used instead.

Sample connection topology for BladeCenter chassis and DS3200 is shown in Figure 5-9 on page 416.

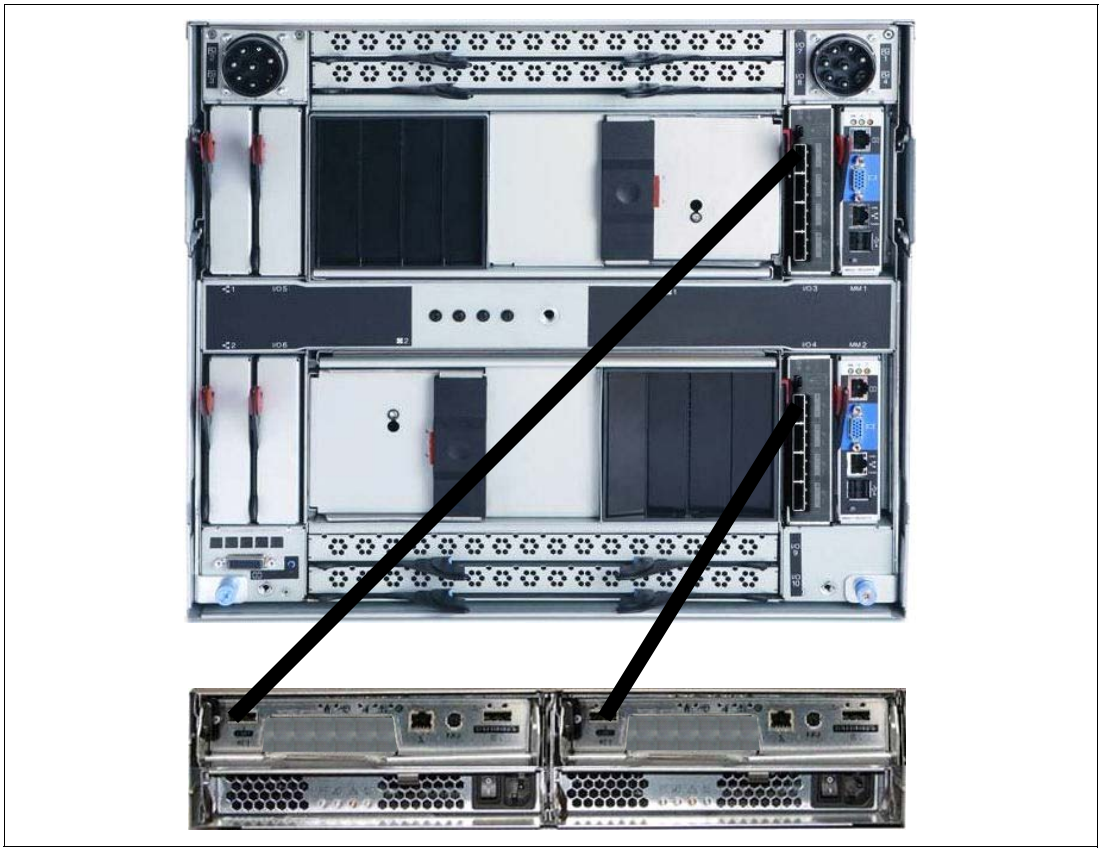

Figure 5-9 Sample connection topology - IBM BladeCenter H and DS3200

Notes:

- The DS3200 single-controller version is not supported.
- Each SAS connectivity module must be attached to one of the controllers in the DS3200, but not to both. That is, having redundant paths from a single SAS Connectivity Module to both DS3200 controllers is not supported.
- Optional DS3200 SAS 2-port Daughter Card (39R6509, feature code 4836) is not supported in this configuration.
- The DS3200 and the Boot Disk System are not supported together when connected to the same chassis.

Refer to *IBM System Storage DS3200 Installation, User, and Maintenance Guide* for additional information:

`http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5069915`

5.9.3 SAS tape storage for IBM BladeCenter

Tape backups are an important part of any solution. With SAS technology, IBM BladeCenter offers an affordable and convenient way to perform backups by attaching external tape drives and autoloaders directly to the BladeCenter chassis.

BladeCenter SAS backup solution components include:

- Chassis
- Blade servers
- SAS Expansion Card (39Y9190, feature code 2979)
- SAS Connectivity Module (39Y9195, feature code 2980)
- Tape backup unit
- SAS cables
- Backup software

Note: SAS Expansion Card (39Y9190, feature code 2979) is the only supported card for SAS-based tape solutions for IBM BladeCenter.

SAS tape backup units supported include devices from both the IBM System Storage division (TS2230, TS2240, TS2340, TS2900, TS3100, TS3200, and TS3310) and the System x division (HH LTO3 and HH LTO4):

- IBM Half High LTO Generation 3 SAS Tape Drive, 43W8478
- IBM Half High LTO Generation 4 SAS Tape Drive, 44E8895
- IBM System Storage TS2230 Half-High LTO 3 SAS Tape Drive, 3580S3E
- IBM System Storage TS2240 Half-High LTO 4 SAS Tape Drive, 3580S4E
- IBM System Storage TS2340 Full-High LTO 4 SAS Tape Drive, 3580S4X
- IBM System Storage TS2900 Tape Autoloader with:
  - Half-high LTO 4 SAS Tape Drive, 3572S4R
  - Half-high LTO 3 SAS Tape Drive, 3572S3R
- IBM System Storage TS3100 Tape Library with:
  - Ultrium 3 SAS Tape, 3573S32
  - Ultrium 4 SAS Tape Drive, 3573S42, 3573E42, or 3573S4S
- IBM System Storage TS3200 Tape Library with:
  - Ultrium 3 SAS Tape, part numbers 3573S34, and
  - Ultrium 4 SAS Tape Drive, part numbers 3573S44, 3573E44, or 3573S4H
- IBM System Storage TS3310 Tape Library, model 3576, with LTO 4 SAS tape drives.

Note: The TS3310 Tape Library is not available through the System x channel. It should be ordered through the System Storage channel.

The HH LTO3 and HH LTO 4 are supported in external tape enclosures only.

The TS3200 Tape Library with Full High Ultrium 4 SAS Tape (3573S4H) and TS3310 (3576) are the only devices from the preceding list that support path failover for x86 blades with Windows/Linux operating systems only. An additional license (feature code 1682) is required to support this feature. The feature 1682 is not available through the System x channel; it must be ordered through the System Storage channel if required.

Refer to Table 5-3 on page 418 for BladeCenter SAS tape backup solutions compatibility information (chassis, blades, tape units, cables, and network operating systems).

Table 5-3 BladeCenter SAS tape-based backup solutions compatibility

Name	Part number	HH LTO3	HH LTO4	TS 2230	TS 2240	TS 2340	TS 2900	TS3100/ TS3200	TS 3310
External tape enclosures									
Half-high tabletop drive enclosure	8767	Yes	Yes	No	No	No	No	No	No
1 U rack mount drive enclosure	8765	Yes	Yes	No	No	No	No	No	No
4 U rack mount drive enclosure	8766	Yes	Yes	No	No	No	No	No	No
Cables									
2.0 m SAS/Mini-SAS 1x Cable	95P4711	No	No	No	No	No	No	No	No
5.5 m SAS/Mini-SAS 1x Cable	95P4712	No	No	No	No	No	No	No	No
2.0 m Mini-SAS/Mini-SAS 1x cable	95P4713	No	No	Yes	Yes	Yes	Yes	Yes	Yes
5.5 m Mini-SAS/Mini-SAS 1x cable	95P4714	No	No	Yes	Yes	Yes	Yes	Yes	Yes
SAS cable (3 m external)[a]	None	Yes	Yes	No	No	No	No	No	No
Chassis									
BladeCenter S	8886	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
BladeCenter E	8677	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
BladeCenter H	8852	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
BladeCenter T	8720	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
BladeCenter T	8730	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
BladeCenter HT	8740	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
BladeCenter HT	8750	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Blade servers									
HS12	8014	No	No	No	No	No	No	No	No
HS12	8028	No	No	No	No	No	No	No	No
HS20	8678	No	No	No	No	No	No	No	No
HS20	8832	No	No	No	No	No	No	No	No
HS20	8843	Yes	Yes	No	No	No	Yes	Yes	Yes
HS20	7981	No	No	No	No	No	No	No	No
HS40	8839	No	No	No	No	No	No	No	No
HS21	8853	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
HS21 XM	7995	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
LS20	8850	No	No	No	No	No	No	Yes	Yes
LS21	7971	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
LS22	7901	No	No	Yes	Yes	No	Yes	Yes	Yes
LS41	7972	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes

Name	Part number	HH LTO3	HH LTO4	TS 2230	TS 2240	TS 2340	TS 2900	TS3100/ TS3200	TS 3310
LS42	7902	No	No	Yes	Yes	No	Yes	Yes	Yes
HC10	7996	No	No	No	No	No	No	No	No
QS20	0200	No	No	No	No	No	No	No	No
QS21	0792	No	No	No	No	No	No	No	No
QS22	0793	No	No	No	No	No	No	No	No
JS12	7998-60X	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
JS20	8842	No	No	No	No	No	No	No	No
JS21	7988	Yes	Yes	No	No	No	No	No	No
JS22	7998	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
Network operating systems									
Microsoft									
Windows Server 2003/2003 R2 x86	None	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Windows Server 2003/2003 R2 x64	None	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Windows Server 2008 x86	None	No	No	No	No	No	Yes	No	No
Windows Server 2008 x64[b]	None	No	No	Yes	Yes	No	Yes	Yes	Yes
Novell									
SLES 9 for x86	None	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SLES 9 for AMD64/EM64T	None	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SLES 9 for IBM POWER[c]	None	No	No	Yes	Yes	Yes	No	Yes	Yes
SLES 10 for x86	None	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SLES 10 for AMD64/EM64T	None	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SLES 10 for IBM POWER	None	No	No	Yes	Yes	Yes	No	Yes	Yes
Red Hat									
RHEL 4 for x86	None	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 4 for AMD64/EM64T	None	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 4 for i/p Series	None	No	No	Yes	Yes	Yes	No	Yes	Yes
RHEL 5 for x86	None	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 5 for AMD64/EM64T	None	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 5 for System i/p	None	No	No	Yes	Yes	Yes	No	Yes	Yes
IBM									
AIX 5.2	None	Yes	Yes	No	No	No	No	No	No
AIX 5.3	None	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
AIX 6.1	None	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes

a. This cable is included with SAS Enclosure Adapter Kit Option, part number 40K2599. This cable cannot be ordered through the System x channel. It is orderable through the service channel only, FRU number 43W8496.
b. If used with TS2230, TS2240, TS3100/TS3200, and TS3310. then it is supported by LS22 (7901) and LS42 (7902) only.
c. Supported by JS12 only.

Note: Not all blades may support the operating systems listed in Table 5-3 on page 418. Check the operating system compatibility information for specific blade servers.

You must take several considerations into account in order to successfully implement a backup solution based on SAS tape devices, as explained here:

- The solution components must be compatible with each other (see Table 5-3 on page 418).
- The backup software used must support the operating environment proposed, including connectivity topologies, tape devices, and operating systems.
- Even if several servers are able to access the same tape device simultaneously because of the SAS architecture, the tape device is unable to handle this situation, so it must be avoided.
- Do not mix both tape devices and disk storage devices on the same HBA port because this may lead to significant performance degradation of backup operations.
- If the SAS connectivity module is used for both disk and tape operations for different servers, then plan to implement zoning to isolate different types of traffic going through.
- LAN-free drive sharing is not supported for SAS-based tape libraries covered in this section. Plan to use dedicated backup servers and LAN-based backup operations.

For more information about tape drive support please refer to *IBM System Storage Products Support - Tape systems* web link:

http://www.ibm.com/systems/support/supportsite.wss/allproducts?brandind=5000034

5.9.4 EXP3000 attachment

With the ServeRAID MR10ie (CIOv) RAID expansion card for BladeCenter servers, it is possible to attach external EXP3000 enclosures to the supported BladeCenter chassis using SAS Connectivity Modules. This provides advanced local storage capability for blades, including increased disk space, battery-backed cache, and a full range of RAID levels including RAID 0, 1, 10, 5, 50, 6, and 60.

IBM System Storage EXP3000, part number 172701X, is an external expansion enclosure used together with supported ServeRAID controllers or DS3000 storage servers to increase the number of drives and disk space available for attached systems. EXP3000 is shown on Figure 5-10 on page 421.

Figure 5-10 IBM System Storage EXP3000

IBM System Storage EXP3000 has the following characteristics:

- Two external x4 3 Gb SAS ports on a single ESM. An additional ESM option is available separately to add a second pair of SAS ports to EXP3000. Currently, only one SAS port is supported to connect EXP3000 to the SAS Connectivity Module.
- Up to 12 hard disk drives.
- Support for SAS and SATA HDDs including intermix.
- Dual redundant power supplies.

The solution components include:

- Supported BladeCenter chassis.
- ServeRAID MR10ie (CIOv) Controller installed into a supported blade server.
- Up to two SAS Connectivity Modules installed into the chassis.
- EXP3000 attached to the external port of the SAS Connectivity Module.

Supported chassis:

- BladeCenter E
- BladeCenter H

Supported blade server:

- HS12 (8028)

Note: The SAS Connectivity Card (43W3974) and the ServeRAID-MR10ie card are both required for the HS12 to connect to the EXP3000 enclosure.

A sample solution is shown in Figure 5-11 on page 422.

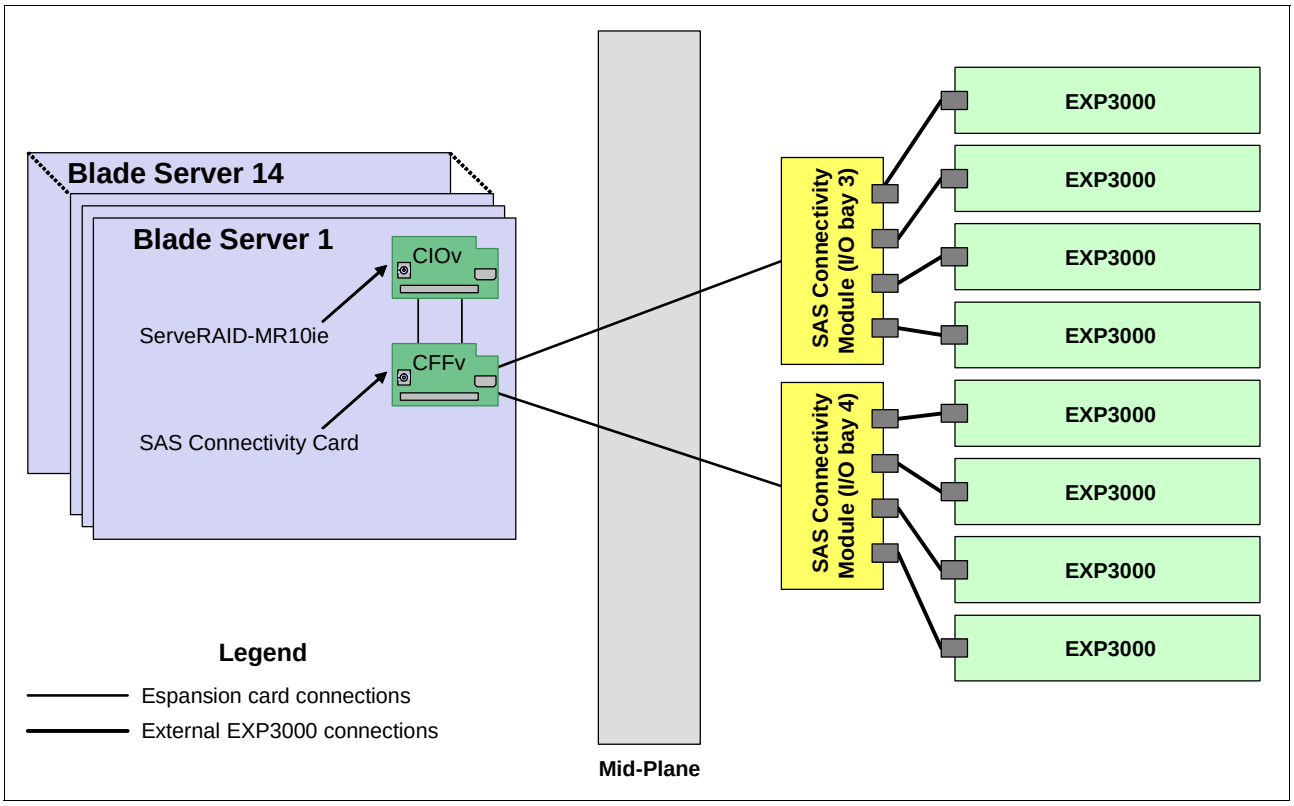

Figure 5-11 Solution topology for EXP3000 attachments

You must take several considerations into account when planning EXP3000 attachments to the BladeCenter:

- Every blade requiring integrated storage connectivity must have one ServeRAID-MR10ie Controller (46C7167) installed.
- At least one SAS Connectivity Module (part number 39Y9195) must be installed into the BladeCenter S chassis. A maximum of two SAS modules are supported.
- IBM System Storage EXP3000 units (up to eight EXP3000s - one per blade) can be attached to the BladeCenter chassis to increase the number of drives per blade.
- A maximum of one EXP3000 is supported per single blade server. Mixing of hard disks in the same array from different EXP3000s is not supported.
- There is no redundancy between SAS Connectivity Modules. Each of them supports their own set of attached EXP3000s.
- All disks in a single EXP3000 can only belong to the single blade server, and the blade itself sees these physical disks as its own “local” hard disk drives connected to the RAID controller.
- The following cables can be used to attach EXP3000 to the SAS Connectivity Module:
  - IBM 3m SAS Cable, 39R6531, feature code 3707
  - IBM 1m SAS Cable, 39R6529, feature code 3708
- The SAS Connectivity Module controls disk assignments by using zoning-like techniques. This results in the SAS module maintaining single, isolated dedicated paths between physical HDDs and the blade server. Configuration of HDD assignments is done by the administrator, and the administrator uses predefined templates or creates custom

configurations. Custom zoning *must* be implemented to ensure that only one blade server has access to one EXP3000.

- RAID functionality is supplied by the ServeRAID-MR10ie controllers installed into blade servers. RAID levels supported by ServeRAID-MR10i are:
  - RAID-0
  - RAID-1
  - RAID-5
  - RAID-6
  - RAID-10
  - RAID-50
  - RAID-60
- The maximum number of drives supported by ServeRAID-MR10ie controller for this solution is 14 (two internal drives plus 12 drives in one EXP3000).

  The maximum number of arrays supported by ServeRAID-MR10ie is 128, and the maximum number of logical disks is 16 per array and 64 per controller.
- Mixing HDDs of different capacities in a single volume is supported. However, the total volume size is aligned with the size of the smallest HDD, and excess space on larger-sized HDDs is not used.
- Both SAS and SATA hot swap HDDs are supported, and intermixing SAS/SATA drives is supported as well. However, each volume must have hard disks of the same type; that is, SAS or SATA.

For more information refer to *ServeRAID-MR10ie Controller Installation and User Guide* at:

`http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5078830`

5.10 IBM BladeCenter S integrated storage

A key feature of the IBM BladeCenter S chassis is support for integrated storage. Currently, there are three ways to implement the integrated storage solution for BladeCenter S:

- 5.10.1, “Basic local storage using SAS Connectivity Module” on page 423
- 5.10.2, “Advanced local storage using ServeRAID MR10ie” on page 426
- 5.10.3, “Advanced shared storage using the SAS RAID Controller Module” on page 429

Note: Feature codes for I/O modules listed in this section are for use when ordering the IBM BladeCenter S chassis.

5.10.1 Basic local storage using SAS Connectivity Module

The main feature of basic local storage is the ability to assign physical disks in disk storage modules (DSMs) to the blade server, and create volumes by using the RAID function of either the onboard RAID controller on the blade itself in conjunction with the SAS Connectivity Card or the SAS Expansion Card installed into blade server.

Table 5-4 lists the basic local storage solution components for BladeCenter S.

Table 5-4 Basic local storage solution components for BladeCenter S

Component description	Part Number	Feature Code	Min/max quantity
Disk Storage Module (DSM)	43W3581	4545	1 / 2

Component description	Part Number	Feature Code	Min/max quantity
SAS Connectivity Module	39Y9195	3267	1 / 2
SAS Expansion Card[a]	39Y9190	2979	1 per blade server
SAS Connectivity Card[a]	43W3974	1591	1 per blade server

a. Either the SAS Connectivity Card or the SAS Expansion Card is required per blade server but not both.

For more information:

- For DSMs, refer to 2.13, “BladeCenter S storage modules” on page 211
- For SAS Connectivity Modules, refer to 2.5, “SAS I/O modules” on page 164
- For the SAS Expansion Card, refer to 3.1.17, “SAS Expansion Card” on page 252
- For the SAS Connectivity Card, refer to 3.1.18, “SAS Connectivity Card” on page 254

A sample connection topology for basic local storage with one SAS Connectivity Module installed is shown on Figure 5-12.

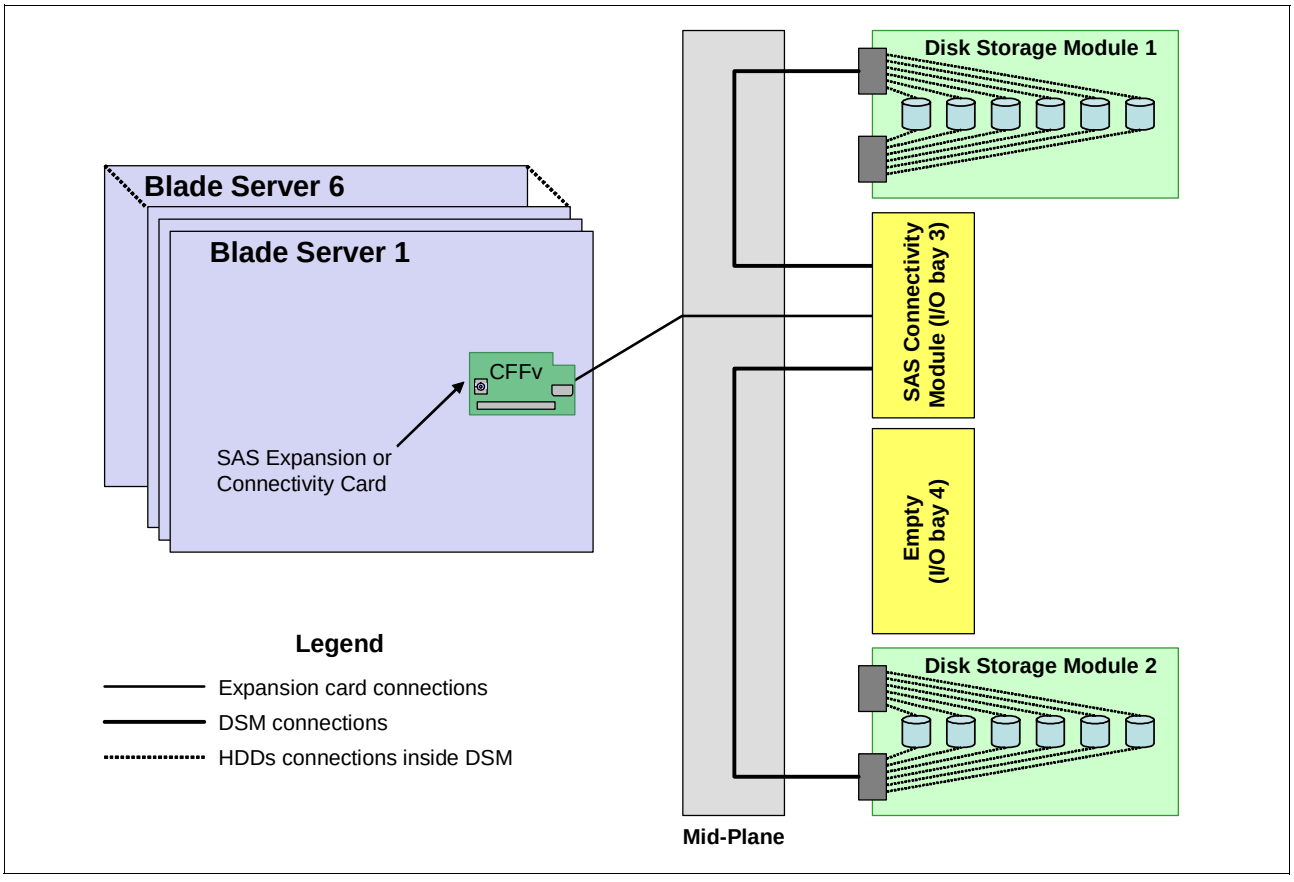

Figure 5-12 SAS I/O connections with one SAS Connectivity Module installed

Figure 5-13 on page 425 shows a sample connection topology for basic local storage with two SAS Connectivity Modules installed.

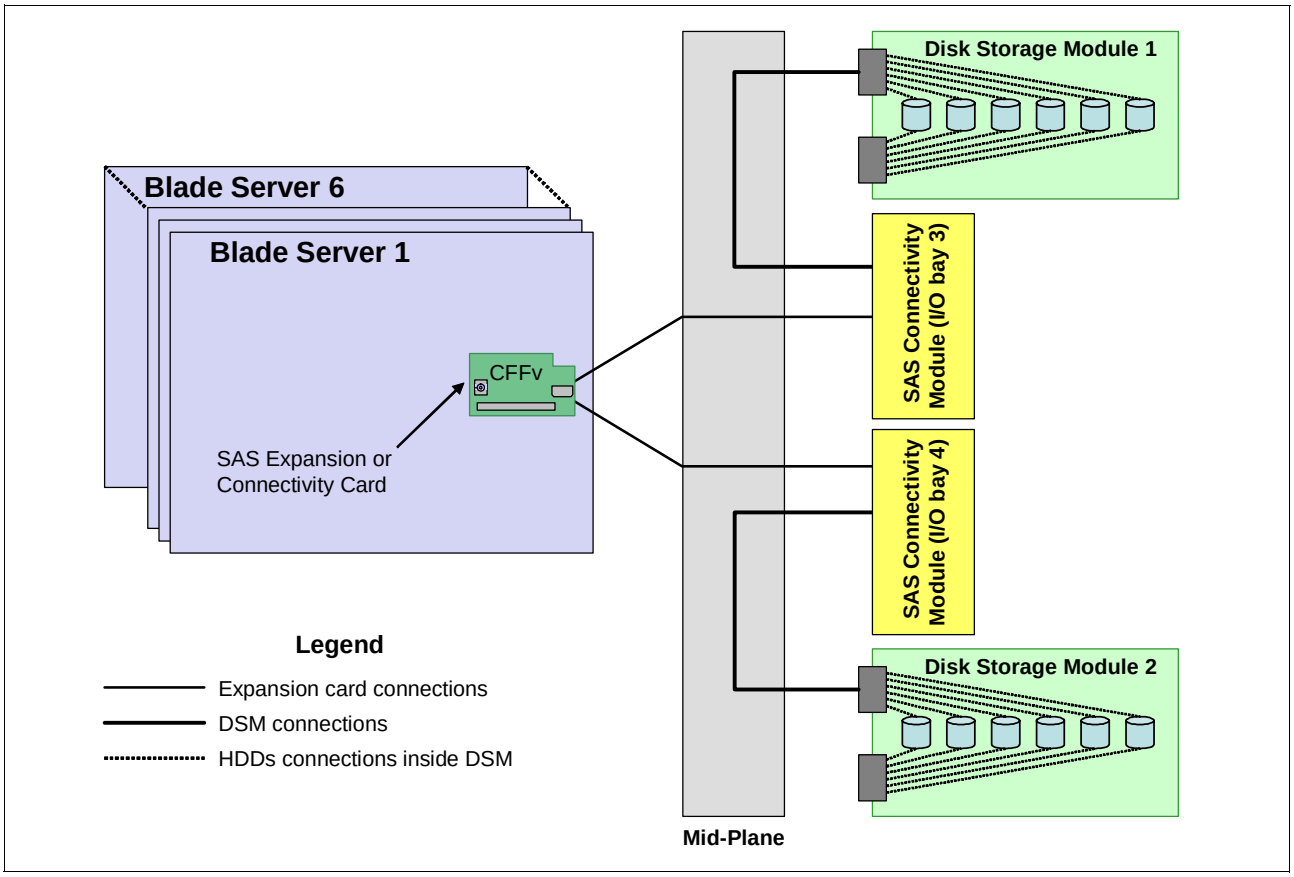

Figure 5-13 SAS I/O connections with two SAS Connectivity Modules installed

You must take several considerations into account when planning BladeCenter S basic local storage implementations:

- Every blade requiring integrated storage connectivity must have one SAS Expansion Card or SAS Connectivity Card installed.
- At least one DSM must be installed into the BladeCenter S chassis; a maximum of two DSMs are supported in one chassis. The IBM BladeCenter S does not ship with any DSMs as standard.
- At least one SAS Connectivity Module must be installed into the BladeCenter S chassis. A maximum of two SAS modules are supported for redundancy and high availability purposes.

  If two SAS connectivity modules are installed, the module in I/O module bay 3 controls access to storage module 1, and the module in I/O module bay 4 controls access to storage module 2.
- Each physical hard disk drive in any DSM can be assigned to the one blade server only, and the blade itself sees this physical disk (or disks) as its own hard disk drives connected via SAS expansion card. That is, there are no LUNs, storage partitions, and so forth as in shared storage systems. Instead, each blade has its own set of physical disks residing in DSM (or DSMs).
- The SAS Connectivity Module controls disk assignments by using zoning-like techniques, which results in the SAS module maintaining single, isolated dedicated paths between physical HDDs and the blade server. Configuration of HDD assignments is done by the

administrator, and the administrator uses predefined templates or creates custom configurations.

- The RAID functionality is supplied by the SAS expansion cards in the blade servers. RAID levels supported by the SAS Expansion Card or by the onboard RAID controller on the blade itself are:
  - RAID-0 (Integrated Striping - IS)
  - RAID-1 (Integrated Mirroring - IM)
  - RAID-1E (Integrated Mirroring Enhanced - IME)
- The maximum number of drives supported by the IM volume is two plus one optional global hot spare. The IME volume supports up to ten HDDs plus two optional hot spares. The IS volume supports up to ten HDDs. The IS volume does not support hot spare drives.
- When creating a RAID-1 array, we recommend that you span both disk storage modules. This maximizes the availability of data if one of the paths to the disks is lost, because there is only one connection to each disk storage module, as shown in Figure 5-13 on page 425.
- Mixing HDDs of different capacities in a single volume is supported. However, the total volume size is aligned with the size of the smallest HDD, and excess space on larger-sized HDDs is not used.

  Supported combinations of volumes include:
  - Two IM or IME volume per blade server
  - One IM or IME volume and IS volume per blade server
  - Two IS volumes per blade server
- Each blade with an SAS expansion card has access to its assigned HDDs in both DSMs, even if only one SAS Connectivity module is present. Potentially, all 12 drives in both DSMs can be assigned to the single blade server. However, only 10 HDDs can be used in a single volume. You can create either two volumes to utilize the capacity of all drives, or designate the remaining two drives as hot spares.
- Both SAS and SATA hot swap HDDs are supported, and an intermix of SAS and SATA drives is supported, as well. However, each volume must have hard disks of the same type; that is, SAS or SATA.
- External disk storage attachments are not supported.

5.10.2 Advanced local storage using ServeRAID MR10ie

The main feature of advanced local storage is the ability to assign physical drives installed into disk storage modules and external disk expansion enclosures to the blade servers, and create RAID arrays and logical drives by using an advanced RAID controller expansion card installed into blade servers. In other words, each blade has its own set of physical disks that are managed by a local RAID controller installed into the blade itself. Table 5-5 lists the advanced local storage solution components for BladeCenter S.

Table 5-5 Advanced local storage solution components for BladeCenter S

Component description	Part Number	Feature Code	Min/max quantity
Disk Storage Module (DSM)	43W3581	4545	1/2
SAS Connectivity Module	39Y9195	3267	1/2
ServeRAID-MR10ie expansion card	46C7167	5490	1 per blade server
SAS Connectivity Card	43W3974	1591	1 per HS12 blade server

Note: Currently, only HS12 (8028) is supported with this solution. The SAS Connectivity Card (43W3974) is required for HS12 to support attachments of EXP3000 to the ServeRAID-MR10ie (CIOv) Controller.

For more information:

- For DSMs, refer to 2.13, “BladeCenter S storage modules” on page 211
- For the SAS Connectivity Module, refer to 2.5, “SAS I/O modules” on page 164
- For the ServeRAID-MR10ie controller, refer to 3.1.19, “ServeRAID MR10ie (CIOv) Controller” on page 255
- For the SAS Connectivity Card, refer to 3.1.18, “SAS Connectivity Card” on page 254.

A sample connection topology for advanced local storage with one SAS Connectivity Module installed is shown in Figure 5-14. Figure 5-15 on page 428.

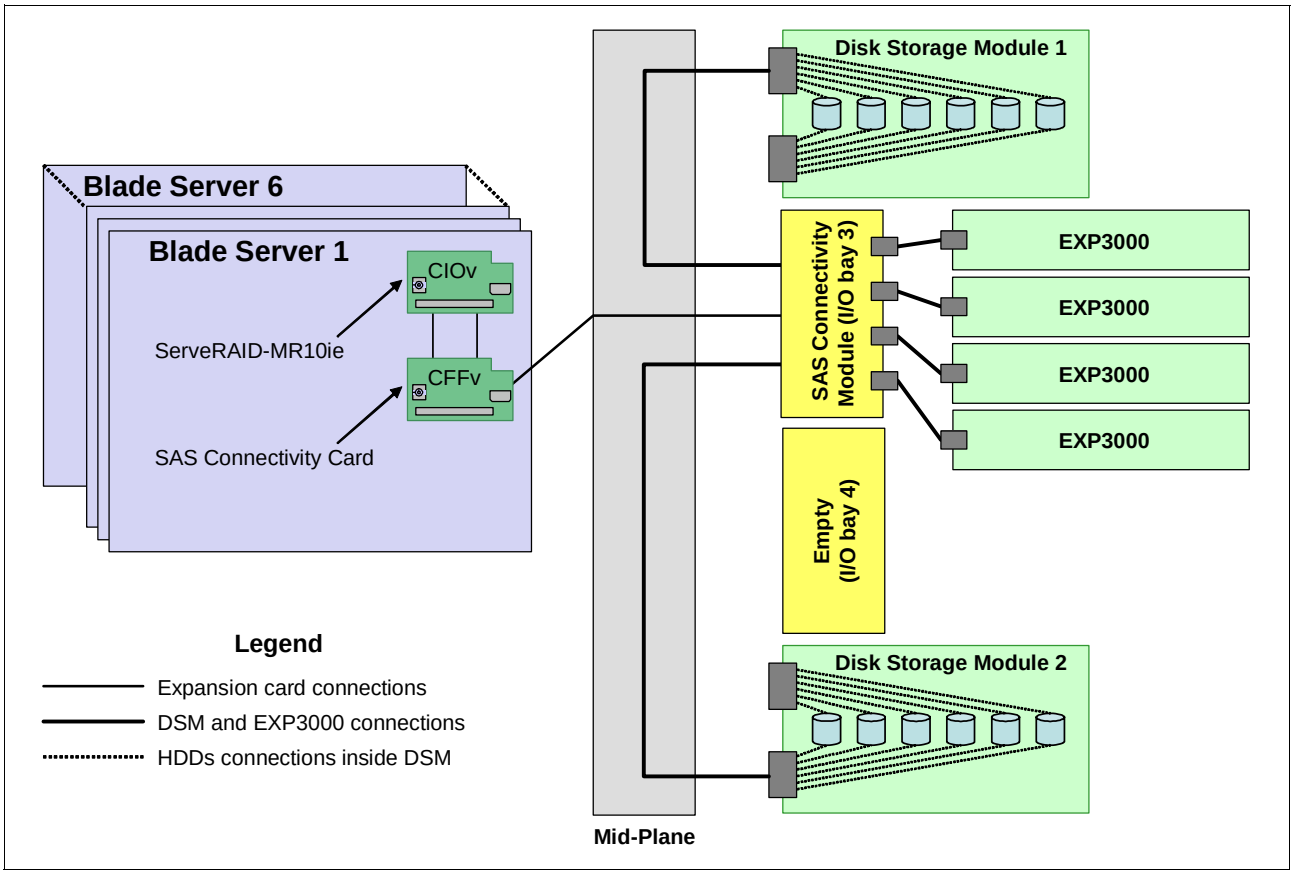

Figure 5-14 BladeCenter S and EXP3000 connections with one SAS Connectivity Module

Figure 5-15 on page 428 shows a sample connection topology for advanced local storage with two SAS Connectivity Modules installed.

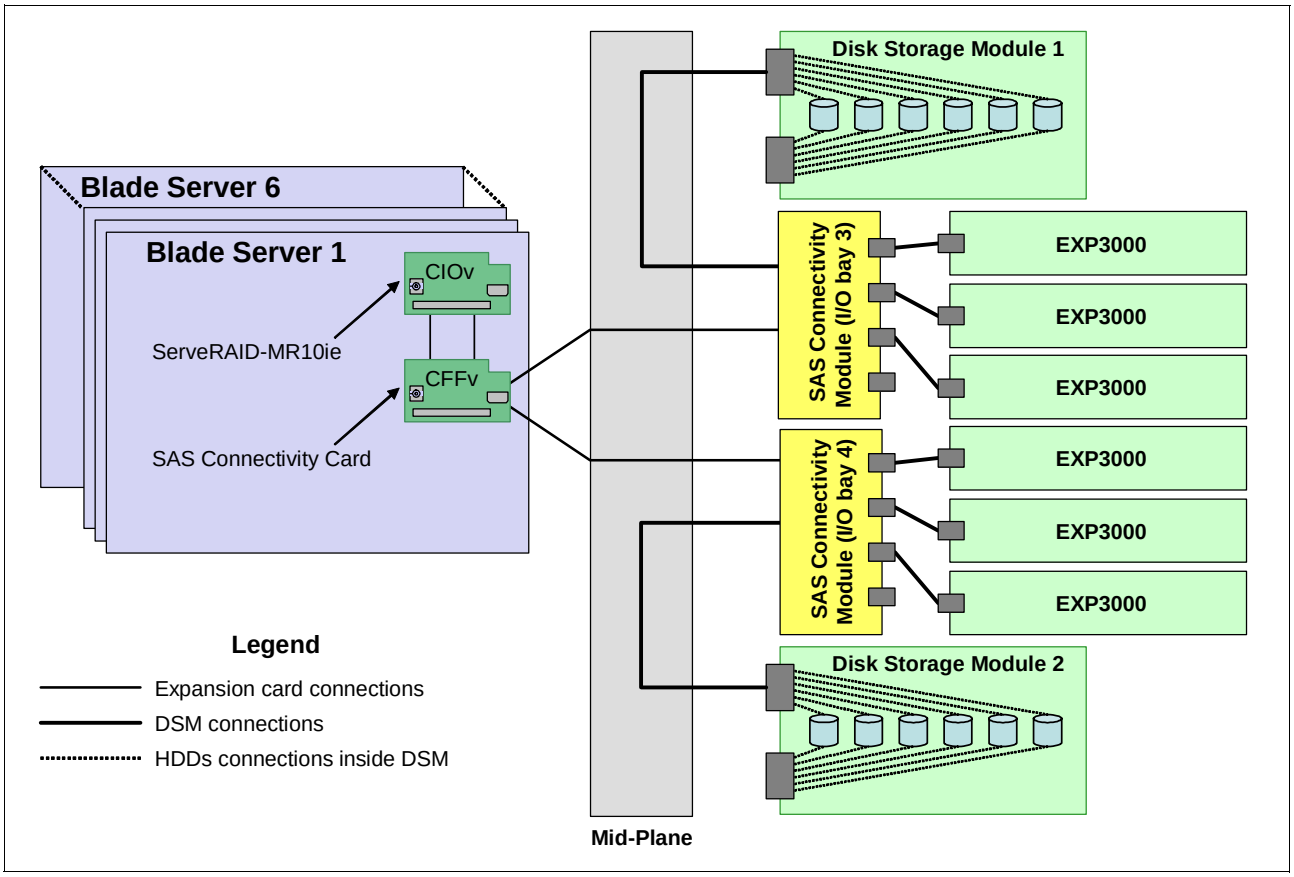

Figure 5-15 BladeCenter S and EXP3000 connections with two SAS Connectivity Module

You must take several considerations into account when planning BladeCenter S advanced local storage implementations:

- Every blade requiring integrated storage connectivity must have one ServeRAID MR10ie (CIOv) Controller installed.
- At least one DSM must be installed into the BladeCenter S chassis; a maximum of two DSMs are supported in one chassis. The IBM BladeCenter S does not ship with any DSMs as standard.
- At least one SAS Connectivity Module must be installed into the BladeCenter S chassis. A maximum of two SAS modules are supported.

  If two SAS connectivity modules are installed, the module in I/O module bay 3 controls access to storage module 1. The module in I/O module bay 4 controls access to storage module 2.
- IBM System Storage EXP3000 units (up to six EXP3000s - one per blade) can be attached to the BladeCenter S chassis to increase the number of drives per blade.
- A maximum of one EXP3000 is supported per single blade server.
- Each physical hard disk drive in any DSM can be assigned to the one blade server only, and the blade itself sees this physical disk (or disks) as its own “local” hard disk drives connected to the RAID controller.

  All disks in a single EXP3000 can only belong to the single blade server.

- The following cables can be used to attach EXP3000 to the SAS Connectivity Module:
  - IBM 3m SAS Cable, part number 39R6531, feature code 3707
  - IBM 1m SAS Cable, part number 39R6529, feature code 3708
- The SAS Connectivity Module controls disk assignments by using zoning-like techniques, which results in the SAS module maintaining single, isolated dedicated paths between physical HDDs and the blade server. Configuration of HDD assignments is done by the administrator, and the administrator uses predefined templates or creates custom configurations. Configuration of custom zoning is required to ensure that only one blade server has access to one EXP3000.
- RAID functionality is supplied by the ServeRAID-MR10ie controller installed into the blade server. RAID levels supported by ServeRAID-MR10i are:
  - RAID-0
  - RAID-1
  - RAID-5
  - RAID-6
  - RAID-10
  - RAID-50
  - RAID-60
- The maximum number of drives supported by the ServeRAID-MR10ie controller is 26 (two internal drives plus 12 drives in two DSMs plus 12 drives in one EXP3000).

  The maximum number of arrays supported by the ServeRAID-MR10ie is 128, and the maximum number of logical disks is 16 per array and 64 per controller.
- Mixing HDDs of different capacities in a single volume is supported. However, the total volume size is aligned with the size of the smallest HDD, and excess space on larger-sized HDDs is not used.
- Both SAS and SATA hot swap HDDs are supported, and intermixing SAS and SATA drives is supported as well. However, each volume must have hard disks of the same type; that is, SAS or SATA.

For more information refer to *ServeRAID-MR10ie (CIOv) Controller Installation and User Guide* at:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5078830

5.10.3 Advanced shared storage using the SAS RAID Controller Module

The main feature of advanced shared storage for BladeCenter S is the ability to create storage pools from hard disks in disk storage modules; create logical volumes in these pools; assign these volumes rather than physical disks to the blade servers; and map a single logical volume to several blade servers simultaneously.

Table 5-6 lists the advanced local storage components for BladeCenter S.

Table 5-6 Advanced local storage solution components for BladeCenter S

Component description	Part number	Feature code	Min/max quantity
Disk Storage Module (DSM)	43W3581	4545	1/2
SAS RAID Controller Module	43W3584	3734	2/2[a]
SAS Expansion Card	39Y9190	2979	1 per blade server

Component description	Part number	Feature code	Min/max quantity
Ethernet Switch Module in I/O bay 1[b]	Varies	Varies	1/1

a. Both two SAS RAID Controller modules *must* be installed into the BladeCenter S chassis.
b. For list of compatible Ethernet switch modules, refer to Table 1-2 on page 3.

Notes:

- The Copper Pass-thru Module (39Y9320, feature code 3219) is not supported by the SAS RAID Controller Module.
- If the Intelligent Copper Pass-thru Module (44W4483, feature code 5865) is used, then ports 7 and 14 *must* be attached to each other by standard Ethernet cable.

For more information:

- For DSMs, refer to 2.13, “BladeCenter S storage modules” on page 211
- For SAS RAID Controller Modules, refer to 2.5.2, “SAS RAID Controller Module” on page 166.
- For SAS expansion cards, refer to 3.1.17, “SAS Expansion Card” on page 252.

A sample topology for BladeCenter S with two SAS RAID Controller Modules is shown in Figure 5-16.

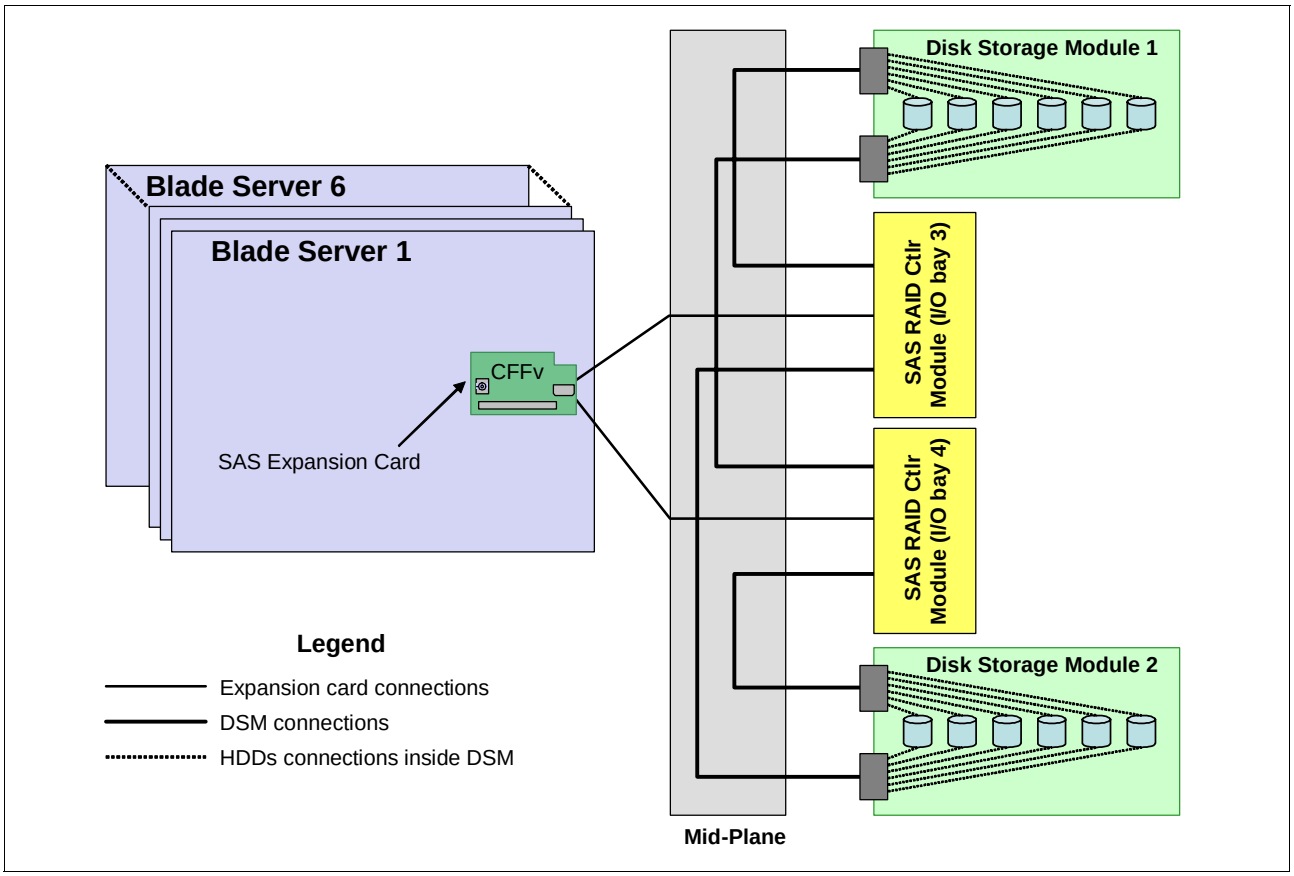

Figure 5-16 BladeCenter S SAS RAID Controller connections topology

You must take several considerations into account when planning BladeCenter S advanced shared storage implementations:

- Every blade requiring integrated storage connectivity must have one SAS Expansion Card installed.
- At least one DSM must be installed into the BladeCenter S chassis; a maximum of two DSMs are supported in one chassis. The IBM BladeCenter S does not ship with any DSMs as standard.
- Two SAS RAID Controller Modules *must* be installed into the BladeCenter S chassis.
- The SAS RAID Controller Module creates storage pools (or arrays), and the RAID level is defined for these storage pools. Logical volumes are created from storage pools. Volumes can be assigned to a specific blade, or can be shared by several blade servers.
- Zoning is supported by the SAS RAID controller module. However, zoning should not be used for regular operations (in other words, for purposes other than troubleshooting).
- RAID functionality is supplied by the SAS RAID Controller Modules installed into the BladeCenter S chassis.
  - RAID levels supported: 0, 1, 5, 10.
  - Maximum volume size is 2 TB.
  - Maximum number of volumes is 8 per blade server (for a total of 48 volumes per chassis).
  - One volume can be mapped to all 6 blades in the chassis.
- Mixing HDDs of different capacities in a single volume is supported. However, the total volume size is aligned with the size of the smallest HDD, and excess space on larger-sized HDDs is not used.
- Both SAS and Near-line SAS (NL SAS) hot swap HDDs are supported, and intermixing SAS/NL SAS drives is supported as well. However, each storage pool must have hard disks of the same type; that is, SAS or NL SAS. SATA drives are not supported by SAS RAID Controller Module.
- Global hot-spare drives are supported. The drive designated as a hot-spare should be as large as, or larger than, other drives in the system.
- Blade boot from logical volume is supported.
- Path failover is supported with IBM Subsystem Device Driver Device Specific Module (SDD DSM) for Windows and Device Mapper Multipath (DMM) for Red Hat/Novell SUSE Linux.
- Currently, external storage attachments are not supported.

Note: If external tape support is required for the SAS RAID Controller Module, then SPORE should be requested.

Ethernet switch modules supported:

- Nortel 1/10Gb Uplink Ethernet Switch Module, 44W4404, feature code 6980
- Nortel Layer 2/3 Copper Gb Ethernet Switch, 32R1860, feature code 3212
- Nortel Layer 2/3 Fiber Gb Ethernet Switch, 32R1861, feature code 3213
- Nortel Layer 2/7 Copper Gb Ethernet Switch, 32R1859, feature code 1494
- Cisco Catalyst Switch Module 3012, 43W4395, feature code 3174
- Server Connectivity Module, 39Y9324, feature code 3220
- Nortel 10G Uplink Ethernet Switch, 32R1783, feature code 3210
- Intelligent Copper Pass-thru Module, 44W4483, feature code 5865

Blade servers supported:

- HS12 (8014)
- HS20 (8843)
- HS21 (8853)
- HS21 XM (7995)
- LS21 (7971)
- LS22 (7901)
- LS41 (7972)
- LS42 (7902)

Network operating systems supported:

- Red Hat Enterprise Linux 5 AS & ES for x86/AMD64/EM64T
- Novell SUSE LINUX Enterprise Server 10 for x86/AMD64/EM64T
- Microsoft Windows Server 2003, Web Edition
- Microsoft Windows Server2003/2003 R2, Standard and Enterprise Editions
- Microsoft Windows Server 2003/2003 R2, Standard x64 and Enterprise x64 Editions
- Microsoft Windows Server 2008 Web, Standard, and Enterprise Editions
- VMware ESX Server 3.5 U2

For more information about the SAS RAID Controller solution, refer to:

- *SAS RAID Controller Installation and User Guide*:

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5078040
- *SAS RAID Controller Module Interoperability Guide*
- *SAS RAID Controller Module Detailed Host Attachment Guide* (Remote Boot included):

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5078491

5.11 InfiniBand and Fibre Channel SANs

If you plan to implement a high-performance enterprise grid computing solution, such as Oracle RAC on BladeCenter, it might require high-bandwidth, low-latency 4X InfiniBand fabric for internode communications to provide a scalable high-performance environment. At the same time, to reduce the costs of implementation and management, you might choose to use the same InfiniBand fabric for access to the network resources on your traditional LAN or SAN. In this case, the QLogic InfiniBand Ethernet Bridge Module or QLogic InfiniBand FC Bridge Module can help you achieve this goal. The IB-to-Ethernet bridging solution is discussed in 4.8, “InfiniBand and Ethernet LANs” on page 392.

One feature of InfiniBand is that it can run multiple upper layer protocols (ULP) like MPI, IPoOB, or SRP, simultaneously. This enables you to use a single InfiniBand fabric for different types of communications, including integration into an existing Fibre Channel network with the QLogic InfiniBand Fibre Channel Bridge Module. In this case, each blade that requires access to the storage resources (such as storage systems or tape devices) can be configured with one or more Virtual HBAs, and these Virtual HBAs provide required communications.

The high-speed InfiniBand fabric itself consists of one ore more high speed-capable chassis (IBM BladeCenter H or BladeCenter HT), blade servers with 4X InfiniBand expansion cards (one card per blade), a Cisco Systems 4X InfiniBand High Speed Switch Module (up to two modules per chassis can be used concurrently and they must be installed into I/O bays 7 and 9 only), and one or more QLogic InfiniBand Fibre Channel Bridge Modules (it depends on your requirements; however, up to two bridge modules per chassis can be used concurrently, and they must be installed into I/O bays 3 and 5).

To be able to communicate with your resources with InfiniBand using ULPs, host-side software called QuickSilver Fabric Access Software is required. This software (as well as the user guides for the software and bridge modules) can be found at the following address:

http://support.qlogic.com/support/oem_detail_all.asp?oemid=377

QuickSilver Fabric Access Software is supported with both Linux and Windows operating systems, but Linux supports more advanced functionality such as Port Recovery, which provides failover capabilities to the Virtual HBAs with BladeCenter's redundant infrastructure.

The ULP called SRP is used for integration with Fibre Channel SANs. It allows you to create virtual HBA interfaces on a server that can be used almost in the same way as physical HBAs. This protocol, in conjunction with the Fibre Channel bridge module, supports the following major features and functions that can be implemented in your solution:

- SRP-to-FCP bridging enables an IB-based blade to access devices on a Fibre Channel-based SAN (such as storage systems or tapes).
- Port Sharing provides the ability to use up to 36 active Virtual HBAs per one Fibre Channel bridge module (up to six external Fibre Channel ports) simultaneously.
- Port Recovery (Linux only) can be used to provide failover capability to the Virtual HBAs including failover to another SRP initiator port, failover to another port of the same bridge module, or failover to another bridge module (in the same or in different chassis).
- Load Balancing and Path Redundancy (Linux only) perform host-based load balancing, allowing you to utilize two or more independent paths to the same target device.
- LUN Masking limits visibility of a device's LUNs to the hosts.
- LUN Mapping provides a means of assigning LUN IDs to hosts independently.

Abbreviations and acronyms

ABR	Automatic BIOS recovery
AC	alternating current
ACL	access control list
AES	Advanced Encryption Standard
AMD	Advanced Micro Devices™
AMM	Advanced Management Module
API	application programming interface
APV	Advanced Power Virtualization
ARP	Address Resolution Protocol
AS	Australian Standards
ASF	Alert Standard Format
ASIC	application-specific integrated circuit
ASR	automatic server restart
BASP	Broadcom Advanced Server Program
BBI	browser-based interface
BCM	Broadcom
BE	Broadband Engine
BGP	Border Gateway Protocol
BIOS	basic input output system
BMC	Baseboard Management Controller
BNT	BLADE Network Technologies, Inc
BOFM	BladeCenter Open Fabric Manager
BPDU	Bridge protocol data unit
BSE	BladeCenter Storage Expansion
BSMP	blade system management processor
BTU	British Thermal Unit
CCDA	Cisco Certified Design Associate
CCNP	Cisco Certified Network Professional
CCSP	Cisco Certified Security Professional
CD-ROM	compact disc read only memory
CDP	Cisco Discovery Protocol
CE	Conformité Européene
CLI	command-line interface
CNA	Cisco Network Assistance
CNS	Cisco Network Services
COG	configuration and option guide
CPM	Copper Pass-thru Module
CPU	central processing unit
CRU	customer replaceable units
CSM	Cluster Systems Management
CTS	clear to send
DASD	direct access storage device
DC	domain controller
DDM	Deployment and Management
DDR	Double Data Rate
DHCP	Dynamic Host Configuration Protocol
DIMM	dual inline memory module
DIP	destination IP
DMAC	destination MAC address
DNS	Domain Name System
DP	dual processor
DPOD	Dynamic Ports on Demand
DSA	Dynamic System Analysis
DSCP	Differentiated Services Code Point
DSM	disk storage module
DSUB	D-subminiature
DTP	Dynamic Trunking Protocol
DVI	Digital Video Interface
DVMRP	Distance Vector Multicast Routing Protocol
DVS	Digital Video Surveillance
ECC	error checking and correcting
EDA	Electronic Design Automation
EIGRP	Enhanced Interior Gateway Routing Protocol
EMC	electromagnetic compatibility
EMEA	Europe, Middle East, Africa
EOT	Enhanced object tracking
EPOW	Early Power Off Warning
ESD	electrostatic discharge
ESM	Ethernet switch modules
ETSI	European Telecommunications Standard Industry
FAN	Fabric Address Notification
FB-DIMM	Fully Buffered DIMMs
FBDIMM	Fully Buffered DIMM
FC	Fibre Channel

FCP	Flow Control Packet
FCSM	Fibre Channel Switch Module
FDD	floppy diskette drive
FDX	full duplex
FSB	front-side bus
FTP	File Transfer Protocol
FTSS	Field Technical Sales Support
GB	gigabyte
GUI	graphical user interface
HA	high availability
HBA	host bus adapter
HCA	host channel adapter
HD	high definition
HDD	hard disk drive
HH	half high
HPC	high performance computing
HS	hot swap
HSDC	high speed daughter card
HSESM	high speed Ethernet switch module
HSFF	high-speed form factor
HSIBPM	high-speed InfiniBand pass-thru module
HSIBSM	high speed InfiniBand switch module
HSRP	Hot Standby Routing Protocol
HT	Hyper-Threading
HTTP	Hypertext Transfer Protocol
I/O	input/output
IB	InfiniBand
IBBM	InfiniBand bridge module
IBM	International Business Machines
ICMP	Internet control message protocol
ICPM	Intelligent Copper Pass-thru Module
ID	identifier
IDE	integrated drive electronics
IEC	International Electro-technical Commission
IEEE	Institute of Electrical and Electronics Engineers
IGESM	Intelligent Gigabit Ethernet Switch Module
IGMP	Internet Group Management Protocol
IGRP	Interior Gateway Routing Protocol
IM	instant messaging
IME	Integrated Mirroring Enhanced
IOS	Internetwork Operating System
IP	Internet Protocol
IPM	Intelligent Pass-thru Module
IPMI	Intelligent Platform Management Interface
IPTV	Internet Protocol Television
IRDP	ICMP Router Discovery Protocol
IS	information store
ISL	Inter-Switch Link
ISMP	Integrated System Management Processor
ISP	Internet service provider
IT	information technology
ITS	IBM Integrated Technology Services
ITSO	International Technical Support Organization
IVM	Integrated Virtualization Manager
KB	kilobyte
KVM	keyboard video mouse
LACP	Link Aggregation Control Protocol
LAN	local area network
LED	light emitting diode
LLDP	Link Layer Discovery Protocol
LPAR	logical partitions
LPH	low profile handle
LR	long range
LTO	Linear Tape-Open
LUN	logical unit number
MAC	media access control
MAN	metropolitan area network
MB	megabyte
MDS	Multilayer DataCenter Switch
MIB	management information base
MIO	Memory and I/O
MM	Management Module
MMF	Multi Mode Fiber
MP	multiprocessor
MPE	Multi Processor Expansion
MPI	Message Passing Interface
MSDP	Multicast Source Discovery Protocol
MSIM	Multi-Switch Interconnect Module
MSTP	Multiple Spanning Tree Protocol

MTM	machine-type-model
MVR	Multicast VLAN registration
NAT	Network Address Translation
NDCLA	Non-Disruptive Code Load Activation
NEBS	Network Equipment Building System
NGN	next-generation network
NIC	network interface card
NMI	non-maskable interrupt
NOS	network operating system
NPIV	N_Port ID Virtualization
NSF	Notes Storage File
NTP	Network Time Protocol
OBFL	On-board failure logging
ODPA	On-Demand Port Activation
OFED	OpenFabrics Enterprise Distribution
OFM	Open Fabric Manager
OPM	Optical Pass-thru Module
OS	operating system
OSPF	Open Shortest Path First
PBR	Policy-based routing
PC	personal computer
PCI	Peripheral Component Interconnect
PDF	Portable Document Format
PDU	power distribution unit
PFA	Predictive Failure Analysis
POST	power-on self test
PPP	point-to-point protocol
PVST	Per-VLAN Spanning Tree
PXE	Preboot Execution Environment
RAC	Real Application Clusters
RADIUS	Remote Authentication Dial In User Service
RAID	redundant array of independent disks
RAM	random access memory
RAS	remote access services; row address strobe
RDAC	Redundant Disk Array Controller
RDC	Remote Desktop Connection
RDIMM	registered DIMM
RDM	Remote Deployment Manager
RDMA	Remote Direct Memory Access
RETAIN	Remote Electronic Technical Assistance Information Network
RHEL	Red Hat Enterprise Linux
RIP	Routing Information Protocol
RMCP	Remote Management Control Protocol
RMON	Remote Monitoring
RP	route processor
RPF	reverse path forwarding
RPM	revolutions per minute
RPQ	Request Per Qualification
RSA	Remote Supervisor Adapter
RSCN	Registered State Change Notification
RSTP	Rapid Spanning Tree Protocol
SAN	storage area network
SAS	Serial Attached SCSI
SASCM	SAS Connectivity Module
SATA	Serial ATA
SBB	Sales Building Block
SCM	Supply Chain Management
SCO	Santa Cruz Operation, Inc
SCSI	Small Computer System Interface
SDD	Subsystem Device Driver
SDK	Software Developers' Kit
SDR	Single Data Rate
SDRAM	static dynamic RAM
SFF	Small Form Factor
SFP	small form-factor pluggable
SIMD	single instruction multiple data
SIO	Storage and I/O
SIP	source IP
SLB	Server Load Balancing
SLES	SUSE Linux Enterprise Server
SM	Subnet Manager
SMAC	source MAC address
SMI-S	Storage Management Initiative - Specification
SMP	symmetric multiprocessing
SMS	System Management Services
SNMP	Simple Network Management Protocol
SOL	Serial over LAN
SPORE	ServerProven Opportunity Request for Evaluation
SR	short range

SRP	Storage RDMA Protocol
SRR	shaped round robin
SSCT	Standalone Solution Configuration Tool
SSD	solid state drive
SSH	Secure Shell
SSL	Secure Sockets Layer
SSP	Serial SCSI Protocol
STP	Spanning Tree Protocol
TACACS	Terminal Access Controller Access Control System
TB	terabyte
TCO	total cost of ownership
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
TDR	Time Domain Reflector
TFTP	Trivial File Transfer Protocol
TOE	TCP offload engine
TX	transmit
UDLD	UniDirectional link detection
UDP	user datagram protocol
ULP	upper layer protocols
URL	Uniform Resource Locator
USB	universal serial bus
UTF	Universal Telco Frame
UTP	unshielded twisted pair
VBS	Virtual Blade Switch
VGA	video graphics array
VIOS	Virtual I/O Server
VLAN	virtual LAN
VLP	very low profile
VM	virtual machine
VMPS	VLAN Membership Policy Server
VNC	Virtual Network Computing
VOIP	Voice over Internet Protocol
VPD	vital product data
VPN	virtual private network
VQP	VLAN Query Protocol
VRRP	virtual router redundancy protocol
VSAN	Virtual Storage Area Network
VT	Virtualization Technology
VTP	VLAN Trunking Protocol
WAN	wide area network
WCCP	Web Cache Communication Protocol
WOL	Wake on LAN
WRR	weighted round-robin
WTD	Weighted tail drop
WWN	World Wide Name
XDR	extreme data rate
XM	extended memory

Related publications

We consider the publications that we list in this section particularly suitable for a more detailed discussion of the topics that we cover in this book.

IBM Redbooks

You can search for, view, or download books, papers, Technotes, draft publications and additional materials, as well as order hardcopy Redbooks, at the IBM Redbooks Web site:

ibm.com/redbooks

Related publications from IBM Redbooks Publications include the following:

- *IBM System Storage Solutions Handbook*, SG24-5250
- *Integrating IBM Director with Enterprise Management Solutions*, SG24-5388
- *Introduction to Storage Area Networks*, SG24-5470
- *IBM System Storage: Implementing an IBM SAN*, SG24-6116
- *The IBM eServer BladeCenter JS20*, SG24-6342
- *IBM eServer xSeries and BladeCenter Server Management*, SG24-6495
- *IBM BladeCenter 4Gb SAN Solution*, SG24-7313
- *Programming the Cell Broadband Engine: Examples and Best Practices*, SG24-7575
- *Nortel Networks L2/3 Ethernet Switch Module for IBM eServer BladeCenter*, REDP-3586
- *IBM BladeCenter Networking with the 4-Port Gigabit Ethernet Switch Module*, REDP-3660
- *IBM eServer BladeCenter Layer 2-7 Network Switching*, REDP-3755
- *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter*, REDP-3869
- *Implementing Cisco InfiniBand on IBM BladeCenter*, REDP-3949
- *IBM BladeCenter iSCSI SAN Solution*, REDP-4153
- *Implementing Sun Solaris on IBM BladeCenter Servers*, REDP-4269
- *Networking with the Nortel 10 Gb Ethernet Switch Module for BladeCenter H*, REDP-4314
- *Implementing the IBM BladeCenter HC10 Workstation Blade*, REDP-4331
- *Implementing the QLogic Intelligent Pass-thru Module for IBM BladeCenter*, REDP-4342
- *Implementing the Brocade Access Gateway for IBM BladeCenter*, REDP-4343
- *Implementing the IBM BladeCenter S Chassis*, REDP-4357

Product publications

This section lists a subset of BladeCenter publications. You can browse for other BladeCenter publications at the following page:

http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/resourceselect?brandind=5000020&taskind=2&pubs=Y

Brocade Switch Modules

- *Brocade 4 Gb SAN Switch Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-62707

- *Brocade Enterprise/Entry SAN Switch Module for IBM eServer BladeCenter Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55363

- *Brocade Fabric OS v4.4.1a* Release Notes

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55419

- *Brocade SAN Switch Modules for IBM eServer BladeCenter: Design, Deployment and Management Guide DDM*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55327

- *Brocade Enterprise/Entry SAN Switch Module Installation Guide - IBM eServer BladeCenter*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55363

- *Multipath Subsystem Device Driver User's Guide* at

  http://www.ibm.com/support/docview.wss?uid=ssg1S7000303

Cisco 4 GB FC Switch Module

- *Cisco Fibre Channel Switch Module set-up information - IBM BladeCenter*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5070625

Cisco 4X InfiniBand Switch Module

- *Cisco 4X InfiniBand Switch Module for IBM BladeCenter User Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-65966

Cisco Topspin InfiniBand Switch Module

- *Topspin InfiniBand Switch Module for BladeCenter User Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-60200

Cisco Systems Fiber Intelligent Gb Ethernet Switch Module

- *Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-63993

- Cisco *Systems Intelligent Gigabit Ethernet Switch Module Software Configuration Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55261

Cisco Systems Intelligent Gb Ethernet Switch Module

- *Release Notes*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55262

- *Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-57858

- *Software Configuration Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55261

- *Command Reference*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55259
- *System Message Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-55259

Cisco Catalyst Switch Module 3110X, 3110G, and 3012

- *Cisco Catalyst Switch Module 3110 and 3012 System Message Guide*
- *Cisco Catalyst Switch Module 3110 and 3012 Software Configuration Guide*
- *Cisco Catalyst Switch Module 3110G, 3110X, and 3012 Hardware Installation Guide*
- *Cisco Catalyst Switch Module 3110G, 3110X, and 3012 Getting Started Guide*
- *Cisco Catalyst Switch Module 3110 and 3012 Command Reference*

All these publications can be downloaded from:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5075938

IBM BladeCenter Boot Disk System Solution

- *IBM BladeCenter Boot Disk System Installation and User Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5072756

IBM 4X InfiniBand Pass-thru Module

- *4X InfiniBand Pass-thru Module Installation and User Guide*:

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5072709

IBM Copper Pass-thru Module

- *Copper Pass-thru Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-57332
- *Copper Pass-thru Module (CPM) Cable Installation Instructions*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-59020

IBM Intelligent Copper Pass-thru Module

- *Intelligent Copper Pass-thru Module Installation Guide*

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5076328

IBM Optical Pass-thru Module

- *Optical Pass-thru Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-52864
- *Troubleshooting Optical Pass-thru Module (OPM) (02R9080) issues*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-61155

IBM Server Connectivity Module

- *IBM Server Connectivity Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-64155
- IBM *Server Connectivity Module User's Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-64190

IBM Multi-Switch Interconnect Module

- *Installation and User's Guide Multi-Switch Interconnect Module - IBM BladeCenter*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5070476

IBM Multi-Switch Interconnect Module for BladeCenter HT

- *Multi-Switch Interconnect Module Installation Instructions - IBM BladeCenter HT*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5073195

IBM BladeCenter Open Fabric Manager

- *IBM BladeCenter Open Fabric Manager Installation and User Guide*:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5073103

IBM SAN10Q Fibre Channel Switch

- *IBM System Storage SAN10Q 4 Gbps 10-Port Fibre Channel Switch Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=ssg1S7001476

- *IBM System Storage SAN10Q 4 Gbps 10-Port Fibre Channel Switch Manager User Guide*

  http://www.ibm.com/support/docview.wss?uid=ssg1S7001475

IBM SAS Connectivity Module

- *SAS Connectivity Module Installation and User Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5072374

IBM SAS RAID Controller Module

- *SAS RAID Controller Module Installation and User Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5078040

Intel Gigabit Ethernet Switch Module for IBM BladeCenter T

- *Intel Gigabit Ethernet Switch Module for IBM eServer BladeCenter Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-56746

- *Intel Gigabit Ethernet Switch Module for IBM eServer BladeCenter User's Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-56745

- *Intel Gigabit Ethernet Switch Module for IBM eServer BladeCenter Command Line Interface (CLI) Reference Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-56744

McDATA Fibre Channel Switch Module

- *McDATA 6-Port Fibre Channel Switch Module for IBM eServer BladeCenter Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-59840

- *McDATA 6-Port Fibre Channel Switch Module for IBM eServer BladeCenter Management Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-59910

Nortel 10 Gigabit Ethernet High Speed Switch Module

- *Installation Guide for Nortel 10 Gb Ethernet Switch Module - IBM BladeCenter*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5070408
- *Nortel 10 Gb Ethernet Switch Module Application Guide - IBM BladeCenter*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5070403

Nortel 10 Gigabit Uplink Ethernet Switch Module

- *Nortel 10 Gb Uplink Ethernet Switch Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-64158
- *Nortel 10 Gb Uplink Ethernet Switch Module Application Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-64687

Nortel Networks L2/3 Copper and Fibre Channel GbE Switch Modules

- *Nortel Networks Layer 2/3 GbE Switch Module Alteon OS 21.0 Release Notes*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58337
- *Nortel Networks Layer 2/3 GbE Switch Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58334
- *Nortel Networks Layer 2/3 GbE Switch Module Alteon OS 21.0 Application Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58325
- *Nortel Networks Layer 2/3 GbE Switch Module Alteon OS 21.0 Command Reference*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58326
- *Nortel Networks Layer 2/3 GbE Switch Module Alteon OS 21.0 BBI Quick Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58323

Nortel Networks Layer 2-7 GbE Switch Module

- *Nortel Networks Layer 2-7 GbE Switch Module for IBM eServer BladeCenter Products: Alteon OS 20.2 Release Notes*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-54999
- *Nortel Networks Layer 2-7 GbE Switch Module for IBM eServer BladeCenter Products: Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-53065
- *Nortel Networks Layer 2-7 GbE Switch Module for IBM eServer BladeCenter Products: Alteon OS 20.2 Application Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-53098
- *Nortel Networks Layer 2-7 GbE Switch Module for IBM eServer BladeCenter Products: Alteon OS 20.2 Browser-Based Interface Quick Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-53102
- *Nortel Networks Layer 2-7 GbE Switch Module for IBM eServer BladeCenter Products: Alteon OS 20.2 Command Reference Edition B*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-53101

Nortel 1/10 Gb Uplink Ethernet Switch Module

- *Nortel 1/10 Gb Uplink Ethernet Switch Module Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5076217

- *Nortel 1/10 Gb Uplink Ethernet Switch Module Application Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5076214

- *Nortel 1/10 Gb Uplink Ethernet Switch Module Command Reference*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5076525

QLogic 4 Gb Intelligent Pass-thru Module

- *QLogic 4 Gb Intelligent Pass-thru Module and 4 Gb SAN Switch Modules Getting Started Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5071313

- *QLogic 4 Gb Intelligent Pass-thru Module and 4 Gb SAN Switch Module for IBM Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5071303

QLogic 4 Gb SAN Switch Module

- *QLogic 4 Gb Intelligent Pass-thru Module and 4 Gb SAN Switch Modules Getting Started Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5071313

- *QLogic 4 Gb Intelligent Pass-thru Module and 4 Gb SAN Switch Module for IBM Installation Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5071303

QLogic 8 Gb SAN Switch Module

- *QLogic 8 Gb Intelligent Pass-thru Module and 20-port 8 Gb SAN Switch Module Installation and User's Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5079302

QLogic 8 Gb Intelligent Pass-thru Module

- *QLogic 8 Gb Intelligent Pass-thru Module and 20-port 8 Gb SAN Switch Module Installation and User's Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5079302

QLogic InfiniBand Ethernet and Fibre Channel Bridge Modules

- *QLogic InfiniBand Ethernet Bridge Module and Fibre Channel Bridge Module Installation Guide - IBM BladeCenter H (Type 8852)*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5070108

- *InfiniBand Fibre Channel Bridge Module for IBM BladeCenter & InfiniBand Ethernet Bridge Module for IBM BladeCenter - Support*

  http://support.qlogic.com/support/oem_detail_all.asp?oemid=377

Management Modules

- *IBM eServer BladeCenter and BladeCenter Management Module User's Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-45153

- *Management Module Command Line Interface Reference Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-54667
- *IBM eServer BladeCenter and BladeCenter Management Module User's Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-45153
- *Management Module Command Line Interface Reference Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-54667
- *Management Processor Command Line Interface Utility Version 3.00*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-54216

Other BladeCenter publications

- *IBM eServer BladeCenter SAN Solutions Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-54854
- *Fibre Channel Switch Interoperability Guide Version 4.00*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58206
- *SAN Solutions Guide version 1.00 - IBM eServer*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-46241

Other publications

These publications are also relevant as further information sources:

- *A Comparison of Single-Core and Dual-Core Opteron Processor Performance for HPC,* Douglas M. Pase and Matthew A. Eckl

  http://www.ibm.com/servers/eserver/opteron/benchmarks/wp.html
- *Systems Management for IBM eServer xSeries and BladeCenter Servers - 2005,* Mark T Chapman

  http://www-306.ibm.com/common/ssi/apilite?infotype=SA&infosubt=WH&lastdays=1825&hitlimit=200&ctvwcode=US&pubno=XSWO*USEN,XSLO*USEN&appname=STG_XS_USEN_WH&additional=summary&contents=keeponlit

  or

  ftp://ftp.software.ibm.com/common/ssi/rep_wh/n/XSW00128OUSEN/XSW00128OUSEN.PDF
- *IBM BladeCenter Deployment Guide*, Chris Gillen, Sam Litenatsky, Kevin Conover

  http://www.ibm.com/support/docview.wss?uid=tss1wp100564
- *BladeCenter processor blades, I/O expansion adapters, and units,* J. E. Hughes, M. L. Scollard, R. Land, J. Parsonese, C. C. West, V. A. Stankevich, C. L. Purrington, D. Q. Hoang, G. R. Shippy, M. L. Loeb, M. W. Williams, B. A. Smith, D. M. Desaiet al, IBM Journal of Research & Development, Vol. 49, No. 6, 837-859, November 2005

  http://www.research.ibm.com/journal/rd49-6.html
- *BladeCenter T system for the telecommunications industry*, S. L. Vanderlinden, B. O. Anthony, G. D. Batalden, B. K. Gorti, J. Lloyd, J. Macon, Jr., G. Pruett, B. A. Smith, IBM Journal of Research & Development, Vol. 49, No. 6, 873-885, November 2005

  http://www.research.ibm.com/journal/rd49-6.html
- *BladeCenter chassis management*, T. Brey, B. E. Bigelow, J. E. Bolan, H. Cheselka, Z. Dayar, J. M. Franke, D. E. Johnson, R. N. Kantesaria, E. J. Klodnicki, S. Kochar, S. M.

Lardinois, C. A. Morrell, M. S. Rollins, R. R. Wolford, D. R. Woodha, IBM Journal of Research & Development, Vol. 49, No. 6, 873-885, November 2005

http://www.research.ibm.com/journal/rd49-6.html

- *BladeCenter systems management software*, G. Pruett A. Abbondanzio, J. Bielski, T. D. Fadale, A. E. Merkin, Z. Rafalovich, L. A. Riedle, J. W. Simpson, IBM Journal of Research & Development, Vol. 49, No. 6, 873-885, November 2005

  http://www.research.ibm.com/journal/rd49-6.html

- *BladeCenter networking*, S. W. Hunter, N. C. Strole, D. W. Cosby, D. M. Green, IBM Journal of Research & Development, Vol. 49, No. 6, 873-885, November 2005

  http://www.research.ibm.com/journal/rd49-6.html

- *BladeCenter packaging, power, and cooling*, M. J. Crippen, R. K. Alo, D. Champion, R. M. Clemo, C. M. Grosser, N. J. Gruendler, M. S. Mansuria, J. A. Matteson, M. S. Miller, B. A. Trumbo, IBM Journal of Research & Development, Vol. 49, No. 6, 887-903, November 2005

  http://www.research.ibm.com/journal/rd49-6.html

Online resources

These Web sites are also relevant as further information sources.

Product information

- IBM BladeCenter

  http://ibm.com/bladecenter

- Standalone Solutions Configuration Tool (SSCT)

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-62168

- *Configuration and Options Guide (COG)*

  http://www.ibm.com/support/docview.wss?uid=psg1SCOD-3ZVQ5W

- *BladeCenter Interoperability Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5073016

- Benchmarks for the BladeCenter

  http://www.ibm.com/servers/eserver/bladecenter/benchmarks/

- BladeCenter announcement letters

  http://www-306.ibm.com/common/ssi/OIX.wss

- pSeries White papers

  http://www.ibm.com/servers/eserver/pseries/hardware/whitepapers/power4.html

- *The Necessity of Change Management in the Intel Server Space*

  http://www.ibm.com/servers/eserver/xseries/literature/wp_lit.html

- Microsoft's licensing policy on logical processors

  http://www.microsoft.com/licensing/highlights/multicore.mspx

- SCO certified and compatible hardware site

  http://wdb1.caldera.com/chwp/owa/hch_search_wizard.screen

- Altiris Web site deployment solutions and alliance

  http://www.altiris.com/products/deploymentsol/
  http://www.altiris.com/partners/alliances/ibm/
- Cluster System Management documents

  http://publib.boulder.ibm.com/clresctr/windows/public/clusterbooks.html
- IBM Dynamic System Analysis (DSA)

  http://www.ibm.com/systems/management/dsa.html
- IBM Storage

  http://www.ibm.com/servers/storage/tape/
- Telcordia

  http://www.telcordia.com
- European Telecommunication Standards Institute (ETSI)

  http://www.etsi.org
- HS40 onboard Intel NICs with Linux OS, use bonding driver

  http://downloadfinder.intel.com/scripts-df-external/Detail_Desc.aspx?agr=Y&DwnldID=5600&ProductID=61
- *Using NIC Teaming Adapters with Network Load Balancing May Cause Network Problems*

  http://go.microsoft.com/fwlink/?LinkId=18366
- QLogic SAN Interoperability

  http://www.qlogic.com/interopguide/
- McData compatibility documents

  http://www.mcdata.com/support/sil
- Brocade SAN Interoperability

  http://www.brocade.com/products/interop/san-interop.jsp
- IBM Total Storage DS4000, refer to DS4000 interoperability matrix

  http://www.ibm.com/servers/storage/disk/ds4000/interop-matrix.html
- *Boot from SAN in Windows Server 2003 and Windows 2000 Server*

  http://www.microsoft.com/windowsserversystem/wss2003/techinfo/plandeploy/bootfromsaninwindows.mspx
- xSeries® configurator tools

  http://www.pc.ibm.com/us/eserver/xseries/library/configtools.html

Technology

- Intel Xeon DP

  http://www.intel.com/design/xeon/prodbref
- Intel Xeon Processor MP

  http://www.intel.com/design/xeon/xeonmp/prodbref
- Intel NetBurst architecture

  http://www.intel.com/cd/ids/developer/asmo-na/eng/44004.htm

- Windows-Based Servers and Intel Hyper-Threading Technology

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

- HyperTransport and the HyperTransport Consortium

  http://www.hypertransport.org/

- Intel 64 Technology (EM64T)

  http://www.intel.com/technology/64bitextensions/
  ftp://ftp.software.ibm.com/eserver/benchmarks/wp_Dual_Core_072505.pdf

- 64-bit technology and performance

  http://www-03.ibm.com/servers/eserver/xseries/benchmarks/related.html

- PowerPC 970 and PowerPC 970FX microprocessors

  http://www.chips.ibm.com

- *BCM570X Broadcom NetXtreme Gigabit Ethernet Teaming* white paper

  http://www.broadcom.com/collateral/wp/570X-WP100-R.pdf

Services

For information about other available IBM Global Services and ITS services

http://www.ibm.com/services/
http://www.ibm.com/services/us/index.wss/gen_it
http://www.pc.ibm.com/ww/eserver/xseries/services/index.html

- Sales Configuration Aid

  http://www.pc.ibm.com/support?page=MIGR-41411

- Warranty Service Upgrade options

  http://www.ibm.com/services/us/index.wss/offering/its/a1001480

- Maintenance Agreement options

  http://www.ibm.com/services/us/index.wss/offering/its/a1001505

- SPORE program

  http://exist.raleigh.ibm.com/spore/

- IBM Business Continuity And Consulting Services

  http://www.ibm.com/services/us/index.wss/of/bcrs/a1000387

Product support

- IBM support for employees

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

- IBM support for Business Partners

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

- IBM support for Asia Pacific

  http://w3.ibm.com/support/ap/

- xSeries Solution Assurance Trigger criteria

  http://w3.ibm.com/support/assure/assur30i.nsf/PubAllNum/SA200

- Software and device drivers matrix (chassis and blade servers)

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-63017
- Advanced Management Module firmware release matrix:

  http://www.ibm.com/support/docview.wss?uid=psg1SERV-AMM
- Troubleshooting Optical Pass-thru Module issues

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-61155
- NEBS compatibility list

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-54720
- *BladeCenter - Serial over LAN Setup Guide* and *Enabling and configuring SOL (Serial over LAN) on an 8842 Blade Server* White paper

  http://www.ibm.com/servers/eserver/support/bladecenter/js20/installinghwonly.html
- UpdateXpress

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-53046
- IBM ServerGuide Scripting Toolkit

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-53564
- *IBM Remote Deployment Manager 4.20 Compatibility Guide* for compatibility

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-56662
- Remote Deployment Manager White papers

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-53487
- *ServerGuide Scripting Toolkit User Reference Version 1.2.01*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-60190
- MegaRAID Configuration Console (HS20)

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-52761
- ServeRAID Manager

  http://www.ibm.com/support/docview.wss?uid=psg1SERV-RAID
- IBM Director upward integration modules

  http://www.ibm.com/support/docview.wss?uid=psg1SERV-DIRECT
- Electronic Service Agent portal site xSeries - Director Extension

  http://www.ibm.com/services/electronic
- Intel Advanced Networking Services

  http://support.intel.com/support/network/sb/CS-009747.htm
- VMware support

  http://www.vmware.com/vmtn/resources/esx_resources.html
- Broadcom Advanced Server Program (BASP) is Broadcom Teaming driver

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-43815
- Intel PROSet is Intel teaming driver

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-54793
- *Introduction to iSCSI in BladeCenter*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-60300

- *IBM eServer BladeCenter Switch Interoperability Guide*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-58206
- *Remote SAN Boot for HS20 (8678/8832) and HS40 (8839)*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-57563
- *IBM eServer BladeCenter Remote SAN Boot - JS20*

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-57235

ServerProven

- Certified PCI adapters for the PCI Expansion Unit

  http://www.ibm.com/servers/eserver/serverproven/compat/us/peu/9OP3721.html
- ServerProven

  http://www.ibm.com/servers/eserver/serverproven/compat/us/
- TotalStorage Proven

  http://www.storage.ibm.com/proven

Training

- Education classes on IBM BladeCenter products

  http://www.ibm.com/servers/eserver/education/cust/blades/custblades.html
- IBM eServer xSeries Servers classes

  http://www.pc.ibm.com/ww/eserver/xseries/education/
- IBM employees training courses

  http://lt.be.ibm.com/stg

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

A

B

C

D

E

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

IBM BladeCenter Products and Technology

IBM BladeCenter Products and Technology

Redbooks®

Describes the BladeCenter chassis and blade server technology

Provides details of available I/O modules and expansion options

Explains networking and storage configurations

IBM BladeCenter remains an innovative solution to running business solutions. IBM BladeCenter builds on the IBM commitment to integrating server, storage and networking functionality with technology exchange, and heterogeneous management. IBM BladeCenter offers the ease, density, availability, affordability, and scalability that are central to the blade technology promise.

Blade servers have captured industry focus because of their modular design, which can reduce cost with a more efficient use of valuable floor space, and its simplified management, which can help to speed up such tasks as deploying, reprovisioning, updating, and troubleshooting hundreds of blade servers. In addition, blade servers provide improved performance by doubling current rack density. By integrating resources and sharing key components, not only are costs reduced, but also availability is increased.

This IBM Redbooks publication describes IBM BladeCenter and discusses the technology and features of the different chassis, blade server models, and connectivity options. We go into details about every major component and provide guidance as to networking and storage connectivity.

This fourth edition includes new products that were announced in the second half of 2008.

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

SG24-7523-03 ISBN 0738432504