


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



Ilya Krutov
Alexander Grechnev

Redbooks



International Technical Support Organization

IBM BladeCenter Products and Technology

| November 2013

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

Twelfth Edition (November 2013)

This edition applies to the following products:

- ▶ IBM BladeCenter E, 8677
- ▶ IBM BladeCenter H, 8852
- ▶ IBM BladeCenter HT, 8740 and 8750
- ▶ IBM BladeCenter S, 8886
- ▶ IBM BladeCenter HS12, 8028
- ▶ IBM BladeCenter HS22, 7870
- ▶ IBM BladeCenter HS23, 7875, E5-2600
- ▶ IBM BladeCenter HS23, 7875, E5-2600 v2
- ▶ IBM BladeCenter HX5, 7873
- ▶ IBM BladeCenter PS700, 8406-70Y
- ▶ IBM BladeCenter PS701, 8406-71Y
- ▶ IBM BladeCenter PS702, 8406-71Y +FC 8358
- ▶ IBM BladeCenter PS703, 7891-73X
- ▶ IBM BladeCenter PS704, 7891-74X

© Copyright International Business Machines Corporation 2007 - 2013. All rights reserved.

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

Contents

Notices	ix
Trademarks	x
Preface	xi
Authors	xi
Now you can become a published author, too!	xii
Comments welcome	xii
Stay connected to IBM Redbooks	xii
Summary of changes	xiii
November 2013, Twelfth Edition	xiii
Chapter 1. Product overview	1
1.1 Support matrixes	2
1.1.1 Chassis-server compatibility	2
1.1.2 Chassis-I/O module compatibility	3
1.1.3 Server-expansion card compatibility	4
1.1.4 I/O module-expansion card compatibility	6
1.1.5 Other compatibility tables	10
1.1.6 Supported operating systems	10
1.2 BladeCenter chassis	16
1.2.1 BladeCenter E	16
1.2.2 BladeCenter H	18
1.2.3 BladeCenter HT	20
1.2.4 BladeCenter S	23
1.3 Blade servers	25
1.3.1 BladeCenter HS12	25
1.3.2 BladeCenter HS22	26
1.3.3 BladeCenter HS23 (E5-2600)	27
1.3.4 BladeCenter HS23 (E5-2600 v2)	28
1.3.5 BladeCenter HS23E	29
1.3.6 BladeCenter HX5	30
1.3.7 IBM BladeCenter PS700, PS701, and PS702	31
1.3.8 BladeCenter PS703 and PS704	34
Chapter 2. IBM BladeCenter chassis and infrastructure configuration	37
2.1 IBM BladeCenter E	38
2.1.1 IBM BladeCenter E product information	38
2.1.2 BladeCenter E I/O topology	41
2.1.3 Blower module for BladeCenter E	43
2.1.4 Acoustic Attenuation Module	44
2.1.5 Power modules for BladeCenter E	45
2.1.6 Product publications	46
2.2 IBM BladeCenter H	46
2.2.1 IBM BladeCenter H product information	47
2.2.2 BladeCenter H I/O topology	50
2.2.3 Blower modules for BladeCenter H	52
2.2.4 Power modules for IBM BladeCenter H	53
2.2.5 Serial Port Breakout Cable	56

2.2.6	Product publications	56
2.3	IBM BladeCenter HT	56
2.3.1	BladeCenter HT product information	57
2.3.2	BladeCenter HT interposers	61
2.3.3	BladeCenter HT I/O topology	61
2.3.4	Fan modules for BladeCenter HT	64
2.3.5	Power modules for BladeCenter HT	65
2.3.6	Serial Port Breakout Cable	66
2.3.7	Product publications	67
2.4	IBM BladeCenter S	67
2.4.1	BladeCenter S product information	67
2.4.2	BladeCenter S I/O topology	71
2.4.3	Fan modules for BladeCenter S	73
2.4.4	Power modules for BladeCenter S	73
2.4.5	BladeCenter S disk storage modules	75
2.4.6	BladeCenter S Serial Pass-thru Module	78
2.4.7	BladeCenter S Office Enablement Kit	78
2.4.8	Product publications	79
2.5	Chassis I/O bay and expansion card port mappings	80
2.6	I/O modules overview	83
2.7	Ethernet Switch Modules	84
2.7.1	Cisco Catalyst Switch Module 3012	88
2.7.2	Cisco Catalyst Switch Module 3110G	96
2.7.3	Cisco Catalyst Switch Module 3110X	104
2.7.4	Cisco Nexus 4001I Switch Module	115
2.7.5	IBM Server Connectivity Module	121
2.7.6	IBM Layer 2/3 Copper GbE Ethernet Switch Module	124
2.7.7	IBM Layer 2/3 Fiber GbE Ethernet Switch Module	127
2.7.8	IBM Layer 2-7 Gigabit Ethernet Switch Module	131
2.7.9	IBM 1/10 Gb Uplink Ethernet Switch Module	134
2.7.10	IBM Virtual Fabric 10 Gb Switch Module	139
2.7.11	Brocade Converged 10 GbE Switch Module	145
2.8	InfiniBand switch module	150
2.8.1	Voltaire 40 Gb InfiniBand Switch Module	151
2.9	Fibre Channel switch modules	153
2.9.1	Brocade Enterprise 20-port, 20-port, and 10-port 8 Gb SAN Switch Modules	154
2.9.2	Cisco Systems 4 Gb 20-port and 10-port Fibre Channel Switch Modules	158
2.9.3	QLogic 20-port 8 Gb and 4/8 Gb SAN Switch Modules	160
2.9.4	QLogic 8 Gb and 4/8 Gb Intelligent Pass-thru Modules	163
2.9.5	QLogic Virtual Fabric Extension Module	165
2.10	SAS I/O modules	168
2.10.1	SAS Connectivity Module	169
2.10.2	SAS RAID Controller Module	171
2.11	Pass-through and interconnect modules	173
2.11.1	Intelligent Copper Pass-thru Module for IBM BladeCenter	174
2.11.2	Multi-Switch Interconnect Module	176
2.11.3	Multi-Switch Interconnect Module for BladeCenter HT	181
2.11.4	10 Gb Ethernet Pass-thru Module	184
2.12	Advanced Management Module	185
2.12.1	Advanced Management Module	187
2.13	Direct serial connections	190
2.13.1	Serial Port Breakout Cable	190
2.13.2	BladeCenter S Serial Pass-thru Module	191

2.14 Installation and physical site plans	193
2.14.1 Chassis dimensions and physical planning.	193
2.14.2 Rack installation	196
2.14.3 Power considerations	197
2.14.4 Cooling considerations	201
2.14.5 Power cable considerations	204
2.14.6 Console Managers and SOL console connectivity	206
Chapter 3. Blade server hardware configuration	211
3.1 Expansion cards	212
3.1.1 Form factors	213
3.1.2 Comparing the features of the 10 Gb Ethernet expansion cards	220
3.1.3 Ethernet Expansion Card (CIOv)	221
3.1.4 The 2/4 Port Ethernet Expansion Card (CFFh).	222
3.1.5 Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh).	223
3.1.6 Broadcom 10 Gb Gen 2 4-port Ethernet Expansion Card (CFFh)	225
3.1.7 Broadcom 2-port 10 Gb Virtual Fabric Adapter for IBM BladeCenter	226
3.1.8 Emulex 10 GbE Virtual Fabric Adapter II (CFFh)/Advanced II (CFFh)	229
3.1.9 Emulex 10 GbE Virtual Fabric Adapter II/Advanced II for HS23.	231
3.1.10 QLogic 10Gb Virtual Fabric Adapter / QLogic 10Gb Virtual Fabric Converged Network Adapter (CFFh).	234
3.1.11 Intel 10 Gb 2-port Ethernet Expansion Card (CFFh).	237
3.1.12 QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)	238
3.1.13 QLogic 4 Gb Fibre Channel Expansion Card (CIOv)	239
3.1.14 QLogic 8 Gb Fibre Channel Expansion Card (CIOv)	240
3.1.15 Emulex 8 Gb Fibre Channel Expansion Card (CIOv)	242
3.1.16 QLogic 2-port 10 Gb Converged Network Adapter (CFFh).	243
3.1.17 Brocade 2-port 10 GbE Converged Network Adapter (CFFh)	244
3.1.18 The 2-port 40 Gb InfiniBand Expansion Card (CFFh).	246
3.1.19 SAS Expansion Card (CFFv)	248
3.1.20 SAS Connectivity Card (CIOv)	249
3.1.21 ServeRAID H1135 (CIOv) Controller	250
3.2 IBM BladeCenter HS12.	252
3.2.1 Processor	254
3.2.2 Memory	255
3.2.3 Onboard network adapters	256
3.2.4 Integrated systems management processor.	256
3.2.5 Local storage options	257
3.2.6 I/O expansion options	259
3.3 IBM BladeCenter HS22.	260
3.3.1 Model information	262
3.3.2 Chassis support	263
3.3.3 Processor	264
3.3.4 Memory	269
3.3.5 Standard onboard features	275
3.3.6 Local storage	276
3.3.7 Integrated virtualization.	279
3.3.8 I/O expansion	280
3.4 BladeCenter HS23 (E5-2600)	283
3.4.1 Overview	284
3.4.2 Model information	286
3.4.3 Chassis support	288
3.4.4 Processor	289

3.4.5	Memory	291
3.4.6	Standard onboard features	295
3.4.7	Local storage	296
3.4.8	Integrated virtualization	299
3.4.9	I/O expansion	299
3.5	BladeCenter HS23 (E5-2600 v2)	304
3.5.1	Overview	304
3.5.2	Model information	306
3.5.3	Chassis support	307
3.5.4	Processor	308
3.5.5	Memory	310
3.5.6	Standard onboard features	314
3.5.7	Local storage	316
3.5.8	Integrated virtualization	318
3.5.9	I/O expansion	319
3.6	BladeCenter HS23E	324
3.6.1	Overview	324
3.6.2	Model information	326
3.6.3	Chassis support	327
3.6.4	Processor	328
3.6.5	Memory	330
3.6.6	Standard onboard features	334
3.6.7	Local storage	336
3.6.8	Integrated virtualization	338
3.6.9	I/O expansion	339
3.7	IBM BladeCenter HX5	341
3.7.1	Introduction	341
3.7.2	Models	344
3.7.3	System architecture	346
3.7.4	Chassis support	349
3.7.5	IBM MAX5 V2 for BladeCenter	350
3.7.6	Scalability	351
3.7.7	Processors	355
3.7.8	Memory	358
3.7.9	Storage	364
3.7.10	I/O expansion cards	367
3.7.11	Standard onboard features	370
3.7.12	Integrated virtualization	371
3.7.13	Partitioning capabilities	372
3.8	IBM BladeCenter PS700, PS701, and PS702	377
3.8.1	Overview	377
3.8.2	Supported BladeCenter chassis	383
3.8.3	Minimum features	385
3.8.4	Processor	386
3.8.5	Memory	386
3.8.6	I/O	387
3.8.7	Storage	388
3.8.8	Standard onboard features	388
3.9	IBM BladeCenter PS703 and PS704	389
3.9.1	Overview	390
3.9.2	Supported BladeCenter chassis	395
3.9.3	Minimum features	396
3.9.4	Processor	396

3.9.5 Memory	397
3.9.6 I/O	398
3.9.7 Storage	398
3.9.8 Standard onboard features	399
3.10 Expansion blades	400
3.10.1 PCI Express I/O Expansion Unit	401
3.10.2 PCI Express Gen 2 Expansion Blade/Expansion Blade II	403
3.10.3 IBM BladeCenter GPU Expansion Blade	408
Chapter 4. Network integration	413
4.1 Ethernet switch module selection	414
4.2 VLAN	415
4.3 High availability and redundancy	416
4.3.1 Redundant network topologies	417
4.3.2 Spanning Tree Protocol	417
4.3.3 Trunk Failover with NIC teaming	418
4.3.4 Virtual Router Redundancy Protocol	419
4.3.5 Routing protocols	420
4.4 Performance	420
4.4.1 Link aggregation	420
4.4.2 Jumbo frames	421
4.4.3 NIC teaming	421
4.4.4 Server Load Balancing	421
4.5 Systems management	422
4.5.1 Management interfaces	422
4.5.2 Management network topology	424
4.5.3 Management protocols and tools	426
4.6 IBM switch stacking	430
4.7 Virtual Blade Switch technology	432
4.8 IBM Fabric Manager	433
4.9 BladeCenter Virtual Fabric vNIC solution	439
4.9.1 IBM Virtual Fabric Mode	440
4.9.2 Switch Independent Mode	444
Chapter 5. Storage integration	449
5.1 External storage for IBM BladeCenter	450
5.2 Fibre Channel storage system interoperability	451
5.2.1 Fibre Channel storage support for x86 blade servers	452
5.2.2 Fibre Channel storage support for Power Systems blade servers	453
5.2.3 Fibre Channel switch selection, fabric interoperability rules, NPIV	453
5.2.4 NPIV support on Power Systems blade servers	457
5.3 iSCSI system storage interoperability	457
5.3.1 Hardware-base iSCSI initiator support	459
5.3.2 iSCSI SAN boot	459
5.4 SAS storage system interoperability	461
5.5 High availability and redundancy	463
5.6 Performance	464
5.7 Backup solutions	464
5.7.1 A dedicated server for centralized LAN backup	465
5.7.2 LAN-free backup for blades	466

5.8 SAS tape storage for IBM BladeCenter	468
5.9 External storage expansion (JBOD) attachments	471
5.10 IBM BladeCenter S integrated storage	475
5.10.1 Basic local storage using the SAS Connectivity Module	475
5.10.2 Advanced shared storage using the SAS RAID Controller Module	479
5.11 Converged networks	484
5.11.1 FCoCEE solution using Virtual Fabric Switches and Virtual Fabric Extension Modules	486
5.11.2 FCoCEE solution with the Brocade Converged 10 GbE Switch	488
5.11.3 Emulex adapters in x86 servers	490
5.11.4 Brocade and QLogic adapters in x86 servers	491
5.11.5 Power Systems servers	493
Chapter 6. IBM zEnterprise Extension Model	495
6.1 zBX hardware description	496
6.2 zBX racks	496
6.3 Top-of-rack switches	498
6.4 zBX BladeCenter chassis	499
6.5 zBX blades	500
6.6 Power distribution unit	502
6.7 zBX management	502
6.8 zBX connectivity	503
Abbreviations and acronyms	511
Related publications	515
IBM Redbooks publications	515
Product publications	516
Online resources	516
Help from IBM	516

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 about 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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:

Active Memory™	POWER®	System p®
AIX®	Power Systems™	System Storage®
AIX 5L™	POWER7®	System x®
BladeCenter®	PowerHA®	System z®
Calibrated Vectored Cooling™	PowerVM®	SystemMirror®
DB2®	PureFlex™	Systems Director VMControl™
DS8000®	RackSwitch™	Tivoli®
EnergyScale™	Redbooks®	VMready®
Enterprise Workload Manager™	Redpaper™	WebSphere®
i5/OS™	Redbooks (logo)  ®	Workload Partitions Manager™
IBM®	RETAIN®	XIV®
IBM Flex System™	ServerProven®	z/OS®
NMotion®	System i®	zEnterprise®

The following terms are trademarks of other companies:

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

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

LTO, Ultrium, the LTO Logo and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

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

Preface

Blade servers captured industry focus because of their modular design, which can reduce costs with a more efficient usage of valuable floor space, and their simplified management, which can speed up such tasks as deploying, reprovisioning, updating, and troubleshooting hundreds of blade servers. By integrating resources and sharing key components, costs are reduced and availability is increased.

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

This eleventh edition includes new products that were announced in the first and second quarters of 2013. This book is for clients, IBM Business Partners, and IBM employees who want to know the details about IBM BladeCenter chassis, servers, and options.

Authors

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



Ilya Krutov is a Project Leader at the ITSO Center in Raleigh and has been with IBM since 1998. Before joining the ITSO, Ilya served in IBM as a Team Leader, Portfolio Manager, Brand Manager, IT Specialist, and Certified Instructor. Ilya has expertise in IBM System x®, BladeCenter, and PureFlex™ System products, server operating systems, and networking solutions. He has authored over 130 books, papers, and guides. He has a Bachelor Degree in Computer Engineering from the Moscow Engineering and Physics Institute.



Alexander Grechnev is a Techline System x Specialist at IBM East Europe/Asia and has been with IBM since 2006. He is responsible for pre-sales technical support of IBM employees and IBM Business Partners, and specializes in System x and BladeCenter products. He graduated from the Moscow State Technical University n.a. Bauman in 2002.

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

Kevin Barnes, David Watts, Tamikia Barrow, Rufus Credle, Linda Robinson, Shari Deiana
International Technical Support Organization

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author - all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

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 publications in one of the following ways:

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

ibm.com/redbooks

- ▶ Send your comments in an email 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

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBM-Redbooks>

- ▶ Follow us on twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

<http://www.redbooks.ibm.com/rss.html>

Summary of changes

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

Summary of Changes
for SG24-7523-11
for IBM BladeCenter Products and Technology
as created or updated on November 26, 2013.

November 2013, Twelfth Edition

This edition includes the following new and changed information.

New products and information

- ▶ BladeCenter HS23 (7875, E5-2600 v2)
- ▶ IBM Hybrid hard disk drives
- ▶ S3500 and S3700 solid-state drives for IBM Systems

Updates to existing information

- ▶ Operating system compatibility
- ▶ Storage compatibility matrixes
- ▶ Chassis support, I/O module, and I/O adapter compatibility matrixes

Consult the previous editions of this book for details about removed products. Click the **Additional Material** link at the following web location to download the PDF:

<http://www.redbooks.ibm.com/abstracts/sg247523.html>



Product overview

Blade servers are thin servers that are inserted into a single rack-mounted chassis that 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 the 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" on page 2
- ▶ 1.2, "BladeCenter chassis" on page 16
- ▶ 1.3, "Blade servers" on page 25

1.1 Support matrixes

This section provides compatibility matrixes. It contains the following sections:

- ▶ 1.1.1, “Chassis-server compatibility” on page 2
- ▶ 1.1.2, “Chassis-I/O module compatibility” on page 3
- ▶ 1.1.3, “Server-expansion card compatibility” on page 4
- ▶ 1.1.4, “I/O module-expansion card compatibility” on page 6
- ▶ 1.1.5, “Other compatibility tables” on page 10
- ▶ 1.1.6, “Supported operating systems” on page 10

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	8028	1 slot	Yes	Yes	Yes	Yes	Yes	Yes	Yes
HS22	7870	1 slot	Yes	Yes ^a	No	No	Yes ^a	Yes ^a	Yes ^a
HS22V	7871	1 slot	Yes	Yes ^a	No	No	Yes ^a	Yes ^a	Yes ^a
HS23	7875, E5-2600	1 slot	Yes	Yes ^a	No	No	Yes ^a	Yes ^a	Yes ^a
HS23	7875, E5-2600 v2	1 slot	Yes	Yes ^a	No	No	Yes ^a	Yes ^a	Yes ^a
HS23E	8038	1 slot	Yes	Yes ^a	No	No	Yes	Yes ^a	Yes ^a
HX5	7872	1 or 2 slots	Yes ^b	No	No	No	Yes ^{b,c}	Yes ^b	Yes ^b
HX5	7873	1 or 2 slots	Yes ^b	No	No	No	Yes ^{b,c}	Yes ^b	Yes ^b
PS700	8406-70x	1 slot	Yes	Yes ^d	No	No	Yes	Yes	Yes
PS701	8406-71x	1 slot	Yes	No	No	No	Yes	Yes	Yes
PS702	8406-71x	2 slots	Yes	No	No	No	Yes	Yes	Yes
PS703	7891-73x	1 slot	Yes	No	No	No	Yes	Yes	Yes
PS704	7891-74x	2 slots	Yes	No	No	No	Yes	Yes	Yes

a. Certain rules apply to the installation of this blade server in the chassis, depending on the power supplies installed in the chassis and the TDP watts value of the processors in the servers. See the “Chassis support” subsection of the particular blade for more information.

b. Certain rules apply to the installation of HX5 servers with 130 W processors in the chassis. See 3.7.4, “Chassis support” on page 349.

c. The HX5 installed in the BladeCenter H requires BladeCenter H Enhanced Cooling Modules. See 3.7.4, “Chassis support” on page 349.

d. Only specific models of the BladeCenter E support the PS700. See Table 3-130 on page 384.

Expansion blades support matrix: The support matrix for the expansion blades is shown in 3.10, “Expansion blades” on page 400.

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 (x-config/e-config)	BC S	BC E	BC T	BC H	BC HT	MSIM	MSIM-HT	Page
Ethernet switch modules										
Cisco Catalyst Switch Module 3110G ^a	00Y3254 ^b	A3FD/none	N	Y	Y	Y	Y	Y	Y	88
Cisco Catalyst Switch Module 3110X ^a	00Y3250 ^c	A3FC/none	N	Y	Y	Y	Y	Y	Y	104
Cisco Catalyst Switch Module 3012 ^a	46C9272 ^d	A3FE/none	Y	Y	Y	Y	Y	Y	Y	121
Cisco Nexus 4001I Switch Module	46C9270 ^e	A3FF/none	N	N	N	Y	Y	N	N	115
IBM Server Connectivity Module ^a	39Y9324	1484/3220	Y	Y	Y	Y	Y	Y	N	121
IBM L2/3 Copper GbE Switch Module	32R1860	1495/3212	Y	Y	Y	Y	Y	Y	Y	124
IBM L2/3 Fiber GbE Switch Module	32R1861	1496/3213	Y	Y	Y	Y	Y	Y	Y	127
IBM L2-7Gb Ethernet Switch Module	32R1859	1494/3211	Y	Y	Y	Y	Y	N	N	131
IBM 1/10Gb Uplink ESM ^a	44W4404	1590/1590	Y	Y	Y	Y	Y	Y	Y	134
IBM Virtual Fabric 10Gb Switch Module	46C7191	1639/3248	N	N	N	Y	Y	N	N	139
Brocade Converged 10GbE Switch Module	69Y1909	7656/none	N	N	N	Y	Y	N	N	145
InfiniBand Modules										
Voltaire 40Gb InfiniBand Switch Module	46M6005	0057/3204	N	N	N	Y	N	N	N	151
Fibre Channel Switch Modules										
Brocade Enterprise 20-port 8Gb SAN SM ^a	42C1828	5764/none	N	Y	N	Y	Y	Y	N	154
Brocade 20-port 8Gb SAN Switch Module ^a	44X1920	5481/5869	N	Y	N	Y	Y	Y	Y	154
Brocade 10-port 8Gb SAN Switch Module ^a	44X1921	5483/5045	N	Y	N	Y	Y	Y	Y	154
Cisco 4Gb 20 port FC Switch Module ^a	44E5696 ^f	A3FH/none	N	Y	Y	Y	Y	Y	N	158
Cisco 4Gb 10 port FC Switch Module ^a	44E5692 ^g	A3FG/none	Y	Y	Y	Y	Y	Y	N	158
QLogic 20-Port 8Gb SAN Switch Module ^a	44X1905	5478/3284	N	Y	Y	Y	Y	Y	Y	160
QLogic 20-Port 4/8Gb SAN Switch Module ^a	88Y6406	A24C/none	Y	Y	Y	Y	Y	Y	Y	160
QLogic 8Gb Intelligent Pass-Thru Module ^a	44X1907	5482/5449	Y	Y	Y	Y	Y	Y	Y	163
QLogic 4/8Gb Intelligent Pass-Thru Module ^a	88Y6410	A24D/none	Y	Y	Y	Y	Y	Y	Y	163
QLogic Virtual Fabric Extension Module ^h	46M6172	4799/none	N	N	N	Y	N	N	N	165
SAS modules										
SAS Connectivity Module ^a	39Y9195	2980/3267	Y	Y	Y	Y	Y	N	N	168
SAS RAID Controller Module	43W3584	3734/none	Y	N	N	N	N	N	N	171

I/O module	Part Number	Feature code (x-config/e-config)	BC S	BC E	BC T	BC H	BC HT	MSIM	MSIM-HT	Page
Pass-through and interconnect modules										
Intelligent Copper Pass-Thru Module ^a	44W4483	5452/5452	Y	Y	Y	Y	Y	Y	N	174
QLogic 8Gb Intelligent Pass-Thru Module ^a	44X1907	5482/5449	Y	Y	Y	Y	Y	Y	N	163
QLogic 4/8Gb Intelligent Pass-Thru Module ^a	88Y6410	A24D/none	Y	Y	Y	Y	Y	Y	Y	163
IBM Multi-Switch Interconnect Module	39Y9314	1465/3239	N	N	N	Y	N	N	N	176
Multi-Switch Interconnect Module for BC HT	44R5913	5491/none	N	N	N	N	Y	N	N	181
10Gb Ethernet Pass-Thru Module	46M6181	1641/5412	N	N	N	Y	Y	N	N	184

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

b. Replaces 41Y8523.

c. Replaces 41Y8522.

d. Replaces 43W4395.

e. Replaces 46M6071.

f. Replaces 39Y9280.

g. Replaces 39Y9284.

h. The QLogic Virtual Fabric Extension Module is only supported with the IBM Virtual Fabric 10 Gb Switch Module.

1.1.3 Server-expansion card compatibility

Table 1-3 lists the available expansion cards that can be installed in blade servers and also indicates in which specific servers they are a supported option. For details about an expansion card, go to the page referenced in the Page column.

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

Expansion cards (Y = supported, N = not supported)	Part number	Feature code ^a	HS12 (8028)	HS22 (7870)	HS22V (7871)	HS23 (7875, E5-2600)	HS23 (7875, E5-2600 v2)	HS23E (8038)	HX5 (7872)	HX5 (7873)	PS700/PS701/PS702 (8406)	PS703/PS704 (7891)	Page
Ethernet expansion cards													
Ethernet Expansion Card (CIOv)	44W4475	1039	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	221
2/4 Port Ethernet Exp Card (CFFh)	44W4479	5476	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	222
Mellanox 2-port 10Gb Ethernet (CFFh)	90Y3570	A1NW	N	Y	Y	Y	Y	N	Y	N	N	N	223
Broadcom 10Gb Gen2 4-pt Ethernet (CFFh)	46M6164	0098	Y	Y	Y	Y	Y	Y	Y	Y	N	N	225
Broadcom 2-pt 10Gb Virt. Fbr. Adpt. (CFFh)	81Y3133	A1QR	N	Y	Y	Y	Y	Y	Y	Y	N	N	226
QLogic Eth. and 8 Gb FC (CFFh)	00Y3270 ^b	A3JC	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	238
Emulex Virtual Fabric Adapter II	00Y3266	A3NV	N	Y	Y	N	N	Y	Y	Y	N	N	229

Expansion cards (Y = supported, N = not supported)	Part number	Feature code ^a	HS12 (8028)	HS22 (7870)	HS22V (7871)	HS23 (7875, E5-2600)	HS23 (7875, E5-2600 v2)	HS23E (8038)	HX5 (7872)	HX5 (7873)	PS700/PS701/PS702 (8406)	PS703/PS704 (7891)	Page
Emulex Virtual Fabric Adapter Advanced II	00Y3264	A3NW	N	Y	Y	N	N	Y	Y	Y	N	N	229
Emulex 10GbE VFA II for HS23	81Y3120	A287	N	N	N	Y	Y	N	N	N	N	N	231
Emulex 10GbE VFA Advanced II for HS23	90Y9332	A2ZN	N	N	N	Y	Y	N	N	N	N	N	231
QLogic 2-port 10Gb CNA (CFFh)	00Y3280 ^c	A3JB	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	243
QLogic 10Gb Virtual Fabric Adapter (CFFh)	00Y3332	A4AC	N	N	N	Y	Y	Y	Y	Y	N	N	234
QLogic 10Gb Virtual Fabric CNA (CFFh)	00Y5618	A4AD	N	N	N	Y	Y	Y	Y	Y	N	N	234
Intel 10Gb 2-port Ethernet (CFFh)	42C1810	3593	N	Y	Y	Y	Y	Y	Y	Y	N	N	229
Brocade 2-port 10Gb CNA (CFFh)	81Y1650	5437	N	Y	Y	Y	Y	Y	Y	Y	N	N	244
Fibre Channel expansion cards													
Emulex 8Gb FC (CIOv)	46M6140	3598	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	242
QLogic 4Gb FC (CIOv)	46M6065	3594	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	239
QLogic 8Gb FC (CIOv)	44X1945	1462	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	240
QLogic Ethernet and 8Gb FC (CFFh)	00Y3270 ^b	A3JC	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	238
InfiniBand expansion cards													
2-port 40Gb InfiniBand (CFFh)	46M6001	0056	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	246
SAS expansion cards													
SAS Expansion Card (CFFv)	44E5688 ^d	A3J9	Y	N	N	N	N	N	N	N	N	N	248
SAS Connectivity Card (CIOv)	43W4068	1041	N	Y	Y	Y	Y	N	Y	Y	Y	Y	249
ServeRAID H1135 Controller (CIOv)	90Y4750	A1XJ	N	N	N	N	N	Y	N	N	N	N	250
Interposer cards													
10Gb Interposer Card for HS23	94Y8550	A244	N	N	N	Y	Y	N	N	N	N	N	300

a. Feature codes listed here are for the System x sales channel (x-config). For IBM System p® users (e-config), see the PS700/PS701/PS702, and PS703/704 sections in Chapter 3, "Blade server hardware configuration" on page 211.

b. Replaces 44X1940.

c. Replaces 42C1830.

d. Replaces 39Y9190.

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

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

- ▶ IBM ServerProven®:
<http://www.ibm.com/systems/info/x86servers/serverproven/compat/us/>
- ▶ IBM System x Configuration and Options Guide:
<http://www.ibm.com/systems/xbc/cog/>
- ▶ IBM BladeCenter Interoperability Guide:
<http://www.redbooks.ibm.com/big>

1.1.4 I/O module-expansion card compatibility

Compatibility between I/O modules in the chassis and expansion cards in the servers are as follows:

- ▶ Ethernet compatibility: See Table 1-4 and Table 1-5 on page 8.
- ▶ Fibre Channel compatibility: See Table 1-6 on page 9.
- ▶ InfiniBand compatibility: See Table 1-7 on page 9.
- ▶ SAS compatibility: See Table 1-8 on page 10.

FCoE and vNIC planning details:

- ▶ When you plan for FCoE, see 5.11, “Converged networks” on page 484.
- ▶ When you plan for vNIC, see 4.9, “BladeCenter Virtual Fabric vNIC solution” on page 439.

Table 1-4 1 Gb and 10 Gb Ethernet Switch Module: Expansion card compatibility

	Ethernet Expansion Cards →	1 Gb to the blade			10 Gb to the blade				
Part number	Ethernet Switch Module	44W4475 Ethernet Expansion Card (CIOv)	44W4479 2/4 Port Ethernet (CFFh)	00Y3270 ^a QLlogic Eth. and 8 GB FC (CFFH)	46M6164 Broadcom 10 Gb Gen 2 4-pt Eth. (CFFh)	46M6168 Broadcom 10 Gb Gen 2 2-pt Eth. (CFFh)	81Y3133 Broadcom 2-port 10 Gb VFA (CFFh)	90Y3570 Mellanox 2-port 10 Gb Eth. (CFFh)	42C1810 Intel 10 Gb 2-port Eth. (CFFh)
Gigabit Ethernet Switch Modules (External ports operate at 1 Gbps)									
32R1859	IBM Layer 2-7Gb Ethernet Switch	Y	N	N	N	N	N	N	N
32R1860	IBM Layer 2/3 Copper Gb Switch	Y	Y ^b	Y ^b	N	N	N	N	N
32R1861	IBM Layer 2/3 Fiber Gb Switch	Y	Y ^b	Y ^b	N	N	N	N	N
39Y9324	IBM Server Connectivity Module	Y	Y ^c	Y ^c	N	N	N	N	N
00Y3250 ^d	Cisco Catalyst Switch 3110X	Y	Y ^e	Y ^e	N	N	N	N	N

Part number	Ethernet Expansion Cards →	1 Gb to the blade			10 Gb to the blade				
		44W4475 Ethernet Expansion Card (Clov)	44W4479 2/4 Port Ethernet (CFFh)	00Y3270 ^a QLogic Eth. and 8 GB FC (CFFh)	46M6164 Broadcom 10 Gb Gen 2 4-pt Eth. (CFFh)	46M6168 Broadcom 10 Gb Gen 2 2-pt Eth. (CFFh)	81Y3133 Broadcom 2-port 10 Gb VFA (CFFh)	90Y3570 Mellanox 2-port 10 Gb Eth. (CFFh)	42C1810 Intel 10 Gb 2-port Eth. (CFFh)
00Y3254 ^f	Cisco Catalyst Switch 3110G	Y	Y ^e	Y ^e	N	N	N	N	N
46C9272 ^g	Cisco Catalyst Switch 3012	Y	Y ^b	Y ^b	N	N	N	N	N
44W4483	Intelligent Copper Pass-Thru Module	Y	Y ^c	Y ^c	N	N	N	N	N
10 Gb Ethernet Switch Modules (External ports operate at 10 Gbps)									
44W4404	IBM 1/10Gb Uplink Switch	Y	Y ^b	Y ^b	N	N	N	N	N
46C7191	IBM Virtual Fabric 10Gb Switch	N	Y	N	Y	Y	Y	Y	N
69Y1909	Brocade Converged 10GbE Switch	N	N	N	N	N	N	N	N
46M6181	10Gb Ethernet Pass-Thru Module	N	N	N	Y	Y	Y	N	N
46C9270 ^h	Cisco Nexus 4001I Switch Module	N	Y	N	Y	Y	Y	Y	Y ⁱ

a. Replaces 44X1940.

b. Supported in BladeCenter S in the I/O bay 2 or in BladeCenter H and HT with the MSIM or MSIM-HT installed.

c. Supported in BladeCenter S in the I/O bay 2 or in BladeCenter H with the MSIM installed.

d. Replaces 41Y8522.

e. Supported in BladeCenter H and HT with the MSIM or MSIM-HT installed.

f. Replaces 41Y8523.

g. Replaces 43W4395.

h. Replaces 46M6071.

i. Force switch port to 10 Gb fixed, WoL not supported.

Table 1-5 1 Gb and 10 Gb Ethernet Switch Module: Expansion card compatibility (Continued)

Ethernet Expansion Cards →		10 Gb to the blade								
Part number	Ethernet Switch Module	00Y3266 Emulex VFA II (CFFh)	00Y3264 Emulex VFA Advanced II (CFFh)	81Y3120 Emulex VFA II (CFFh) for HS23	90Y9332 Emulex VFA Adv. II (CFFh) for HS23	94Y8550 10 Gb Interposer Card for HS23 (uses HS23 Integrated Virtual Fabric LOM)	81Y1650 Brocade 2-port 10 GbE CNA (CFFh)	00Y3280 ^a QLogic 2-port 10 Gb CNA (CFFh)	00Y3332 QLogic 10Gb VFA (CFFh)	00Y5618 QLogic 10Gb CNA (CFFh)
Gigabit Ethernet Switch Modules (External ports operate at 1 Gbps)										
32R1859	IBM Layer 2-7Gb Ethernet Switch	N	N	N	N	N	N	N	N	N
32R1860	IBM Layer 2/3 Copper Gb Switch	N	N	N	N	N	N	N	N	N
32R1861	IBM Layer 2/3 Fiber Gb Switch	N	N	N	N	N	N	N	N	N
39Y9324	IBM Server Connectivity Module	N	N	N	N	Y ^b	N	N	N	N
00Y3250 ^c	Cisco Catalyst Switch 3110X	N	N	N	N	N	N	N	N	N
00Y3254 ^d	Cisco Catalyst Switch 3110G	N	N	N	N	N	N	N	N	N
46C9272 ^e	Cisco Catalyst Switch 3012	N	N	N	N	Y ^b	N	N	N	N
44W4483	Intelligent Copper Pass-Thru Module	N	N	N	N	Y ^b	N	N	N	N
10 Gb Ethernet Switch Modules (External ports operate at 10 Gbps)										
44W4404	IBM 1/10Gb Uplink Switch	N	N	N	N	Y ^b	N	N	N	N
46C7191	IBM Virtual Fabric 10Gb Switch	Y	Y	Y	Y	Y	Y	Y	Y	Y
69Y1909	Brocade Converged 10GbE Switch	Y	Y	N	N	Y	Y	Y	Y	Y
46M6181	10Gb Ethernet Pass-Thru Module	Y	Y	Y	Y	Y	N	Y	Y	Y
46C9270 ^f	Cisco Nexus 4001I Switch Module	Y	Y	Y	Y	Y	N	Y	Y	Y

a. Replaces 42C1830.

b. Requires MSIM in BladeCenter H or MSIM-HT in BladeCenter HT.

c. Replaces 41Y8522.

d. Replaces 41Y8523.

e. Replaces 43W4395.

f. Replaces 46M6071.

Table 1-6 4 Gb and 8 Gb Fibre Channel Switch Module: Expansion card compatibility

	Fibre Channel Expansion Cards →	4 Gb Card	8 Gb Cards		
Part number	I/O module	46M6065 QLogic 4 Gb FC (CIOv)	00Y3270 ^a QLogic Eth. and 8 Gb FC (CFFh)	44X1945 QLogic 8 Gb FC (CIOv)	46M6140 Emulex 8 Gb FC (CIOv)
4 Gb Fibre Channel Switch Modules					
44E5696 ^b	Cisco Systems 4Gb 20 port FC Switch Module	Y	Y	Y	Y
44E5692 ^c	Cisco Systems 4Gb 10 port FC Switch Module	Y	Y	Y	Y
8 Gb Fibre Channel Switch Modules					
44X1905	QLogic 20-port 8Gb SAN Switch Module	Y	Y	Y	Y
44X1907	QLogic 8Gb Intelligent Pass-Thru Module	Y	Y	Y	Y
42C1828	Brocade Enterprise 20-port Switch Module	Y	Y	Y	Y
44X1920	Brocade 20-port 8Gb SAN Switch Module	Y	Y	Y	Y
44X1921	Brocade 10-port 8Gb SAN Switch Module	Y	Y	Y	Y
88Y6406	QLogic 20-Port 4/8Gb SAN Switch Module	Y	Y	Y	Y
88Y6410	QLogic 4/8Gb Intelligent Pass-Thru Module	Y	Y	Y	Y

a. Replaces 44X1940.

b. Replaces 39Y9280.

c. Replaces 39Y9284.

Table 1-7 InfiniBand Switch Module: Expansion card compatibility

	InfiniBand Expansion Cards →	QDR
Part number	I/O Module	46M6001 2-pt 40 Gb InfiniBand (CFFh)
46M6005	Voltaire 40Gb InfiniBand Switch Module	Y

Table 1-8 SAS I/O Module: Expansion card compatibility

	SAS Expansion Cards →	CFFv	CIOv	
Part number	I/O Module	44E5688 ^a SAS Exp. Card (CFFv)	43W4068 SAS Connectivity (CIOv)	90Y4750 ServeRAID H1135 (CIOv)
39Y9195	IBM BladeCenter SAS Connectivity Module	Y	Y	Y
43W3584	SAS RAID Controller Module for BladeCenter S Chassis	Y	Y	Y

a. Replaces 39Y9190.

1.1.5 Other compatibility tables

Other compatibility tables in this book are as follows:

- ▶ Expansion blades:
 - Expansion blades and supported server blades: Table 3-142 on page 400
- ▶ Networking:
 - IBM Fabric Manager compatibility: Table 4-7 on page 437
 - BladeCenter Virtual Fabric vNIC solution support matrix: Table 4-9 on page 443
- ▶ Storage:
 - Fibre Channel storage compatibility matrix: Table 5-1 on page 452 and Table 5-2 on page 453.
 - iSCSI storage compatibility matrix: Table 5-5 on page 459
 - SAS storage compatibility matrix: Table 5-7 on page 462
 - IBM SAN boot support with built-in software iSCSI initiator: Table 5-6 on page 460
 - Entry IBM Tape Storage - SAS compatibility matrix: Table 5-8 on page 469
 - IBM System x Tape Drives - SAS compatibility matrix: Table 5-9 on page 470
 - BladeCenter S SAS RAID Controller compatibility matrix: Table 5-18 on page 483
 - FCoCEE storage compatibility matrix: Table 5-20 on page 490, Table 5-21 on page 491, and Table 5-23 on page 493.

1.1.6 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 website for the latest information:

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

The information in Table 1-9 on page 11 through Table 1-14 on page 16 is for general purposes. See the IBM ServerProven website for any additional information and limitations.

The following operating systems are described next:

- ▶ “Microsoft” on page 11
- ▶ “Novell SUSE Linux” on page 13
- ▶ “Red Hat Linux” on page 14
- ▶ “Oracle Solaris” on page 15
- ▶ “VMware” on page 15
- ▶ “IBM” on page 16

Microsoft

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

Table 1-9 Supported Microsoft operating systems

Operating system	HS12 (8028)	HS22 (7870)	HS22V (7871)	HS23 (7875, E5-2600)	HS23 (7875, E5-2600 v2)	HS23E (8038)	HX5 (7872)	HX5 (7873)	PS700/1/2 (8406)	PS703/704 (7891)
Microsoft Windows Server 2012 R2										
All Editions	N	N	N	N	Y	N	N	N	N	N
Microsoft Windows Server 2012										
All Editions	N	Y	N	Y	Y	Y	Y	Y	N	N
Microsoft Windows Server 2008 R2										
All Editions	Y	Y	Y	Y	Y	Y	Y	Y	N	N
Microsoft Windows Server 2008 x64										
Foundation	Y	N	N	N	N	N	N	N	N	N
Web Edition	Y	Y	Y	Y	Y	Y	Y	Y	N	N
Standard Edition	Y	Y	Y	Y	Y	N	Y	Y	N	N
Enterprise Edition	Y	Y	Y	Y	Y	Y	Y	Y	N	N
DataCenter Edition	N	Y	Y	Y	Y	Y	Y	Y	N	N
HPC Edition	N	Y	Y	N	N	Y	Y	Y	N	N
HPC Server 2008	N	Y	N	Y	Y	N	N	N	N	N
Microsoft Windows Server 2008 x86										
Web Edition	Y	Y	Y	N	N	N	N	N	N	N
Standard Edition	Y	Y	Y	N	N	N	N	N	N	N
Enterprise Edition	Y	Y	Y	N	N	N	N	N	N	N
DataCenter Edition	N	Y	Y	N	N	N	N	N	N	N
Microsoft Windows Essentials Business Server 2008										
Standard Edition	Y	Y	Y	N	N	N	N	N	N	N

Operating system	HS12 (8028)	HS22 (7870)	HS22V (7871)	HS23 (7875, E5-2600)	HS23 (7875, E5-2600 v2)	HS23E (8038)	HX5 (7872)	HX5 (7873)	PS700/1/2 (8406)	PS703/704 (7891)
Premium Edition	Y	Y	Y	N	N	N	N	N	N	N
Microsoft Windows Small Business Server 2008										
Standard Edition	Y	Y	Y	Y	N	N	Y	Y	N	N
Premium Edition	Y	Y	Y	Y	N	N	Y	Y	N	N
Microsoft Windows Server 2003/2003 R2 x64										
Standard Edition	Y	Y	N	N	N	N	N	N	N	N
Enterprise Edition	Y	Y	N	N	N	N	N	N	N	N
Enterprise Edition with MSCS	N	Y	N	N	N	N	N	N	N	N
DataCenter Edition	N	Y	N	N	N	N	N	N	N	N
Compute Cluster Edition	Y	Y	N	N	N	N	N	N	N	N
Compute Cluster Server 2003	Y	Y	N	N	N	N	N	N	N	N
Microsoft Windows Server 2003/2003 R2										
Web Edition	Y	Y	N	N	N	N	N	N	N	N
Standard Edition	Y	Y	N	N	N	N	N	N	N	N
Enterprise Edition	Y	Y	N	N	N	N	N	N	N	N
Enterprise Edition with MSCS	N	Y	N	N	N	N	N	N	N	N
DataCenter Edition	N	Y	N	N	N	N	N	N	N	N
Windows Small Business Server 2003/2003 R2										
Standard Edition	Y	Y	N	N	N	N	N	N	N	N
Premium Edition	Y	Y	N	N	N	N	N	N	N	N

Novell SUSE Linux

Table 1-10 summarizes support for SUSE Linux operating systems from Novell.

Table 1-10 Supported SUSE Linux operating systems

Operating system	HS12 (8028)	HS22 (7870)	HS22V (7871)	HS23 (7875, E5-2600)	HS23 (7875, E5-2600 v2)	HS23E (8038)	HX5 (7872)	HX5 (7873)	PS700/1/2 (8406)	PS703/704 (7891)
SUSE Linux Enterprise Server 11										
Real Time for AMD64/EM64T	N	Y	Y	N	N	N	N	N	N	N
With Xen for AMD64/EM64T	Y	Y	Y	Y	Y	Y	Y	Y	N	N
For AMD64/EM64T	Y	Y	Y	Y	Y	Y	Y	Y	N	N
For x86	Y	Y	Y	Y	Y	Y	N	N	N	N
For IBM POWER®	N	N	N	N	N	N	N	N	Y	Y
SUSE Linux Enterprise Server 10										
Real Time for AMD64/EM64T	N	Y	Y	N	N	N	N	N	N	N
With Xen for AMD64/EM64T	Y	Y	Y	N	N	N	N	N	N	N
With Xen for x86	Y	Y	Y	N	N	N	N	N	N	N
For AMD64/EM64T	Y	Y	Y	Y	N	Y	Y	Y	N	N
For x86	Y	Y	N	Y	N	N	N	N	N	N
For IBM POWER	N	N	N	N	N	N	N	N	Y	Y
SUSE Linux Enterprise Server 9										
For AMD64/EM64T	Y	N	N	N	N	N	N	N	N	N
For x86	Y	N	N	N	N	N	N	N	N	N

Red Hat Linux

Table 1-11 summarizes support for Red Hat Linux operating systems from Red Hat.

Table 1-11 Supported Red Hat Linux operating systems

Operating system	HS12 (8028)	HS22 (7870)	HS22V (7871)	HS23 (7875, E5-2600)	HS23 (7875, E5-2600 v2)	HS23E (8038)	HX5 (7872)	HX5 (7873)	PS700/1/2 (8406)	PS703/704 (7891)
Red Hat Enterprise Linux 6										
Server x64 Edition	N	Y	Y	Y	Y	Y	Y	Y	N	N
Server Edition	N	Y	N	Y	Y	Y	N	N	N	N
Red Hat Enterprise Linux 5										
Server x64 Edition	Y	Y	Y	Y	Y	Y	Y	Y	N	N
Server with Xen x64 Edition	Y	Y	Y	Y	Y	Y	Y	Y	N	N
Server Edition	Y	Y	N	Y	Y	Y	N	N	N	N
Server with Xen Edition	Y	Y	Y	N	N	N	N	N	N	N
For IBM POWER	N	N	N	N	N	N	N	N	Y	Y
Red Hat Enterprise Linux 4										
AS for AMD64/EM64T	Y	Y	Y	N	N	N	N	N	N	N
AS for x86	Y	Y	N	N	N	N	N	N	N	N
ES for AMD64/EM64T	Y	Y	Y	N	N	N	N	N	N	N
ES for x86	Y	Y	N	N	N	N	N	N	N	N
WS/HPC for AMD64/EM64T	Y	Y	Y	N	N	N	N	N	N	N
WS/HPC for x86	Y	Y	N	N	N	N	N	N	N	N
Red Hat Enterprise MRG Realtime (x64)										
MRG 1.0 Realtime (x64)	N	Y	Y	N	N	N	N	N	N	N
MRG 2.0 Realtime (x64)	N	Y	Y	N	N	N	Y	Y	N	N

Oracle Solaris

Table 1-12 summarizes support for Solaris 10 operating systems from Oracle.

Table 1-12 Supported Oracle Solaris operating systems

Operating system	HS12 (8028)	HS22 (7870)	HS22V (7871)	HS23 (7875, E5-2600)	HS23 (7875, E5-2600 v2)	HS23E (8038)	HX5 (7872)	HX5 (7873)	PS700/1/2 (8406)	PS703/704 (7891)
Solaris 10	N	Y	Y	Y	Y	N	Y	Y	N	N

VMware

Table 1-13 summarizes support for VMware operating systems from VMware.

Table 1-13 Supported VMware operating systems

Operating system	HS12 (8028)	HS22 (7870)	HS22V (7871)	HS23 (7875, E5-2600)	HS23 (7875, E5-2600 v2)	HS23E (8038)	HX5 (7872)	HX5 (7873)	PS700/1/2 (8406)	PS703/704 (7891)
VMware vSphere 5.5 (ESXi)	N	Y	N	N	Y	Y	Y	Y	N	N
VMware vSphere 5.1 (ESXi)	N	Y	Y	Y	Y	Y	Y	Y	N	N
VMware vSphere 5.0 (ESXi)	Y	Y	Y	Y	N	Y	Y	Y	N	N
VMware ESX 4.1	Y	Y	Y	Y	N	Y	Y	Y	N	N
VMware ESXi 4.1	Y	Y	Y	Y	N	Y	Y	Y	N	N
VMware ESX 4.0	Y	Y	Y	N	N	N	N	N	N	N
VMware ESXi 4.0	Y	Y	Y	N	N	N	N	N	N	N
VMware ESX 3.5	Y	Y	Y	N	N	N	N	N	N	N
VMware ESXi 3.5	Y	Y	Y	N	N	N	N	N	N	N

IBM

Table 1-14 summarizes support for the IBM AIX® and IBM i5/OS™ operating systems.

Table 1-14 Supported IBM AIX and i5/OS operating systems

Operating system	HS12 (8028)	HS22 (7870)	HS22V (7871)	HS23 (7875, E5-2600)	HS23 (7875, E5-2600 v2)	HS23E (8038)	HX5 (7872)	HX5 (7873)	PS700/1/2 (8406)	PS703/704 (7891)
IBM AIX 5L™ V5.3	N	N	N	N	N	N	N	N	Y	N
AIX V6.1	N	N	N	N	N	N	N	N	Y	Y
AIX V7.1	N	N	N	N	N	N	N	N	N	Y
IBM i Operating System 6.1	N	N	N	N	N	N	N	N	Y	Y
IBM i Operating System 7.1	N	N	N	N	N	N	N	N	Y	Y
IBM Virtual I/O Server	N	N	N	N	N	N	N	N	Y	Y
4690 Operating System V6	Y	N	N	Y	Y	N	N	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 16.
- ▶ 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 18.
- ▶ IBM BladeCenter HT models are designed for high-performance flexible telecommunications environments by supporting high-speed Internet working technologies, such as 10Gb Ethernet (10GbE), and provide a robust platform for NGNs. See 1.2.3, “BladeCenter HT” on page 20.
- ▶ 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 1.2.4, “BladeCenter S” on page 23.

All five chassis share a common set of blades and standard switch modules. BladeCenter H and HT offer high-speed I/O bays for high-speed switches, such as IBM 10-port 10 Gb Ethernet 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 SAS for high-speed connectivity to the client's existing network environment.

BladeCenter E also supports a redundant pair of Advanced Management Modules for comprehensive systems management.

The BladeCenter E 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. The BladeCenter E lower power consumption and IBM Calibrated Vectors™ allow more servers to fit in a tight power or cooling environment.

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



Figure 1-1 BladeCenter E front view

Figure 1-2 shows the rear view of an IBM BladeCenter E.



Figure 1-2 BladeCenter E rear view

The key features of IBM BladeCenter E chassis are as follows:

- ▶ A rack-optimized, 7 U modular design enclosure for up to 14 hot-swap blades
- ▶ High-availability mid-plane that supports hot-swap of individual blades
- ▶ Two 2,000 W or 2,320 W, hot-swap power modules (model-dependent) and support for two optional 2,000-watt or 2,320-watt power modules, offering redundancy and power for robust configurations
- ▶ Two hot-swap blowers
- ▶ An Advanced Management Module that provides chassis-level solutions, simplifying deployment and management of your installation
- ▶ Support for up to four network or storage switches or pass-through modules
- ▶ A light path diagnostic panel, and USB 2.0 port
- ▶ Support for UltraSlim enhanced SATA DVD-ROM and multi-burner drives
- ▶ Support for IBM Director and IBM Tivoli® Provisioning Manager for OS Deployments for easy installation and management
- ▶ Energy-efficient design and innovative features to maximize productivity and reduce power usage
- ▶ Extreme density and integration to ease data center space constraints
- ▶ Help in protecting your IT investment through IBM BladeCenter family longevity, compatibility, and innovation leadership in blades
- ▶ Support for the latest generation of IBM BladeCenter blades, helping provide investment protection

For more information about the IBM BladeCenter E chassis, see 2.1, “IBM BladeCenter E” on page 38.

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 an optical drive.

Figure 1-3 shows the front view of an IBM BladeCenter H.



Figure 1-3 BladeCenter H front view

Figure 1-4 shows the rear view of an IBM BladeCenter H.



Figure 1-4 BladeCenter H rear view

The key features of IBM BladeCenter H chassis are as follows:

- ▶ A rack-optimized, 9 U modular design enclosure for up to 14 hot-swap blades
- ▶ High-availability mid-plane that supports hot-swap of individual blades
- ▶ Two 2,900 W or 2,980 W, hot-swap power modules (model-dependent) and support for two optional power modules, offering redundancy and power for robust configurations
- ▶ Two hot-swap redundant blowers, and 6 or 12 supplemental fans with power supplies
- ▶ An Advanced Management Module that provides chassis-level solutions, simplifying deployment and management of your installation
- ▶ Support for up to four network or storage switches or pass-through modules
- ▶ Support for up to four bridge modules
- ▶ Support for up to four high speed switch modules
- ▶ A light path diagnostic panel, and two USB 2.0 ports
- ▶ Serial port breakout connector
- ▶ Support for UltraSlim Enhanced SATA DVD-ROM and Multi-Burner Drives
- ▶ Support for IBM Director and Tivoli Provisioning Manager for OS Deployments for easy installation and management
- ▶ Energy-efficient design and innovative features to maximize productivity and reduce power usage
- ▶ Density and integration to ease data center space constraints
- ▶ Help in protecting your IT investment through IBM BladeCenter family longevity, compatibility, and innovation leadership in blades

For more information about the IBM BladeCenter H chassis, see 2.2, “IBM BladeCenter H” on page 46.

1.2.3 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 10Gb Ethernet installations. This 12U high chassis with DC or AC power supplies provides a cost-effective, high-performance, and high-availability solution for telecommunication network and other *rugged* non-telco environments. The 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 uses the industry know-how of key IBM Business Partners to jointly deliver added value within service provider networks.

Figure 1-5 shows the front view of the BladeCenter HT.




Figure 1-5 BladeCenter HT front view

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

Figure 1-6 shows the rear view of the BladeCenter HT.



Figure 1-6 BladeCenter HT 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 the following items:

- ▶ 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 key features of the BladeCenter HT include:

- ▶ Support for up to 12 blade servers, compatible with the other chassis in the BladeCenter family
- ▶ Four standard and four high-speed I/O module bays, compatible with the other chassis in the BladeCenter family
- ▶ Two media trays at the front (each tray with light path diagnostics, two USB 2.0 ports, and optional compact flash memory module support)
- ▶ Two hot-swap management-module bays (two management modules standard)
- ▶ Four hot-swap power-module bays (four power modules standard, either 2,535 W DC or 3,160 W AC)
- ▶ Serial port for direct serial connection to installed blades
- ▶ Compliance with the NEBS 3 and ETSI core network specifications

For more information about the IBM BladeCenter HT chassis, see 2.3, “IBM BladeCenter HT” on page 56.

1.2.4 BladeCenter S

The BladeCenter S chassis can hold up to six blade servers, and up to 12 hot-swap 3.5-inch SAS or SATA disk drives, or up to 24 hot-swap 2.5-inch SAS disk drives in just 7U of rack space. It can also include up to four C14 950-watt/1450-watt power supplies. The BladeCenter S offers the necessary I/O network switching, power, cooling, and control panel information to support the individual servers.

The IBM BladeCenter S is one of five chassis in the BladeCenter family. The BladeCenter S provides an easy IT solution to the small and medium office and to the distributed enterprise.

Figure 1-7 shows the front view of IBM BladeCenter S.



Figure 1-7 The front of the BladeCenter S chassis

Figure 1-8 shows the rear view of the chassis.



Figure 1-8 The rear of the BladeCenter S chassis

The key features of IBM BladeCenter H chassis are as follows:

- ▶ A rack-optimized, 7U modular design enclosure for up to six hot-swap blades.
- ▶ Two optional Disk Storage Modules for HDDs.
- ▶ High-availability mid-plane that supports hot-swap of individual blades.
- ▶ Two 950/1450-watt, hot-swap power modules and support for two optional 950/1450-watt power modules, offering redundancy and power for robust configurations.
- ▶ Four hot-swap redundant blowers, plus one fan in each power supply.
- ▶ An Advanced Management Module that provides chassis-level solutions, simplifying deployment and management of your installation.
- ▶ Support for up to four network or storage switches or pass-through modules.
- ▶ A light path diagnostic panel, and two USB 2.0 ports.
- ▶ Support for optional UltraSlim Enhanced SATA DVD-ROM and Multi-Burner Drives.
- ▶ Support for SAS RAID Controller Module makes it easy for clients to buy the all-in-one BladeCenter S solution.
- ▶ Support for IBM Director, Storage Configuration Manager (SCM), Start Now Advisor, and Tivoli Provisioning Manager for OS Deployments support for easy installation and management.
- ▶ Energy-efficient design and innovative features to maximize productivity and reduce power usage.
- ▶ Help in protecting your IT investment through IBM BladeCenter family longevity, compatibility, and innovation leadership in blades.
- ▶ Support for the latest generation of IBM BladeCenter blades, helping provide investment protection.

For more information about the IBM BladeCenter S chassis, see 2.4, “IBM BladeCenter S” on page 67.

1.3 Blade servers

IBM BladeCenter servers support a wide selection of processor technologies and operating systems to allow clients to run all of their diverse workloads inside a single architecture. The slim and 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 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 are 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.

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

The following blade servers are covered in this section:

- ▶ 1.3.1, "BladeCenter HS12"
- ▶ 1.3.2, "BladeCenter HS22" on page 26
- ▶ 1.3.3, "BladeCenter HS23 (E5-2600)" on page 27
- ▶ 1.3.4, "BladeCenter HS23 (E5-2600 v2)" on page 28
- ▶ 1.3.5, "BladeCenter HS23E" on page 29
- ▶ 1.3.6, "BladeCenter HX5" on page 30
- ▶ 1.3.7, "IBM BladeCenter PS700, PS701, and PS702" on page 31
- ▶ 1.3.8, "BladeCenter PS703 and PS704" on page 34

1.3.1 BladeCenter HS12

BladeCenter HS12 is a single-socket blade server that supports the economics of application server deployment with power, scalability, control, and serviceability. It is well-suited for web caching, terminal serving, firewalls, dynamic web serving, and virtualization.

Figure 1-9 shows the IBM BladeCenter HS12.



Figure 1-9 IBM BladeCenter HS12

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

- ▶ 3.0 GHz dual-core processor with 6 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 two hot-swap SAS HDDs, and RAID 0 and 1 support with RAID card
- ▶ Support for Ethernet, SAS, or Fibre Channel expansion cards

The supported operating systems for the HS12 blade servers are Microsoft Windows 2003 and 2003 R2 x86 and x64, Microsoft Windows Server 2008 x86 and x64 and Windows Server 2008 R2, Red Hat Enterprise Linux 4 and 5 x86 and x64, SUSE Linux Enterprise Server 10 and 11 for x86 and for AMD64/EM64T, and VMware ESX 3.5, 4.x and vSphere 5.0. See the ServerProven website for further supported network operating system (NOS) information:

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

For more details about the server, see 3.2, “IBM BladeCenter HS12” on page 252.

1.3.2 BladeCenter HS22

IBM BladeCenter HS22 is a two-socket blade server running the latest Intel Xeon processors. It is ideal for infrastructure, virtualization, and enterprise business applications.

The HS22 supports up to two Intel Xeon Processor 5600 series multicore processors, 12 DIMMs modules, two hot-swap drives, two PCI Express connectors, and one internal USB connector.

Figure 1-10 shows the IBM BladeCenter HS22 server.



Figure 1-10 IBM BladeCenter HS22

The HS22 server has the following features:

- ▶ Supports up to two Intel Xeon 5600 series multi-core processors
- ▶ Twelve DIMM slots, for up to 192 GB of memory (using 16 GB DIMMs)
- ▶ Two hot-swap disk drive bays
- ▶ Internal USB port for embedded hypervisor support

- ▶ Onboard dual 1 Gb Ethernet ports
- ▶ Two PCI Express connectors for attaching various I/O expansion cards

The supported operating systems for the HS22 blade servers are Microsoft Windows 2003 and 2003 R2 x86 and x64, Microsoft Windows Server 2008 x86 and x64, 2008 R2, and 2012, Red Hat Enterprise Linux 4, 5, and 6 x86 and x64, SUSE Linux Enterprise Server 10 and 11 for x86 and for AMD64/EM64T, VMware ESX 3.5, 4.x and vSphere 5.0, and Oracle Solaris 10. See the following ServerProven page for further supported NOS information:

<http://www.ibm.com/systems/info/x86servers/serverproven/compat/us/nos/ematrix.shtml>

For more details about the server, see 3.3, “IBM BladeCenter HS22” on page 260.

1.3.3 BladeCenter HS23 (E5-2600)

IBM BladeCenter HS23 (E5-2600) is a two-socket blade server running the Intel Xeon processor E5-2600 product family. With its industry-leading RAS features, energy efficiency, outstanding performance, flexible and scalable I/O, and complete systems management, the HS23 offers a robust platform optimized for your mission-critical applications. The standard 30-mm single-wide form-factor protects your investments by providing compatibility with the IBM BladeCenter H, E, S, and HT chassis.

Figure 1-11 shows the HS23.



Figure 1-11 The HS23 type 7875 (E5-2600)

The HS23 (E5-2600) has the following features:

- ▶ Supports up to two Intel Xeon processor E5-2600 product family CPUs
- ▶ Sixteen DIMM slots for up to 512 GB of memory (using 32 GB DIMMs)
- ▶ Up to 1600 MHz memory speeds
- ▶ Supports up to two hot-swap HDDs
- ▶ Internal USB port for embedded hypervisor support
- ▶ Integrated dual-port Gigabit Ethernet and dual-port 10 Gb Ethernet Virtual Fabric
- ▶ Up to 18 I/O ports per blade with virtual NICs
- ▶ PCI Express connectors for attaching various I/O expansion cards
- ▶ Interposer Card connector

The supported operating systems for the HS23 (E5-2600) blade servers are Microsoft Windows Server 2012, 2008 x64, and 2008 R2, Red Hat Enterprise Linux 5 and 6, SUSE Linux Enterprise Server 10 and 11, and VMware ESX and ESXi 4.1 and vSphere 5.

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, “BladeCenter HS23 (E5-2600)” on page 283.

1.3.4 BladeCenter HS23 (E5-2600 v2)

IBM BladeCenter HS23 (E5-2600 v2) is a next-generation two-socket blade server running the Intel Xeon processor E5-2600 v2 product family. With its industry-leading RAS features, energy efficiency, outstanding performance, flexible and scalable I/O, and complete systems management, the HS23 offers a robust platform optimized for your mission-critical applications. The standard 30-mm single-wide form-factor protects your investments by providing compatibility with the IBM BladeCenter H, E, S, and HT chassis.

Figure 1-11 shows the HS23.




Figure 1-12 The HS23 type 7875 (E5-2600 v2)

The HS23 (E5-2600 v2) has the following features:

- ▶ Supports up to two Intel Xeon processor E5-2600 product family CPUs
- ▶ Sixteen DIMM slots for up to 512 GB of memory (using 32 GB DIMMs)
- ▶ Up to 1866 MHz memory speeds
- ▶ Supports up to two hot-swap HDDs
- ▶ Internal USB port for embedded hypervisor support
- ▶ Integrated dual-port Gigabit Ethernet and dual-port 10 Gb Ethernet Virtual Fabric
- ▶ Up to 18 I/O ports per blade with virtual NICs
- ▶ PCI Express connectors for attaching various I/O expansion cards
- ▶ Interposer Card connector

The supported operating systems for the HS23 (E5-2600 v2) blade servers are Microsoft Windows Server 2012 R2, 2012, 2008 x64, and 2008 R2, Red Hat Enterprise Linux 5 and 6, SUSE Linux Enterprise Server 11, and VMware vSphere (ESXi) 5.1 and 5.5.

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 3.5, “BladeCenter HS23 (E5-2600 v2)” on page 304.

1.3.5 BladeCenter HS23E

IBM BladeCenter HS23E is a versatile, dual-socket blade server running the Intel Xeon processor E5-2400 family. The server offers performance for value with new levels of memory capacity, processor performance, and flexible configuration options. A standard 30 mm single-wide form factor protects your investments by providing compatibility with the IBM BladeCenter H, E, S, and HT chassis.

Figure 1-13 shows the HS23E.




Figure 1-13 The HS23E Type 8038

The HS23E has the following features:

- ▶ Support for up to two Intel Xeon E5-2400 family processors, or one Intel Xeon E5-1410, or one Intel Pentium 1400 family processor (Intel Pentium 1403 or Intel Pentium 1407)
- ▶ Twelve DIMM slots for up to 192 GB of memory (using 32 GB DIMMs)
- ▶ Support for up to two hot-swap HDDs
- ▶ Internal USB port for embedded hypervisor support
- ▶ Integrated dual-port Gigabit Ethernet
- ▶ PCI Express connectors for attaching various I/O expansion cards

The supported operating systems for the HS23E blade servers are Microsoft Windows Server 2012, 2008 x64, and 2008 R2, Red Hat Enterprise Linux 5 x64 and 6, SUSE Linux Enterprise Server 10 for AMD64/EM64T and 11, and VMware ESX and ESXi 4.1 and vSphere 5.

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 3.6, “BladeCenter HS23E” on page 324.

1.3.6 BladeCenter HX5

IBM BladeCenter HX5 type 7873 is a scalable two-socket blade server running the latest Intel Xeon processor E7-8800, E7-4800, or E7-2800 families. The IBM BladeCenter HX5 type 7872 is a scalable two-socket blade server running the Intel Xeon 7500 or 6500 series processors. The HX5 is ideal for large virtualization and consolidation solutions and for enterprise business applications that require high performance. The HX5 is supported in the BladeCenter H, HT (AC version only), and S chassis.

The HX5 has 16 DIMM module slots, three PCI Express connectors, and one internal USB connector. Two single HX5 blade servers can be joined to form one four-socket server with up to 32 DIMM modules.

The MAX5 memory expansion blade can be attached to the HX5, which increases the memory capacity to 40 DIMM sockets for a two-socket HX5 server. The combined HX5 and MAX5 then take up two slots in the BladeCenter chassis.

Figure 1-14 shows the two HX5 blade servers joined.



Figure 1-14 HX5 (two servers shown, connected to form one single-image four-socket server)

The HX5 has the following features:

- ▶ Supports up to two Intel Xeon processors E7 family
- ▶ Supports up to two 1.8-inch SSDs
- ▶ Sixteen DIMM slots for up to 256 GB of memory (using 16 GB DIMMs)
- ▶ Integrated dual-port Gigabit Ethernet
- ▶ Two PCIe expansion connectors
- ▶ Up to 12 I/O ports per blade

The supported operating systems for the HX5 blade servers are Microsoft Windows Server 2012, 2008 x64, and 2008 R2, Red Hat Enterprise Linux 5 and 6 x64, SUSE Linux Enterprise Server 10 and 11 for AMD64/EM64T, and VMware ESX and ESXi 4.x and vSphere 5.

For more details about the server, see 3.7, “IBM BladeCenter HX5” on page 341.

1.3.7 IBM BladeCenter PS700, PS701, and PS702

The PS700, PS701, and PS702 blades are high-performance blades for 64-bit applications. They are designed to minimize complexity, improve efficiency, automate processes, reduce energy consumption, and scale easily. The IBM POWER7® processor-based blades support AIX, IBM i, and Linux operating systems. Their ability to coexist in the same chassis with other IBM BladeCenter blades servers enhances the ability to deliver rapid return of investment demanded by clients and businesses.

Figure 1-15 shows the IBM BladeCenter PS700, PS701, and PS702 blade servers.



Figure 1-15 The IBM BladeCenter PS702, BladeCenter PS701, and BladeCenter PS700

The PS700 blade server

The PS700 blade server (8406-70Y) is a single-socket, single-wide 4-core 3.0 GHz POWER7 processor-based server. The POWER7 processor is a 64-bit, 4-core with 256 KB L2 cache per core and 4 MB L3 cache per core.

The PS700 blade server has eight DDR3 memory DIMM slots. The industry standard VLP DDR3 Memory DIMMs are either 4 GB or 8 GB running at 1066 MHz. The memory is supported in pairs. Thus, the minimum memory required for PS700 blade server is 8 GB (two 4 GB DIMMs). The maximum memory that can be supported is 64 GB (eight 8 GB DIMMs).

It has two Host Ethernet Adapter (HEA) 1 GB integrated Ethernet ports that are connected to the BladeCenter chassis fabric (midplane). The PS700 has an integrated SAS controller that supports local (onboard) storage, an integrated USB controller, and Serial over LAN console access through the service processor, and the BladeCenter Advance Management Module.

It supports two onboard disk drive bays. The onboard storage can be one or two 2.5-inch SAS HDD or SSD drives. The integrated SAS controller supports RAID 0, RAID 1, and RAID10 hardware when two HDDs or solid-state drives (SSDs) are used.

The PS700 also supports one PCIe CIOv expansion card slot and one PCIe CFFh expansion card slot.

The PS701 blade server

The PS701 blade server (8406-71Y) is a single socket, single-wide 8-core 3.0 GHz POWER7 processor-based server. The POWER7 processor is a 64-bit, 8-core with 256 KB L2 cache per core and 4 MB L3 cache per core.

The PS701 blade server has 16 DDR3 memory DIMM slots. The industry standard VLP DDR3 memory DIMMs are either 4 GB or 8 GB running at 1066 MHz. The memory is supported in pairs. Thus, the minimum memory required for the PS701 blade server is 8 GB (two 4 GB DIMMs). The maximum memory that can be supported is 128 GB (16 x 8 GB DIMMs).

The PS701 blade server has two HEA 1 GB integrated Ethernet ports that are connected to the BladeCenter chassis fabric (midplane). The PS701 also has an integrated SAS controller that supports local (onboard) storage, an integrated USB controller, and Serial over LAN console access through the service processor, and the BladeCenter Advance Management Module.

The PS701 has one onboard disk drive bay. The onboard storage can be one 2.5-inch SAS HDD or SSD drive. The PS701 also supports one PCIe CIOv expansion card slot and one PCIe CFFh expansion card slot.

The PS702 blade server

The PS702 blade server (8406-71Y +FC 8358) is a two-socket, double-wide 16-core 3.0 GHz POWER7 processor-based server. The POWER7 processor is a 64-bit, 8-core with 256 KB L2 cache per core and 4 MB L3 cache per core.

The PS702 combines a single-wide base blade (PS701) and an expansion unit (feature 8358), referred to as double-wide blade, which occupies two adjacent slots in the IBM BladeCenter chassis.

The PS702 blade server has 32 DDR3 memory DIMM slots. The industry-standard VLP DDR3 memory DIMMs are either 4 GB or 8 GB running at 1066 MHz. The memory is supported in pairs. Thus, the minimum memory required for the PS702 blade server is 8 GB (two 4 GB DIMMs). The maximum memory that can be supported is 256 GB (32 x 8 GB DIMMs).

The PS702 blade server has four HEA 1 GB integrated Ethernet ports that are connected to the BladeCenter chassis fabric (midplane). The PS702 also has an integrated SAS controller that supports local (onboard) storage, an integrated USB controller, and Serial over LAN console access through the service processor, and the BladeCenter Advance Management Module.

The PS702 blade server has two disk drive bays, one on the base blade and one on the expansion unit. The onboard storage can be one or two 2.5-inch SAS HDD or SSD drives. The integrated SAS controller supports RAID 0, RAID 1, and RAID 10 hardware when two HDDs or SSDs are used.

The PS702 supports two PCIe CIOv expansion card slots and two PCIe CFFh expansion card slots.

Figure 1-16 shows the internals of the BladeCenter PS702 blade server.



Figure 1-16 The internals of the BladeCenter PS702 blade server (showing the top half)

The PS700, PS701, and PS702 include the following features and benefits:

- ▶ The efficient and flexible design of IBM blades:
 - Densely pack more servers into a smaller space.
 - Tailor the system to meet varied business requirements with a choice of the BladeCenter chassis.
 - Lower acquisition cost and energy consumption versus traditional 1U or 2U rack servers.
 - Integrate a networking switch infrastructure for improved cabling and data center maintenance.
 - Deploy in virtually any office environment for quiet, highly secure, and contaminant-protected operation.
- ▶ Pioneering IBM EnergyScale™ technology and Systems Director Active Energy Manager software:
 - Generate less heat by managing application utilization and server energy consumption.
 - Use less energy to cool the system.
- ▶ Industry-leading IBM PowerVM® virtualization technology:
 - Reduce infrastructure costs by doing more with fewer servers.
 - Simplify IT operations to use storage, network, and computing resources to control costs and be more responsive.

- ▶ Innovative reliability features and systems management:
 - Expedite hardware repairs and reduce service time.
 - Enable scheduled maintenance with proactive monitoring of critical system components to help reduce unplanned failures.
- ▶ Support for IBM AIX, IBM i, or Linux operating systems:
Standardize on a single platform that runs the large and varied portfolio of applications that support your business.

The list of supported operating systems include AIX, IBM i, SLES for IBM POWER, and RHEL for IBM POWER.

See the following ServerProven web page for further supported NOS information:

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

For more information about these POWER7 processor-based blade servers, see 3.8, “IBM BladeCenter PS700, PS701, and PS702” on page 377.

1.3.8 BladeCenter PS703 and PS704

The BladeCenter PS703 and BladeCenter PS704 are new servers based on POWER7 processors. These new blades offer processor scalability from 16 cores to 32 cores:

- ▶ IBM BladeCenter PS703: Single-wide blade with two 8-core processors
- ▶ IBM BladeCenter PS704: Double-wide blade with four 8-core processors

The POWER7 processor-based PS703 and PS704 blades support the AIX, IBM i, and Linux operating systems. Their ability to coexist in the same chassis with other IBM BladeCenter blades servers enhances their ability to deliver rapid return of investment demanded by clients and businesses.

Figure 1-17 shows the IBM BladeCenter PS703 and PS704 blade servers.



Figure 1-17 The IBM BladeCenter PS703 (right) and BladeCenter PS704 (left)

The PS703 blade server

The IBM BladeCenter PS703 (7891-73X) is a single-wide blade server with two 8-core POWER7 processors with a total of 16 cores. The processors are 64-bit, 8-core 2.4 GHz processors with 256 KB L2 cache per core and 4 MB L3 cache per core.

The PS703 blade server has 16 DDR3 memory DIMM slots. The industry standard VLP DDR3 memory DIMMs are either 4 GB or 8 GB running at 1066 MHz. The minimum memory required for PS703 blade server is 16 GB. The maximum memory that can be supported is 128 GB (16 x 8 GB DIMMs).

The PS703 blade server supports optional IBM Active Memory™ Expansion, which is a POWER7 technology that allows the effective maximum memory capacity to be much larger than the true physical memory. Innovative compression and decompression of memory content using processor cycles can allow memory expansion up to 100%. This expansion can allow an AIX V6.1, or later, partition to do more work with the same physical amount of memory or a server to run more partitions and do more work with the same physical amount of memory.

The PS703 blade server has two onboard 1 Gb integrated Ethernet ports that are connected to the BladeCenter chassis fabric (midplane). The PS703 also has an integrated SAS controller that supports local (onboard) storage, integrated USB controller and Serial over LAN console access through the service processor, and the BladeCenter Advance Management Module.

The PS703 has one onboard disk drive bay. The onboard storage can be one 2.5-inch SAS HDD or two 1.8-inch SATA SSD drives (with the addition of an SSD interposer tray). The PS703 also supports one PCIe CIOv expansion card slot and one PCIe CFFh expansion card slot.

The PS704 blade server

The IBM BladeCenter PS704 (7891-74X) is a double-wide blade server with four 8-core POWER7 processors with a total of 32 cores. The processors are 64-bit 8-core 2.4 GHz processors with 256 KB L2 cache per core and 4 MB L3 cache per core.

The PS704 is a double-wide blade, which means that it occupies two adjacent slots in the IBM BladeCenter chassis.

The PS704 blade server has 32 DDR3 memory DIMM slots. The industry standard VLP DDR3 memory DIMMs are either 4 GB or 8 GB running at 1066 MHz. The minimum memory required for PS704 blade server is 32 GB. The maximum memory that can be supported is 256 GB (32x 8 GB DIMMs).

The PS704 blade server supports optional Active Memory Expansion, which is a POWER7 technology that allows the effective maximum memory capacity to be much larger than the true physical memory. Innovative compression and decompression of memory content using processor cycles can allow memory expansion up to 100%. This expansion can allow an AIX V6.1, or later, partition to do more work with the same physical amount of memory or a server to run more partitions and do more work with the same physical amount of memory.

The PS704 blade server has four onboard 1 Gb integrated Ethernet ports that are connected to the BladeCenter chassis fabric (midplane). The PS704 also has an integrated SAS controller that supports local (onboard) storage, integrated USB controller and Serial over LAN console access through the service processor, and the BladeCenter Advance Management Module.

The PS704 blade server has two disk drive bays, one on the base blade and one on the expansion unit. The onboard storage can be one or two 2.5-inch SAS HDD or up to four 1.8-inch SSD drives. The integrated SAS controller supports RAID 0, 10, 5, or 6, depending on the numbers of HDDs or SSDs installed.

The PS704 supports two PCIe CIOv expansion card slots and two PCIe CFFh expansion card slots.

For more information about these POWER7 processor-based blade servers, see 3.9, “IBM BladeCenter PS703 and PS704” on page 389.



IBM BladeCenter chassis and infrastructure configuration

This chapter contains information about infrastructure options available for IBM BladeCenter, such as management, power, cooling, network, and storage connectivity modules.

This chapter includes the following topics:

- ▶ 2.1, “IBM BladeCenter E” on page 38
- ▶ 2.2, “IBM BladeCenter H” on page 46
- ▶ 2.3, “IBM BladeCenter HT” on page 56
- ▶ 2.4, “IBM BladeCenter S” on page 67
- ▶ 2.5, “Chassis I/O bay and expansion card port mappings” on page 80
- ▶ 2.6, “I/O modules overview” on page 83
- ▶ 2.7, “Ethernet Switch Modules” on page 84
- ▶ 2.8, “InfiniBand switch module” on page 150
- ▶ 2.9, “Fibre Channel switch modules” on page 153
- ▶ 2.10, “SAS I/O modules” on page 168
- ▶ 2.11, “Pass-through and interconnect modules” on page 173
- ▶ 2.12, “Advanced Management Module” on page 185
- ▶ 2.13, “Direct serial connections” on page 190
- ▶ 2.14, “Installation and physical site plans” on page 193

2.1 IBM BladeCenter E

This section describes the following topics:

- ▶ IBM BladeCenter E product information
- ▶ BladeCenter E I/O topology
- ▶ Blower module for BladeCenter E
- ▶ Acoustic Attenuation Module
- ▶ Power modules for BladeCenter E
- ▶ Product publications

2.1.1 IBM BladeCenter E product information

Figure 2-1 shows the IBM BladeCenter E front view with its key chassis features.




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

The following components are at the front of the BladeCenter chassis:

- ▶ A media tray with an optical drive, one USB V2.0 port, and a system status LED panel
- ▶ Fourteen hot-swap blade server bays supporting different blade server types
- ▶ A system service card next to blade bay 14

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

Figure 2-2 shows the rear view of the IBM BladeCenter E with its key chassis features.




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

The back of the BladeCenter E chassis has the following components:

- ▶ Two hot-swap blower modules
- ▶ Four hot-swap power supply bays
- ▶ Two hot-swap management module bays
- ▶ Four hot-swap I/O bays for standard form-factor I/O modules
- ▶ Rear system status LED panel

Supported I/O modules: For a list of I/O modules supported in the BladeCenter E, see Table 2-88 on page 185.

Table 2-1 lists the major features of the IBM BladeCenter E.

Table 2-1 BladeCenter E major features

Feature	Machine type/model 8677-3Tx ^{a,b}	Machine type/model 8677-4Tx ^a
Rack form factor (H x D)	7U x 28 in. (711 mm)	7U x 28 in. (711 mm)
Optical drives in Media Tray standard	None ^c	None ^c
Diskette drives in Media Tray standard	Not available ^d	Not available ^d
Number of blade server slots (30 mm)	14	14
Number of management modules (std/max ^e)	1/2	1/2
Number of standard I/O module slots	4	4
Switch modules standard	None	None
Number of power supplies (std/max)	2/4	2/4
Power supply size standard	2000 W AC	2320 W AC
Number of blowers (std/max)	2/2	2/2
Dimensions	Height: 12.0 in. (305 mm) Width: 17.5 in. (4429 mm) Depth: 28.0 in. (711 mm)	Height: 12.0 in. (305 mm) Width: 17.5 in. (4429 mm) Depth: 28.0 in. (711 mm)

a. The x is the country-specific letter (for example, EMEA MTM is 86774TG, and the US MTM is 86774TU).

b. This model is not available in EMEA.

c. The optical drive is optional, as shown in Table 2-3.

d. There is no diskette drive bay in media tray.

e. Standard/maximum (std/max)

Optical drive: The optical drive is optional on all currently available models of BladeCenter E (models 8677-3Tx and 8677-4Tx).

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 to other configuration rules that are described in later chapters of this document.

The BladeCenter E chassis includes, as standard, one Advanced Management Module (AMM). This module manages the chassis, and provides the local keyboard, video, mouse (KVM) function.

An optional redundant AMM 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.

Advanced Management Module: The BladeCenter E chassis started included, as standard, the AMM in 1Q06. Before 1Q06, it included the original Management Module.

Further details regarding the features and functions of the management modules can be found in 2.12, “Advanced Management Module” on page 185.

The BladeCenter E does not include, as standard, 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 depend on the I/O Expansion Card installed in the blade servers.

Table 2-2 lists the part numbers to use when ordering the IBM BladeCenter E chassis.

Table 2-2 IBM BladeCenter E part numbers for ordering

Description	Part number	Machine type/model
IBM BladeCenter E	86773Tx ^{a,b}	8677-3Tx ^{a,b}
IBM BladeCenter E	86774Tx ^a	8677-4Tx ^a

a. The x is the country-specific letter (for example, the EMEA part number is 86774TG, and the US part number is 86774TU).

b. This model is not available in EMEA.

Table 2-3 lists optional optical drives that can be ordered for the BladeCenter E chassis and that are not standard with the optical drive.

Table 2-3 Optical drives for BladeCenter E

Description	Part number	Feature code
IBM UltraSlim Enhanced SATA DVD-ROM	46M0901	4161
IBM UltraSlim Enhanced SATA Multi-Burner	46M0902	4163

The BladeCenter E chassis includes, as standard, the following items:

- ▶ One AMM
- ▶ Two hot-swap power supply modules (2000 W or 2320 W - model-dependent)
- ▶ Two hot-swap blower modules
- ▶ Two intra-rack IEC 320 C19 to C20 power cables
- ▶ Rack mount kit
- ▶ Publications/CD package

Guidelines for how to install I/O module options are in 2.1.2, “BladeCenter E I/O topology” on page 41.

2.1.2 BladeCenter E I/O topology

The IBM BladeCenter E chassis provides connection paths between server blade bays and I/O or switch bays (Figure 1-1 on page 17 and Figure 1-2 on page 17) through a hardwired dual redundant midplane.

The I/O topology is shown in Figure 2-3 on page 42. 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).

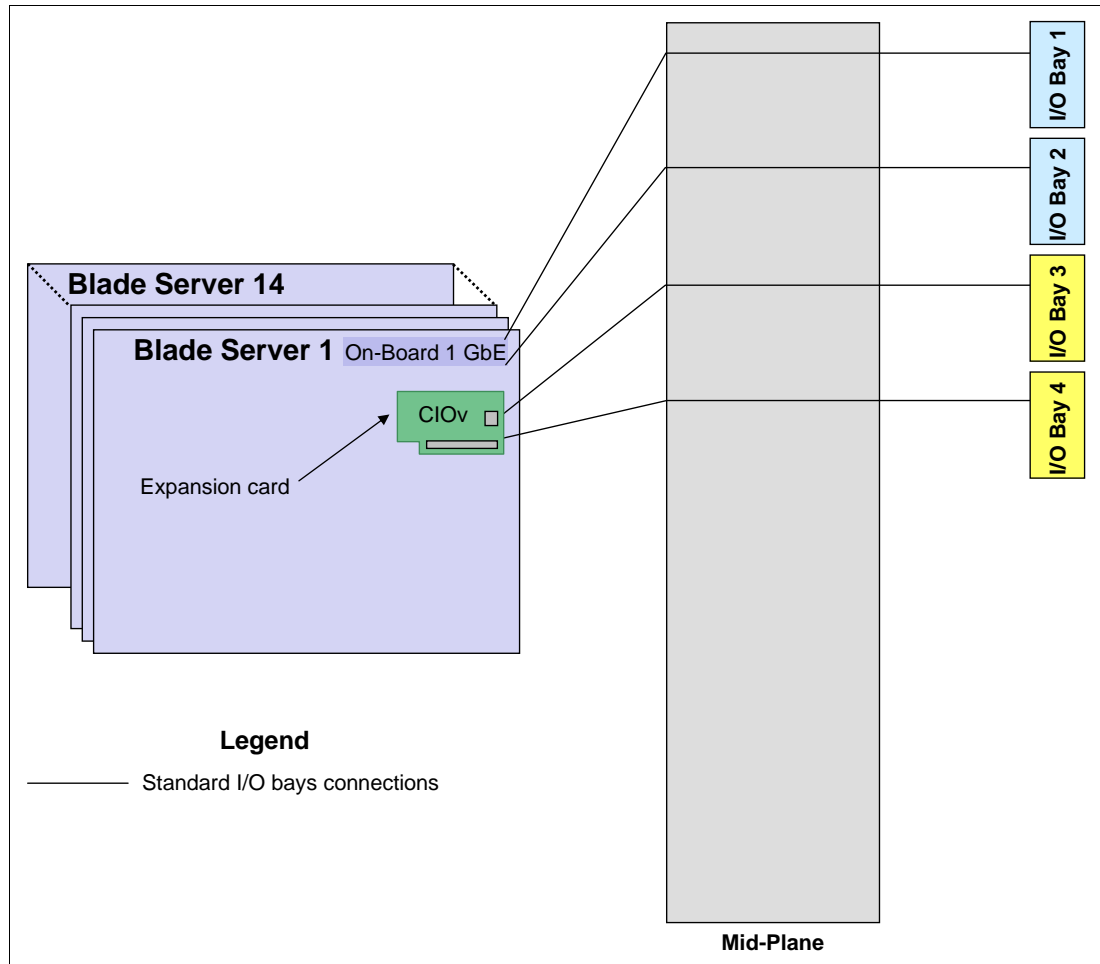


Figure 2-3 IBM BladeCenter E I/O topology

Each blade server has at least two Ethernet ports (Ethernet 1 and Ethernet 2). Ethernet 1 is routed to I/O bay 1 and Ethernet 2 is routed to I/O bay 2, 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 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 (dual-port Gigabit Ethernet or dual-port Fibre Channel) is to avoid a single point of failure 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 the CFFv or CIOv slot of a blade server, the I/O modules 3 and 4 must also be based on Fibre Channel. The reverse is also true. If you install Fibre Channel switches in bays 3 and 4, then any expansion card installed in the CFFv or CIOv in all blade servers in the chassis must be Fibre Channel cards.

See Table 2-31 on page 80 to match the expansion card ports with the corresponding I/O bay numbers. See Table 2-32 on page 82 to match the I/O bay number with the corresponding I/O module that can be installed into this bay.

2.1.3 Blower module for BladeCenter E

The BladeCenter E chassis includes two hot-swap blowers for 1+1 cooling redundancy.

Figure 2-4 shows an IBM BladeCenter E blower. Figure 2-2 on page 39 shows the locations of the blowers in the chassis.



Figure 2-4 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 run at 30% of their maximum rotational speed, increasing their speed as required to control the internal BladeCenter temperature.
- ▶ If the ambient temperature is above 72°F, the blowers runs at 80% of their maximum rotational speed, increasing their speed as required to control the 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.

For information about blowers, see *Installation and User's Guide - IBM BladeCenter E*:

<http://www.redbooks.ibm.com/abstracts/tips0756.html>

In noise-sensitive environments, you can use the Acoustic Attenuation Module (for more information, see 2.1.4, “Acoustic Attenuation Module” on page 44). The other way of limiting noise level is to use the “acoustic mode” setting in the AMM. With this mode, the AMM throttles the processor speeds of the blades to stay within noise limits.

For more information, see *IBM BladeCenter Advanced Management Module User's Guide*:

<http://www.redbooks.ibm.com/abstracts/tips0756.html>

2.1.4 Acoustic Attenuation Module

The Acoustic Attenuation Module, part number 39M4674 (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 (Figure 2-5).




Figure 2-5 IBM BladeCenter E Acoustic Attenuation Module

For more information, see *Acoustic Attenuation Module User's Guide*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-45155>

Muffler support: There is no Acoustic Attenuation Module for the BladeCenter H, HT, T, or S chassis.

2.1.5 Power modules for BladeCenter E

There are two power modules that are standard with the IBM BladeCenter E, with support for a maximum of four.

Figure 2-6 shows a power module.



Figure 2-6 Power module for IBM BladeCenter E

The IBM BladeCenter E unit includes one pair of 2000 W or 2320 W (model-dependent) 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 - 6.

The BladeCenter E unit supports a second pair of power modules in power module bays 3 and 4 that provide power to blade bays 7 - 14 (Figure 2-2 on page 39).

Table 2-4 summarizes the application for each power module.

Table 2-4 Power module functions

Bays	Power module function
1 and 2	Provides power to all modules and to blade bays 1 - 6.
3 and 4	Provides power to blade bays 7 - 14.

Power modules are not needed in bays 3 and 4 until you begin installing blade servers and have options for blade bays 7 - 14.

Table 2-5 lists the power supply options available for the IBM BladeCenter E chassis.

Table 2-5 Power supply options for IBM BladeCenter E

Part number	Feature code	Description
39M4675	2105	IBM BladeCenter E 2000 W Power Supply Modules ^a
46M0508	2106	IBM BladeCenter E 2320 W Power Supply Option ^{a,b}

a. This option contains two power supplies and two IEC 320-C19 to C20 power cables.

b. This option requires the AMM.

Storage expansion unit option: If you install a blade server that has a storage expansion unit option attached in blade bay 6, the option uses blade bay 7. Power modules are 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, 1800, or 2000 W), it is important to ensure that the pair of power modules in the same power domain has the same capacity.

When using 2320 W power modules in a BladeCenter E Type 8677 unit, all power modules must be 2320 W modules.

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 power distribution unit (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.

For more information about power requirements for BladeCenter E, see 2.14.3, "Power considerations" on page 197.

2.1.6 Product publications

The following product publications are available for the BladeCenter E:

- ▶ *Installation and User's Guide - IBM BladeCenter (Type 8677):*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-45152>
- ▶ *Hardware Maintenance Manual and Troubleshooting Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-50053>
- ▶ *Acoustic Attenuation Module User's Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-45155>

2.2 IBM BladeCenter H

This section covers the following topics:

- ▶ IBM BladeCenter H product information
- ▶ BladeCenter H I/O topology
- ▶ Blower modules for BladeCenter H
- ▶ Power modules for IBM BladeCenter H
- ▶ Serial Port Breakout Cable
- ▶ Product publications

2.2.1 IBM BladeCenter H product information

Figure 2-7 shows the front view of the IBM BladeCenter H with its key chassis features (model 8852-5Tx).




Figure 2-7 BladeCenter H front view showing the key features of the chassis

The key features on the front of the BladeCenter H (model 8852-5Tx) are as follows:

- ▶ A media tray at the front right, with an optical drive, two USB 2.0 ports, and a system status LED panel.
- ▶ One pair of 2980 W power modules. An additional power module option (containing two 2980 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.

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

Figure 2-8 shows the rear view of the IBM BladeServer H with its key chassis features (model 8852-5Tx).




Figure 2-8 BladeCenter H rear view showing the key features of the chassis

The following components are at the rear of the BladeCenter H chassis (8852-5Tx):

- ▶ Two hot-swap blower modules
- ▶ Two hot-swap management-module bays
- ▶ Six I/O bays for standard form factor I/O modules (two of them are dedicated to bridge modules only)
- ▶ Four I/O bays for high speed form factor I/O modules
- ▶ Serial port breakout connector to give direct serial connection to installed blades (for those blades with the functionality)
- ▶ Rear system status LED panel
- ▶ Two power connectors

Supported I/O modules: For a list of I/O modules supported in the BladeCenter H, see Table 1-2 on page 3.

Table 2-6 lists the major features of the IBM BladeCenter H.

Table 2-6 BladeCenter H features at a glance

Feature	Specifications
Machine type	8852-5Tx ^a
Rack form factor (H x D)	9U x 28 in. (711 mm)
Optical drives in media tray standard	None ^b
Diskette drives in media tray standard	Not applicable ^c
Number of 30-mm blade server slots	14
Number of management modules (std/max)	1 / 2
Number of standard I/O module slots	6
Number of high-speed I/O module slots	4
Switch modules standard	None
Number of power supplies (std/max)	2 / 4
Power supply size standard	2980 W AC
Number of blowers (std/max)	2 / 2
Dimensions	Height: 15.75 in. (400 mm) Width: 17.4 in. (442 mm) Depth: 28.0 in. (711 mm)

a. The x is the country-specific letter (for example, EMEA MTM is 88525TG, and the US MTM is 88525TU).

b. The optical drive is optional.

c. There is no diskette drive bay in the media tray.

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 model 8852-5Tx includes, as standard, one AMM. This module manages the chassis, and also provides the local KVM function.

The optional redundant AMM 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 model 8852-5Tx does not include, as standard, any I/O modules. You choose these I/O modules based on your connectivity needs. An ESM is required in I/O module bays 1 and 2 to enable the usage of both Ethernet ports on a blade server. The I/O modules that are required in I/O module bays 3 and 4 depend on the I/O card installed in the blade servers. The I/O modules that are required in the high speed I/O module bays 7, 8, 9, and 10 depends on the I/O cards installed in the blade servers.

Adapter support: The adapters installed in the blades servers must be supported by corresponding I/O Modules.

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

Power cords: The BladeCenter H chassis does not include a power cord. Power cords must be ordered separately. See Table 2-12 on page 55.

Table 2-7 lists the part number to use to order the IBM BladeCenter H chassis.

Table 2-7 IBM BladeCenter H part numbers for ordering

Description	Part number	Machine type
IBM BladeCenter H	88525Tx ^a	8852-5Tx ^a

a. The x is the country-specific letter (for example, the EMEA part number is 88525TG, and the US part number is 88525TU).

Table 2-8 lists optional optical drives that can be ordered for the BladeCenter H chassis.

Table 2-8 Optical drives for BladeCenter H

Description	Part number	Feature code
IBM UltraSlim Enhanced SATA DVD-ROM	46M0901	4161
IBM UltraSlim Enhanced SATA Multi-Burner	46M0902	4163

The BladeCenter H chassis includes, as standard, the following items:

- ▶ One AMM
- ▶ Two hot-swap power supply modules 2980 W AC
- ▶ Two hot-swap blower modules
- ▶ Rack mount kit
- ▶ Publications/CD package

Guidelines about how to install I/O module options are described in 2.2.2, “BladeCenter H I/O topology” on page 50.

2.2.2 BladeCenter H I/O topology

The IBM BladeCenter H chassis has two types of fabrics inside:

- ▶ Standard fabric (similar to the BladeCenter E chassis)
- ▶ High-speed fabric, which can carry 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 4.

The IBM BladeCenter H chassis has a total of 10 I/O bays (Figure 2-8 on page 48). Each blade bay has a total of eight dedicated connection paths to the I/O modules (Figure 2-9).

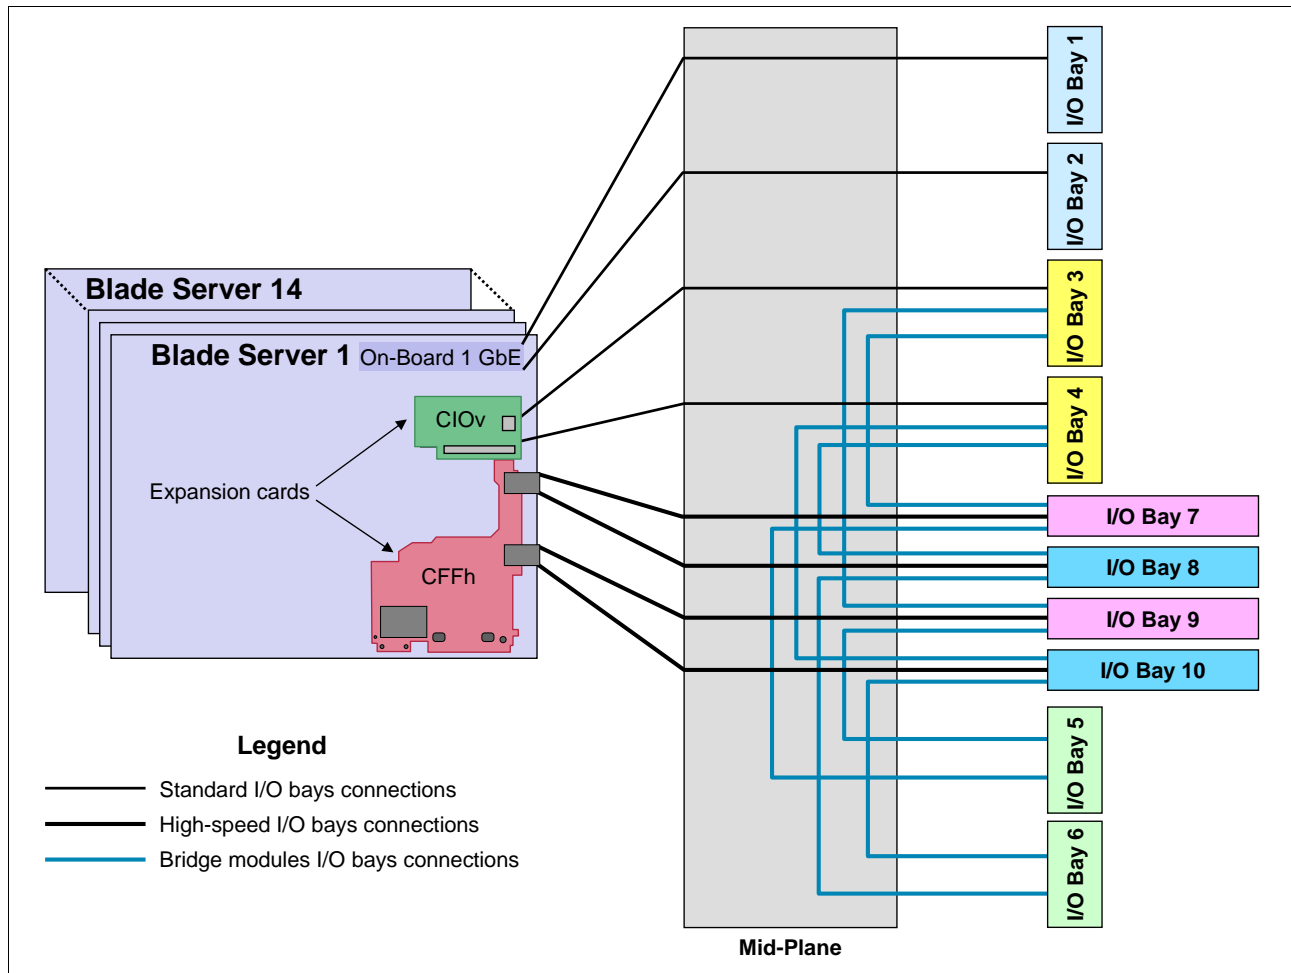


Figure 2-9 IBM BladeCenter H I/O topology

The bays are as follows:

- Bays 1 and 2 only support standard Ethernet-compatible I/O modules. These bays are routed internally to the onboard Ethernet controllers on the blades (and slot 1 in the SIO expansion blade).
- Bays 3 and 4 can be used either for standard switch or pass-through modules (such as Fibre Channel connectivity or additional Ethernet ports) or for bridge modules. These bays 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 an IBM Virtual Fabric 10 Gb Switch Module or Cisco Nexus 4001I Switch Module. These bays are routed internally through midplane connectors on certain expansion cards to the PCI Express connector on blades that have it.

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, an FC switch module), 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.

See Table 2-31 on page 80 to match the expansion card ports with the corresponding I/O bay numbers. See Table 2-32 on page 82 to match the I/O bay number with the corresponding I/O module that can be installed into this bay.

2.2.3 Blower modules for BladeCenter H

The BladeCenter H chassis includes two hot-swap blowers for 1+1 cooling redundancy (Figure 2-10).



Figure 2-10 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 run at their minimum rotational speed, increasing their speed as required to control the internal BladeCenter temperature.
- ▶ If the ambient temperature is above 25°C (77°F), the blowers run faster, increasing their speed as required to control the internal BladeCenter H unit temperature.

If a blower fails, the remaining blower runs at 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.

Two types of blower modules are available:

- ▶ Standard blowers: Standard in model 8852-4Sx and older models.
- ▶ Enhanced Cooling Module (also referred to as the *enhanced blower*): Standard in model 5Tx and some older models 4Tx, 91x, 92x, 93x, 94x, 95x and 96x. Optional in all other BladeCenter H models.

Upgrades: The Enhanced Cooling Module must support the 130 W Intel processors available for the HS22, HS22V, and HX5 blade servers. The IBM BladeCenter H chassis model 8852-5Tx and older models 8852-94x, 8852-95x, 8852-96x, 8852-4Tx, 8852-91x, 8852-92x, 8852-93x have the Enhanced Cooling Modules standard. Other BladeCenter H models support the Enhanced Cooling Module as an upgrade.

To determine whether you need Enhanced Cooling modules, consult the “Chassis support” subsection of the relevant blade server section in Chapter 3, “Blade server hardware configuration” on page 211.

Table 2-9 lists the optional enhanced blowers that can be ordered for the IBM BladeCenter H chassis.

Table 2-9 Enhanced blowers for IBM BladeCenter H

Description	Part number	Feature code
IBM BladeCenter H Enhanced Cooling Modules (two blowers) ^a	68Y6650	0724

a. The IBM BladeCenter H Enhanced Cooling Modules option contains two blowers. Only one option should be ordered to upgrade a single chassis.

Blower support: The IBM BladeCenter H chassis supports only the pair of blowers of the same type, either standard or enhanced. You cannot mix standard and enhanced blowers in the same chassis.

For more information, see *IBM BladeCenter H - Installation and User's Guide*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-63306>

2.2.4 Power modules for IBM BladeCenter H

The IBM BladeCenter H unit includes two 2,980 W hot-swap power modules. Figure 2-11 shows the power module for IBM BladeCenter H.



Figure 2-11 IBM BladeCenter H Power Supply

The standard power modules in bays 1 and 2 are used to power blade servers in blade bays 1 - 7 and I/O modules in I/O module bays 1 - 4 and 7 - 10.

Optional power modules are needed in power module bays 3 and 4 if you install blade servers in blade bays 8 - 14, or if you install I/O modules in any of I/O module bays 5 - 10 (Table 2-10).

Table 2-10 Devices powered by each power module

Devices ^a	Power module 1	Power module 2	Power module 3	Power module 4
Blades 1 - 7	Yes	Yes		
Blades 8 - 14			Yes	Yes
I/O modules 1 and 2 (SM 1 and 2) ^b	Yes	Yes	Yes	Yes
I/O modules 3 and 4 (SM 3 and 4/BM 3 and 4)	Yes	Yes		
I/O modules 5 and 6 (BM 1 and 2)			Yes	Yes
I/O modules 7 and 8 (HSSM 1 and 3)	Yes		Yes	
I/O modules 9 and 10 (HSSM 2 and 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.

Tip: Management module bays 1 and 2 and I/O modules 1 and 2 are powered from any or all power. Any of four power supplies can supply power to them.

Two types of power supply modules are available:

- ▶ 2,900 W AC power module: Standard in older BladeCenter H models
- ▶ 2,980 W AC power module: Standard in BladeCenter H model 8852-5Tx and older models 8852-94x, 8852-95x, 8852-96x, 8852-4Tx, 8852-91x, 8852-92x, 8852-93x; optional in other BladeCenter H models

Upgrades: The 2,980 W AC power module provides additional power capacity to the chassis, which is especially important when using the 130 W Intel processors available for the HS22, HS22V, and HX5 blade servers. The IBM BladeCenter H chassis model 8852-5Tx and older models 8852-94x, 8852-95x, 8852-96x, 8852-4Tx, 8852-91x, 8852-92x, 8852-93x have the 2,980 W AC power modules standard. Other BladeCenter H models support the 2,980 W AC power module as an upgrade.

To determine whether you need 2,980 W power supplies, consult the “Chassis support” subsection of the relevant blade server section in Chapter 3, “Blade server hardware configuration” on page 211.

Table 2-11 lists the power supply option available for IBM BladeCenter H chassis. Single power supply option part number contains two power supplies and two fan packs (one fan pack per one power supply).

Table 2-11 Power supply option for IBM BladeCenter H

Part number	Feature code	Description
68Y6601	2143	IBM BladeCenter H 2980 W AC Power Modules (2) 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 three-fan pack used for power supply cooling. The power supply itself has no external power connectors. Two special power connectors are at the rear of the chassis (see Figure 2-8 on page 48).

No power cables are included with either the BladeCenter H chassis or the power module option. Order them separately.

Table 2-12 lists the cable options available for BladeCenter H. You need two cables for each BladeCenter H chassis. For more information about power considerations for IBM BladeCenter H, see 2.14.3, "Power considerations" on page 197.

Connectors: The photographs in Table 2-12 show BladeCenter chassis connectors with silver rings. These connectors 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 2-12 Power cable options for IBM BladeCenter H

Part number	Feature code	Description	Photo
25R5783	6270	4.3 m 208 V Double 30A NEMA L6-30P	
25R5784	6271	4.3 m 230 V Dual 32A IEC 309 P+N+G/16A IEC 320-C20	
25R5785	6226	2.8 m 200-240 V Triple 16A IEC 320-C20	

Part number	Feature code	Description	Photo
25R5811	6273	4.3 m 220 V Double 30A KSC 8305 (for South Korea)	
25R5812	6272	4.3 m 230 V Dual 32A AS/NZS 3112/16A IEC 320-C20 (for Australia/NZ)	

For more information about power requirements for BladeCenter H, see 2.14.3, “Power considerations” on page 197.

2.2.5 Serial Port Breakout Cable

The Serial Port Breakout Cable provides 14 serial connections for terminal access, one to each supported blade server. See 2.13.1, “Serial Port Breakout Cable” on page 190 for more information.

2.2.6 Product publications

The following product publications are available for the BladeCenter H:

- ▶ *Installation and User's Guide - IBM BladeCenter H:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-63306>
- ▶ *Blade Center H - Problem Determination and Service Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-63570>

2.3 IBM BladeCenter HT

This section covers the following topics:

- ▶ BladeCenter HT product information
- ▶ BladeCenter HT interposers
- ▶ BladeCenter HT I/O topology
- ▶ Fan modules for BladeCenter HT
- ▶ Power modules for BladeCenter HT
- ▶ Serial Port Breakout Cable
- ▶ Product publications

2.3.1 BladeCenter HT product information

Figure 2-12 show the front view of the BladeCenter HT with the key chassis features.




Figure 2-12 The BladeCenter HT chassis

The following components are at the front of the BladeCenter HT:

- ▶ Twelve hot-swap 30-mm blade server bays
- ▶ Two hot-swap AMM bays
- ▶ Four hot-swap standard form-factor I/O module bays
- ▶ Four hot-swap high-speed form-factor I/O module bays
- ▶ Four hot-swap power module bays
- ▶ Two hot-swap media tray bays

Supported blade servers: For a list of blade servers supported in the BladeCenter HT, see Table 1-2 on page 3.

For a list of I/O modules supported in the BladeCenter HT, see Table 1-1 on page 2.

Figure 2-13 shows the rear view of the BladeCenter HT with the key chassis features.




Figure 2-13 Diagram of BladeCenter HT (AC model) rear view

The following components are at the rear of BladeCenter HT:

- ▶ Four hot-swap blower modules
- ▶ Two network clock bays
- ▶ Two multiplexer expansion module bays
- ▶ Alarm panel module bay
- ▶ Serial breakout connector
- ▶ Four power connectors

Table 2-13 lists the major features of the BladeCenter HT.

Table 2-13 BladeCenter HT features at a glance

Feature	Specifications	
Machine type	8740-2Rx ^a	8750-2Rx ^a
Rack form factor (H x D)	12U x 27.8 in. (706 mm)	12U x 27.8 in. (706 mm)
Number of media trays (std/max)	2 / 2	2 / 2
Optical drives in media tray standard	Not applicable ^b	Not applicable ^b
Diskette drives in media tray standard	Not applicable ^c	Not applicable ^c

Feature	Specifications	
Number of 30-mm Blade Server slots	12	12
Number of management modules (std/max)	2 / 2	2 / 2
Number of standard I/O module slots	4	4
Number of high-speed I/O module slots	4	4
Switch modules standard	None	None
Number of power supplies (std/max)	4 / 4	4 / 4
Power supply size standard	2535 W DC	3160 W AC
Number of blowers (std/max)	4 / 4	4 / 4
Dimensions	Height: 21.0 in. (528 mm) Width: 17.4 in. (442 mm) Depth: 27.8 in. (706 mm)	Height: 21.0 in. (528 mm) Width: 17.4 in. (442 mm) Depth: 27.8 in. (706 mm)

- a. The x is the country-specific letter (for example, EMEA MTM is 87402RG, and the US MTM is 87402RU).
- b. There is no optical drive bay in the media tray.
- c. There is no diskette drive bay in the media tray.

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 does not include, as standard, any I/O modules. You need to choose these I/O modules depending on your connectivity needs. An 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 - 10, and are used together with high-speed expansion cards installed into blade servers. I/O bays 1 - 4 can also be used for bridge modules.

BladeCenter HT includes, as standard, two multiplexer expansion modules for redundancy. These modules control USB connectors, video signals, console redirection, and status LEDs.

BladeCenter HT also includes an alarm panel module that provides telecom relay and LED alarm status, and a serial breakout connector for serial console functionality.

Table 2-14 shows the part numbers to use to order the BladeCenter HT chassis.

Table 2-14 BladeCenter HT model numbers for ordering

Description	Part number	Machine type
IBM BladeCenter HT	87402Rx ^a	8740-2Rx ^a
IBM BladeCenter HT	87502Rx ^a	8750-2Rx ^a

a. The x is the country-specific letter (for example, the EMEA part number is 87402RG, and the US part number is 87402RU).

The BladeCenter HT chassis includes, as standard, the following items:

- ▶ Two AMMs.
- ▶ Two media trays. Each media tray contains two USB 2.0 ports, status LEDs, and two CompactFlash bays.
- ▶ Two multiplexer expansion modules.
- ▶ Alarm panel module.
- ▶ Four 2,535 W DC (model 8740-2Rx) or 3,160 W AC (model 8750-2Rx) hot-swap power supply modules.
- ▶ Four hot-swap blower modules.
- ▶ Publications/CD package.

Notes:

- ▶ Advance Management Modules and trays: Models 8740-1Rx and 8750-1Rx included one AMM and one media tray.
- ▶ Internal power wiring: The internal power wiring, as shipped, differs inside the two chassis. Therefore, it is impossible to convert an existing chassis to the other *type* of power supply. An AC chassis cannot be changed into a DC chassis. A DC chassis cannot be changed in an AC chassis.

The BladeCenter HT does not include, as standard, any rails. Rails must be ordered separately.

Table 2-15 shows the part numbers to use to order the rail kit and other BladeCenter HT-related options if required.

Table 2-15 BladeCenter HT-related options part numbers and feature codes for ordering

Description	Part number	Feature code
IBM BladeCenter HT 4 Post Rack Mount Kit	42C5284	None
IBM BladeCenter HT Redundant Media Tray ^a	42C5305	4817
IBM BladeCenter HT 4 GB Compact Flash Option	42C5310	4822
IBM BladeCenter HT Bezel (incl Filter and Cable Mgmt Collar)	42C5278	4816

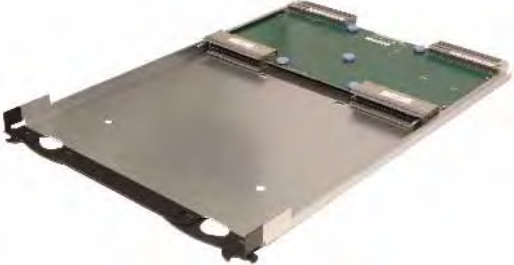
a. Only for models 8740-1Rx and 8750-1Rx.

Guidelines about how to install I/O module options are described in detail in 2.3.3, "BladeCenter HT I/O topology" on page 61.

2.3.2 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 extend the connections from the midplane to the rear of the module. Table 2-16 lists the possible options.

Table 2-16 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	

2.3.3 BladeCenter HT I/O topology

The 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 on the blade itself. For information about servers that support these cards, see Table 1-3 on page 4.

The BladeCenter HT chassis has a total of eight I/O bays (Figure 2-12 on page 57). Each blade bay has a total of eight dedicated connection paths to the I/O modules (Figure 2-14).




Figure 2-14 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-through 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 (Figure 2-14).
- ▶ 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 used for interswitch links.

The use of any of I/O modules in the BladeCenter HT chassis requires additional interposers. See 2.3.2, "BladeCenter HT interposers" on page 61 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 ISLs 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.

ISLs 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. For details about interposers, see 2.3.2, “BladeCenter HT interposers” on page 61.

ISL considerations:

- ▶ Currently, only the IBM Ethernet switches support the ISLs in BladeCenter HT (also require 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.

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 the CFFv or CIOv slot of a blade server, the I/O modules 3 and 4 must also be Fibre Channel-based. The reverse is also true: If you install Fibre Channel switches in bays 3 and 4, then any expansion card installed in the CFFv or CIOv in all blade servers in the chassis must be Fibre Channel cards.

See Table 2-31 on page 80 to match the expansion card ports with the corresponding I/O bay numbers. See Table 2-32 on page 82 to match the I/O bay number with the corresponding I/O module that can be installed into this bay.

Multiswitch interconnect modules: The BladeCenter H and BladeCenter HT chassis have different multiswitch interconnect modules (MSIMs). The MSIM for BladeCenter HT is called MSIM-HT, part number 44R5913, feature code 5491.

Speed considerations

The following speed considerations are for Fibre Channel switch modules.

4 Gb Fibre Channel switch modules speed considerations

Consider the following information:

- ▶ 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. 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 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. 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.

8 Gb Fibre Channel switch modules speed considerations

If 8 Gb Fibre Channel switch modules are installed into standard bays (bays 3 and 4) of a BladeCenter HT chassis or into MSIM-HT I/O bays (bays 8 and 10), then internal connections between blades and these switch modules operate at speeds up to 4 Gbps. 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.

2.3.4 Fan modules for BladeCenter HT

The BladeCenter HT unit includes, as standard, four hot-swap fan modules for N+1 cooling redundancy. The fan module for BladeCenter HT is shown in Figure 2-15.



Figure 2-15 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 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 run faster, increasing their speed as required to control the internal BladeCenter HT unit temperature.

For more information, see *Installation and User's Guide - IBM BladeCenter HT* at the following address:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5071317>

2.3.5 Power modules for BladeCenter HT

The BladeCenter HT chassis is available in two types with either DC power (8740) or AC power (8750) capabilities.

Power supplies: After a decision is made for one type of power supply, a conversion to the other type of power supply is impossible. If an initial configuration includes AC power supplies, it cannot be changed to DC power.

Both AC and DC power modules for BladeCenter HT look the same (Figure 2-16).



Figure 2-16 IBM BladeCenter HT Power Supply

The BladeCenter HT unit includes two (model 1Rx) or four (model 2Rx) DC (8740) or AC (8750) hot-swap power modules. The model 1Rx 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 2-17 on page 66).

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 three-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 (upper right connector) supplies power to power module bay 1, power connector 2 (lower left connector) supplies power to power module 2, power connector 3 (upper left connector) supplies power to the power module 3, and power connector 4 (lower right connector) supplies power to the power module 4 (Figure 2-13 on page 58).

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 - 6 and I/O modules in I/O module bays 1 - 4. Optional power modules are needed in power module bays 3 and 4 if you install blade servers in blade bays 7 - 12 or if you install I/O modules in high-speed I/O module bays 7 - 10 (Table 2-17).

Table 2-17 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 - 4, and to blade bays 1 - 6
3 and 4	Provides power to high-speed I/O modules 7 - 10, and blade bays 7 - 12

Table 2-18 lists the power supply options available for the IBM BladeCenter HT chassis.

Table 2-18 Power supply options for BladeCenter HT (models 8740-1Rx and 8750-1Rx only)

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 2.8 m 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.

For more information about power considerations for BladeCenter HT, see 2.14.3, "Power considerations" on page 197.

2.3.6 Serial Port Breakout Cable

The Serial Port Breakout Cable provides 12 serial connections for terminal access, one to each supported blade server. For details, see 2.13.1, "Serial Port Breakout Cable" on page 190.

2.3.7 Product publications

The following product publications are available for the BladeCenter HT:

- ▶ *Installation and User's Guide - IBM BladeCenter HT:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5071317>
- ▶ *IBM BladeCenter HT - Problem Determination and Service Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5071591>

2.4 IBM BladeCenter S

This section covers the following topics:

- ▶ BladeCenter S product information
- ▶ BladeCenter S I/O topology
- ▶ Fan modules for BladeCenter S
- ▶ Power modules for BladeCenter S
- ▶ BladeCenter S disk storage modules
- ▶ BladeCenter S Serial Pass-thru Module
- ▶ BladeCenter S Office Enablement Kit
- ▶ Product publications

2.4.1 BladeCenter S product information

Figure 2-17 shows the front view of the BladeCenter S with the key chassis features.




Figure 2-17 BladeCenter S front view showing the key features of the chassis

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

- ▶ A media tray at the front right, with an open bay for optical drive, two USB 2.0 ports, and a system status LED panel.
- ▶ 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.
- ▶ Two bays for battery modules used with SAS RAID Controller Module.

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

Figure 2-18 shows the rear view of the BladeCenter S with the key chassis features.




Figure 2-18 BladeCenter S rear view showing the key features of the chassis

The following components are at the rear of the BladeCenter S chassis:

- ▶ One pair of 950/1450 W power modules. An additional power module option (containing two 950/1450 W power modules) is available.
- ▶ Four hot-swap blower modules.
- ▶ One hot-swap management-module bay.
- ▶ Four hot-swap I/O bays for standard form factor I/O modules.
- ▶ One hot-swap bay for Serial Pass-thru Module.
- ▶ Rear system status LED panel.

Support I/O modules: For a list of I/O modules supported in the BladeCenter S, see Table 1-2 on page 3.

Table 2-19 lists the major features of the BladeCenter S.

Table 2-19 BladeCenter S features at a glance

Feature	Specifications
Machine type	8886-1Tx ^a
Rack form factor (H x D)	7U x 28.9 in. (733.4 mm)
Optical drives in media tray standard	None ^b
Diskette drives in media tray standard	Not applicable ^c
Number of 30-mm Blade Server slots	6
Number of management modules (std/max)	1 / 1
Number of standard I/O module slots	4
Switch modules standard	None
Number of power supplies (standard/maximum)	2 / 4
Power supply size standard	950/1450 W AC
Input Power Connector	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.)

a. The x is the country-specific letter (for example, the EMEA MTM is 8886-1TG, and the US MTM is 8886-1TU).

b. The optical drive is optional.

c. There is no diskette drive bay on the media tray.

The BladeCenter S chassis allows 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 includes, as standard, an AMM. This module manages the chassis, and provides the local KVM function. Only one AMM is supported with a BladeCenter S chassis.

The BladeCenter S chassis does not include, as standard, any I/O modules. You choose these I/O modules based on your connectivity needs. An ESM is required in I/O module bay 1, to enable the use of both Ethernet ports on a blade server. The I/O modules that are 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 used with an Ethernet-compatible I/O module when 2 / 4 port Gigabit Ethernet Expansion Cards are installed in blade servers.

The chassis does not include any storage modules.

Guidelines about how to install I/O module options are in 2.4.2, “BladeCenter S I/O topology” on page 71.

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

Table 2-20 shows the part numbers to use to order the BladeCenter S chassis.

Table 2-20 IBM BladeCenter S part numbers for ordering

Description	Part number	Machine type
IBM BladeCenter S	88861Tx ^a	8886-1Tx ^a

a. The x is the country-specific letter (for example, the EMEA part number is 88861TG, and the US part number is 88861TU).

The BladeCenter S chassis includes, as standard, the following items:

- ▶ One AMM
- ▶ Two hot-swap power supply modules (950/1450 W AC)
- ▶ Four hot-swap blower modules
- ▶ Rack mount kit
- ▶ Publications/CD package
- ▶ Two 2.5 m rack jumper power cords (IEC 320 C19 - C20)
- ▶ Four country-specific power cords (except EMEA models, which do not ship with country-specific power cords)

The 8886-1Tx model does not include any optical drive. An optical drive must be ordered separately if required. Table 2-21 lists the part numbers to use to order the optical drive.

Table 2-21 Optical drive part numbers and feature codes for ordering

Description	Part number	Feature code
UltraSlim Enhanced SATA Multi-burner	46M0902	4163
UltraSlim Enhanced SATA CD-RW / DVD-ROM Combo	46M0901	4161

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

2.4.2 BladeCenter S I/O topology

The BladeCenter S chassis provides connection paths between server blade bays and I/O or switch bays (Figure 1-1 on page 17 and Figure 1-2 on page 17) through a hardwired dual redundant midplane.

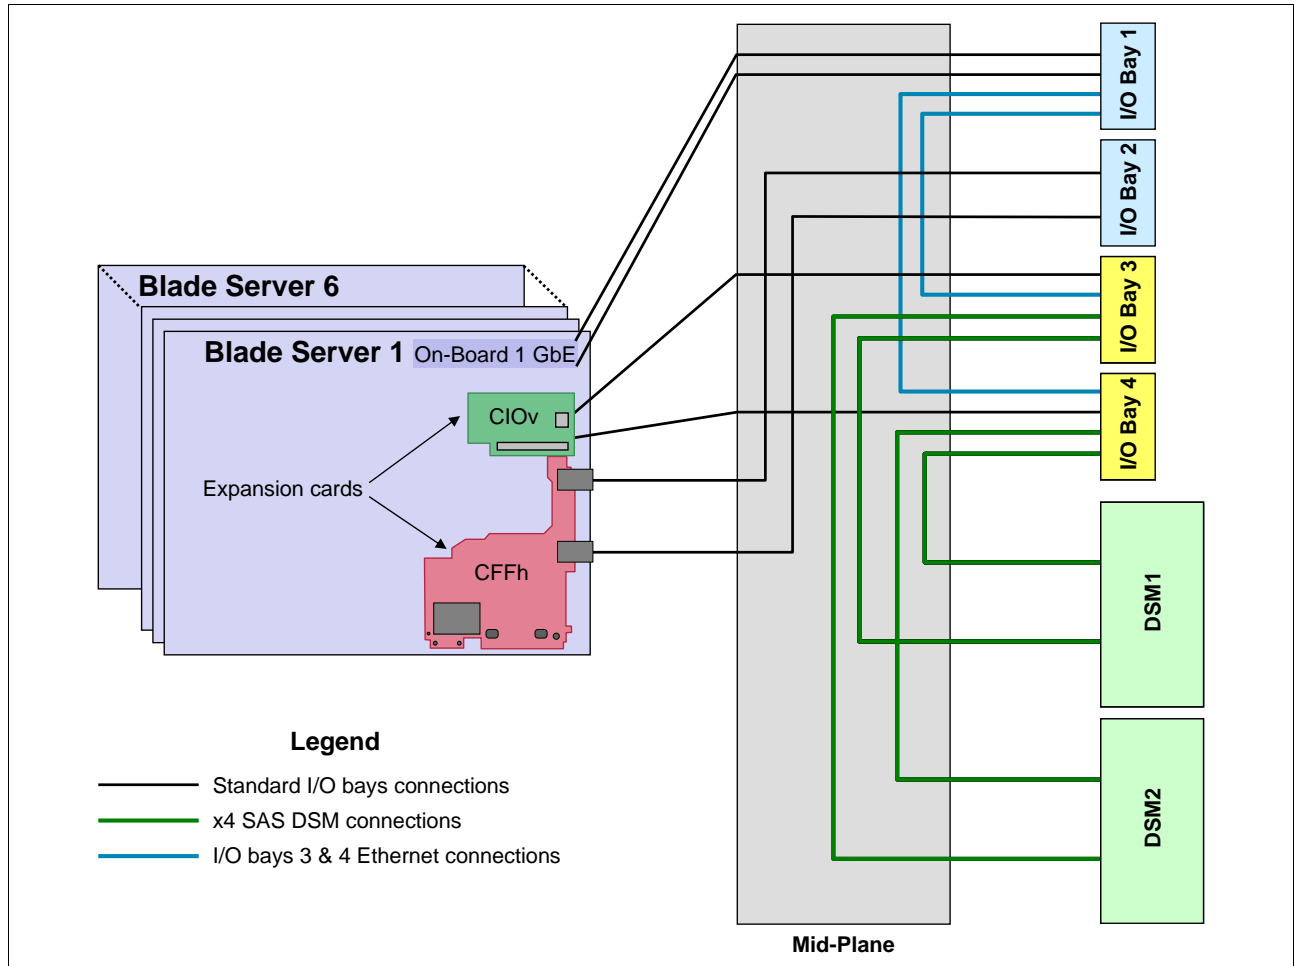


Figure 2-19 BladeCenter S I/O connection paths

The I/O topology is shown in Figure 2-19. Each of six blade bays has 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 six 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, see 5.10, "IBM BladeCenter S integrated storage" on page 475.

Each blade server has at least two Ethernet ports (Ethernet 1 and Ethernet 2). Both of them are routed to I/O bay 1 (which 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 that 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 a 2/4 Port Ethernet Expansion Card (CFFh) is installed into a 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.

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

Table 2-22 lists the mappings of onboard Ethernet ports to the ports of switch module in I/O bay 1.

Table 2-22 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 the CFFv or CIOv slot of a blade server, the I/O modules 3 and 4 must also be Fibre Channel-based. The reverse is also true: If you install Fibre Channel switches in bays 3 and 4, then any expansion card installed in the CFFv or CIOv in all blade servers in the chassis must be Fibre Channel cards.

See Table 2-31 on page 80 to match the expansion card ports with corresponding I/O bay numbers. See Table 2-32 on page 82 to match the I/O bay number with the corresponding I/O module that can be installed into this bay.

Unsupported I/O modules:

- ▶ Cisco Catalyst 3110G and 3110X ESMs are not supported in BladeCenter S.
- ▶ Twenty-port versions of FC switches (except Intelligent Optical Pass-thru modules) are not supported in BladeCenter S.

2.4.3 Fan modules for BladeCenter S

The BladeCenter S chassis includes 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 in the power supply modules.

Each fan module contains two fans. Figure 2-20 shows the fan module for BladeCenter S.




Figure 2-20 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.

For more information, see *Installation and User's Guide - IBM BladeCenter S*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5073635>

2.4.4 Power modules for BladeCenter S

The BladeCenter S chassis supports up to four auto-sensing power modules, which can support either 110 V or 220 V AC power. Two power modules are standard, and a maximum of four power modules are supported.

Figure 2-21 shows one of the power modules.




Figure 2-21 BladeCenter S power module

There are two power supplies for BladeCenter S based on the input connector (Table 2-23).

Table 2-23 Power supply connectors by chassis model

BladeCenter S model	Input connector	Availability
8886-1MY	IEC 320 C20	Worldwide
8886-1NG	IEC 320 C14	EMEA
8886-1TY	IEC 320 C14	Worldwide

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 on) 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. For more information about this topic, see “IBM BladeCenter S specific considerations” on page 200.

Use the BladeCenter Power Configurator, located at the following address, 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 is 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. 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 do not illuminate. The AMM also posts messages to the error log.

The BladeCenter S 8886-1MY includes, as standard, two power supplies 2.5 m IEC 320 C19 - C20 power jumper cords and four country-specific power cords (except for model 8886-1MG for the EMEA geographical area, which does not include country-specific power cords).

BladeCenter S 8886-1NG and 8886-1TY includes two 2 m IEC 320 C13 - C14 power jumper cords.

Table 2-24 lists power supply and power cord options available for BladeCenter S models.

Table 2-24 Power supply and power cable options for IBM BladeCenter S

Part number	Feature code	Description
43W3582	1992	950 W/1450 W Auto-Sensing Power Supplies 3 and 4 ^a
46C7438	2102	C14 950 W/1450 W Auto-Sensing Power Supplies 3 and 4 ^b
40K9766	None	C19 4.3 m Power Cord: Europe
40K9767	None	C19 4.3 m Power Cord: UK
40K9768	None	C19 4.3 m Power Cord: Italy
40K9769	None	C19 4.3 m Power Cord: Switzerland/Denmark
40K9770	None	C19 4.3 m Power Cord: South Africa
40K9771	None	C19 4.3 m Power Cord: Israel
39Y7917	None	2.8 m, 10A/230V, C13 to CEE7-VII (Europe) Power Cord
39Y7922	None	2.8 m, 10A/230V, C13 to SABS 164 (South Africa) Power Cord
39Y7923	None	2.8 m, 10A/230V, C13 to BS 1363/A (UK) Power Cord
39Y7921	None	2.8 m, 10A/230V, C13 to CEI 23-16 (Italy) Power Cord
39Y7920	None	2.8 m, 10A/230V, C13 to SI 32 (Israel) Power Cord
39Y7918	None	2.8 m, 10A/230V, C13 to DK2-5a (Denmark) Power Cord
39Y7919	None	2.8 m, 10A/230V, C13 to SEV 1011 (Sws) Power Cord

a. This option includes two 2.5 m IEC 320 C19 - C20 power jumper cords.

b. This option includes two 2.0 m IEC 320 C13 - C14 power jumper cords.

2.4.5 BladeCenter S disk storage modules

The BladeCenter S chassis supports up to two disk storage modules (DSM). These modules provide integrated SAS storage functionality to the BladeCenter S chassis. No disk storage modules are standard with the BladeCenter S chassis.

The storage module is a collection of disk drives, which is made accessible to blade servers through a SAS I/O module (see 2.10, “SAS I/O modules” on page 168) and SAS cards installed in the blades.

Two disk storage modules are available:

- ▶ 6-disk storage module (Figure 2-22 on page 76): Contains up to six 3.5-inch hot-swap hard disk drives, for a total of 12 internal drives with two 6-disk storage modules installed. The storage module supports SAS, SATA, and nearline SAS (NL SAS) drives. Intermixing different types of drives within the same storage module is supported.
- ▶ 12-disk storage module: Contains up to twelve 2.5-inch hot-swap hard disk drives, for a total of 24 internal drives with two 12-disk storage modules installed. The storage module supports only SAS drives.

Drive support: The SAS RAID Controller Module (43W3584) supports only SAS and NL SAS drives. The SAS Connectivity Module (39Y9195) supports all drive types.

The BladeCenter S chassis can accommodate up to six 3.5-inch disk drives in each 6-disk storage module or up to twelve 2.5-inch disk drives in each 12-disk storage module. You can easily and quickly assign the drives directly to blades using built-in predefined configurations or through user-defined custom configurations.

Figure 2-22 shows the IBM BladeCenter S 6-Disk Storage Module.



Figure 2-22 BladeCenter S 6-Disk Storage Module

Table 2-25 shows BladeCenter S 6-Disk Storage Module ordering information.

Table 2-25 IBM BladeCenter S 6-Disk Storage Module part numbers for ordering

Description	Part number	Feature code
IBM BladeCenter S 6-Disk Storage Module	43W3581	4545

Table 2-26 lists hard disk drives supported by BladeCenter S BladeCenter S 6-Disk Storage Modules and available for order.

Table 2-26 Ordering part numbers for disk drives supported by BladeCenter S 6-disk storage module

Description	Part number	Feature code
300 GB 15 K 6 Gbps SAS 3.5-inch Hot-Swap HDD	44W2234	5311
450 GB 15 K 6 Gbps SAS 3.5-inch Hot-Swap HDD	44W2239	5312
600 GB 15 K 6 Gbps SAS 3.5-inch Hot-Swap HDD	44W2244	5313
IBM 1 TB 7200 Dual Port SATA 3.5-inch HS HDD ^a	43W7630	5561
IBM 1 TB 7.2K 6 Gbps NL SAS 3.5-inch HS HDD	42D0777	5418
IBM 2 TB 7.2K 6 Gbps NL SAS 3.5-inch HS HDD	42D0767	5417

a. This drive is not supported by SAS RAID Controller Module (43W3584, feature code 3734).

Drive support: 6-Disk Storage Module supports 3Gb SAS topology only. 6Gb SAS drives are supported, but will run at 3Gb.

Figure 2-23 shows the 12-Disk Storage Modules installed in the BladeCenter S chassis.



Figure 2-23 BladeCenter S with two 1-Disk Storage Modules

Table 2-27 shows BladeCenter S 12-Disk Storage Module ordering information.

Table 2-27 IBM BladeCenter S 12-Disk Storage Module part numbers for ordering

Description	Part number	Feature code
IBM BladeCenter S 12-Disk Storage Module	49Y3234	A3KS

Table 2-28 lists hard disk drives supported by BladeCenter S BladeCenter S 12-Disk Storage Modules and available for order.

Table 2-28 Ordering part numbers for disk drives supported by BladeCenter S 12-disk storage module

Description	Part number	Feature code
IBM 300GB 2.5in SFF 10K 6Gbps HS SAS HDD	90Y8877	A2XC
IBM 600GB 2.5in SFF 10K 6Gbps HS SAS HDD	90Y8872	A2XD
IBM 900GB 2.5in SFF HS 10K 6Gbps SAS HDD	81Y9650	A282
IBM 300GB 2.5in SFF HS 15K 6Gbps SAS HDD	81Y9670	A283

Drive and DSM support:

- ▶ Drive support: 12-Disk Storage Module supports 3Gb SAS topology only. 6Gb SAS drives are supported, but will run at 3Gb.
- ▶ DSM support: The SAS RAID Controller Module (43W3584) supports 6-disk storage modules and 12-disk storage modules. The SAS Connectivity Module (39Y9195) supports 6-disk storage modules only.

For more information about BladeCenter S integrated storage planning, see 5.10, “IBM BladeCenter S integrated storage” on page 475.

2.4.6 BladeCenter S Serial Pass-thru Module

The Serial Pass-thru Module provides six serial port connectors that can be used to directly attach to each blade server in the BladeCenter S chassis through a four-wire serial RJ-45 connector. Port connector links bypass the AMM and provide a dedicated link directly to each blade. See 2.13.1, “Serial Port Breakout Cable” on page 190 for details.

2.4.7 BladeCenter S Office Enablement Kit

The BladeCenter S Office Enablement Kit is an enclosure designed for the BladeCenter S chassis for use in offices without a dedicated server room, or where the dust level is high. Figure 2-24 shows the enclosure with the BladeCenter S chassis and the Flat Panel Monitor kit installed.



Figure 2-24 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 setup leaves an extra 4U of space to hold other rack devices.

Table 2-29 shows BladeCenter S Office Enablement Kit ordering information.

Table 2-29 BladeCenter S Office Enablement Kit part numbers for ordering

Description	Part number	Feature code
BladeCenter S Office Enablement Kit	201886X	None

The Office Enablement Kit features:

- ▶ An acoustical module

The Office Enablement Kit includes an acoustical module that helps to quiet the BladeCenter S 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 includes 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 includes 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 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).

Table 2-30 shows options related to the Office Enablement Kit.

Table 2-30 BladeCenter S Office Enablement Kit related options

Description	Part number	Feature code
IBM BladeCenter Airborne Contaminant Filter	43X0340	4024
IBM BladeCenter Airborne Contaminant Replacement Filter (4-Pack)	43X0437	4025

The enclosure has the following approximate dimensions:

- ▶ Height: 24 in.
- ▶ Width: 24 in.
- ▶ Depth: 42 in.

2.4.8 Product publications

The following product publications are available for the BladeCenter S:

- ▶ *Planning Guide - IBM BladeCenter S:*

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5073632>

- ▶ *Installation and User's Guide - IBM BladeCenter S:*

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5073635>

- ▶ *Problem Determination and Service Guide:*

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5076785>

- ▶ *Office Enablement Kit Installation and User's Guide:*

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5073634>

2.5 Chassis I/O bay and expansion card port mappings

The various BladeCenter chassis have different numbers and types of I/O bays, as follows:

- ▶ BladeCenter S, E, and T have four standard I/O bays (1, 2, 3, and 4).
- ▶ BladeCenter H has four standard I/O bays (1, 2, 3, and 4), two bridge bays (5 and 6), and four high-speed bays (7, 8, 9, and 10).
- ▶ BladeCenter HT has four standard I/O bays (1, 2, 3, and 4) and four high-speed bays (7, 8, 9, and 10).

Table 2-31 matches the expansion card ports with the corresponding I/O bay numbers when used in different BladeCenter chassis. For a list of I/O modules compatible with the expansion card installed in a blade server, see the following tables:

- ▶ Ethernet compatibility: See Table 1-4 on page 6 and Table 1-3 on page 4.
- ▶ Fibre Channel compatibility: See Table 1-6 on page 9.
- ▶ InfiniBand compatibility: See Table 1-7 on page 9.
- ▶ SAS compatibility: See Table 1-8 on page 10.

Table 2-31 Matching I/O bay numbers with I/O expansion card ports

Part number	Feature code	Expansion card	I/O bay number									
			1	2	3	4	5	6	7	8	9	10
Ethernet expansion cards												
None	None	Integrated dual-port Gigabit Ethernet ^a	Y	Y	N	N	N	N	N	N	N	N
None	None	Integrated dual-port 10 Gb Ethernet	N	N	N	N	N	N	Y	N	Y	N
44W4475	1039	Ethernet Expansion Card (CIOv)	N	N	Y	Y	N	N	N	N	N	N
44W4479	5476	2/4 Port Ethernet Exp. Card (CFFh) ^b	N	Y	N	N	N	N	Y	Y	Y	Y
00Y3270 ^c	A3JC	QLogic Ethernet and 8 Gb FC (CFFh) ^d	N	N	N	N	N	N	Y	Y	Y	Y
00Y3280 ^e	A3JB	QLogic 2-port 10 Gb CNA (CFFh)	N	N	N	N	N	N	Y	N	Y	N
00Y3332	A4AC	QLogic 10Gb Virtual Fabric Adapter (CFFh)	N	N	N	N	N	N	Y	N	Y	N
00Y5618	A4AD	QLogic 10 Gb Virtual Fabric CNA (CFFh)	N	N	N	N	N	N	Y	N	Y	N
46M6164	0098	Broadcom 10 Gb Gen 2 4-port Eth. (CFFh)	N	N	N	N	N	N	Y	Y	Y	Y
81Y3133	A1QR	Broadcom 2-pt 10 Gb Virtual Fabric Adpt. (CFFh)	N	N	N	N	N	N	Y	N	Y	N
81Y1650	5437	Brocade 2-port 10 Gb CNA (CFFh)	N	N	N	N	N	N	Y	N	Y	N
00Y3266	A3NV	Emulex Virtual Fabric Adapter II (CFFh)	N	N	N	N	N	N	Y	N	Y	N
00y3264	A3NW	Emulex Virtual Fabric Adpt. Advanced II (CFFh)	N	N	N	N	N	N	Y	N	Y	N
81Y3120	A287	Emulex 10 GbE VFA II (CFFh) for HS23	N	N	N	N	N	N	N	Y	N	Y
90Y9332	A2ZN	Emulex 10 GbE VFA Adv. II (CFFh) for HS23	N	N	N	N	N	N	N	Y	N	Y
42C1810	3593	Intel 10 Gb 2-port Ethernet (CFFh)	N	N	N	N	N	N	Y	N	Y	N
90Y3570	A1NW	Mellanox 2-port 10 Gb Ethernet (CFFh)	N	N	N	N	N	N	Y	N	Y	N

Part number	Feature code	Expansion card	I/O bay number									
			1	2	3	4	5	6	7	8	9	10
Fibre Channel Expansion Cards												
43W6859	2994	Emulex 4 Gb FC Expansion Card (CFFv)	N	N	Y	Y	N	N	N	N	N	N
46M6065	3594	QLogic 4 Gb FC Expansion Card (CIOv)	N	N	Y	Y	N	N	N	N	N	N
44X1945	1462	QLogic 8 Gb FC Expansion Card (CIOv)	N	N	Y	Y	N	N	N	N	N	N
46M6140	3598	Emulex 8 Gb FC Expansion Card (CIOv)	N	N	Y	Y	N	N	N	N	N	N
00Y3270 ^c	A3JC	QLogic Ethernet and 8 Gb FC (CFFh) ^d	N	N	N	N	N	N	Y	Y	Y	Y
SAS Expansion Cards												
44E5688 ^f	A3J9	SAS Expansion Card (CFFv)	N	N	Y	Y	N	N	N	N	N	N
43W4068	1041	SAS Connectivity Card (CIOv)	N	N	Y	Y	N	N	N	N	N	N
InfiniBand Expansion Cards												
46M6001	0056	2-Port 40 Gb InfiniBand Exp. Card (CFFh)	N	N	N	N	N	N	Y	N	Y	N

- a. In BladeCenter S chassis, both ports on this expansion card are routed to I/O bay 1.
b. I/O bay 1 is only supported when this card is installed in a blade server installed in BladeCenter S chassis.
c. Replaces 44X1940.
d. Requires MSIM (39Y9314, feature code 1465) or MSIM-HT (44R5913, feature code 5491) to be installed in high-speed I/O bays 7 and 8 or 9 and 10 or both.
e. Replaces 42C1830.
f. Replaces 39Y9190.

Table 2-32 on page 82 matches the I/O bay number with the corresponding I/O module that can be installed into this bay. For a list of I/O modules compatible with the expansion card installed in a blade server, see the following tables:

- Ethernet compatibility: Table 1-4 on page 6 and Table 1-3 on page 4
- Fibre Channel compatibility: Table 1-6 on page 9
- InfiniBand compatibility: Table 1-7 on page 9
- SAS compatibility: Table 1-8 on page 10

Table 2-32 Matching I/O bay numbers with supported I/O modules

Part number	Feature code	I/O Module	I/O bay number									
			1	2	3	4	5	6	7	8	9	10
Ethernet Switch Modules												
46C9272 ^a	A3FE	Cisco Catalyst Switch Module 3012 ^b	Y	Y	Y	Y	N	N	Y	Y	Y	Y
00Y3254 ^c	A3FD	Cisco Catalyst Switch Module 3110G ^{b,d}	Y	Y	Y	Y	N	N	Y	Y	Y	Y
00Y3250 ^e	A3FC	Cisco Catalyst Switch Module 3110X ^{b,d}	Y	Y	Y	Y	N	N	Y	Y	Y	Y
46C9270 ^f	A3FF	Cisco Nexus 4001I Switch Module	N	N	N	N	N	N	Y	Y	Y	Y
39Y9324	1484	IBM Server Connectivity Module ^b	Y	Y	Y	Y	N	N	Y	Y	Y	Y
32R1860	1495	IBM L2/3 Copper Gigabit Ethernet Switch ^b	Y	Y	Y	Y	N	N	Y	Y	Y	Y
32R1861	1496	IBM L2/3 Fiber Gigabit Ethernet Switch ^b	Y	Y	Y	Y	N	N	Y	Y	Y	Y
32R1859	1494	IBM Layer 2-7 Gigabit Ethernet Switch	Y	Y	Y	Y	N	N	N	N	N	N
44W4404	1590	IBM 1/10Gb Uplink Ethernet Switch ^b	Y	Y	Y	Y	N	N	Y	Y	Y	Y
46C7191	1639	IBM Virtual Fabric 10Gb Switch Module	N	N	N	N	N	N	Y	Y	Y	Y
69Y1909	7656	Brocade Converged 10GbE Switch ^g	N	N	N	N	N	N	Y		Y	
Fibre Channel switch modules												
42C1828	5764	Brocade Enterprise 20-port 8Gb SAN SM ^{d,h}	N	N	Y	Y	N	N	N	Y	N	Y
44X1920	5481	Brocade 20-port 8Gb SAN Switch Module ^{d,h}	N	N	Y	Y	N	N	N	Y	N	Y
44X1921	5483	Brocade 10-port 8Gb SAN Switch Module ^{d,h}	N	N	Y	Y	N	N	N	Y	N	Y
44E5696 ⁱ	A3FH	Cisco Systems 20-port 4Gb FC Switch ^{d,h}	N	N	Y	Y	N	N	N	Y	N	Y
44E5692 ^j	A3FG	Cisco Systems 10-port 4Gb FC Switch ^h	N	N	Y	Y	N	N	N	Y	N	Y
44X1905	5478	QLogic 20-Port 8Gb SAN Switch Module ^h	N	N	Y	Y	N	N	N	Y	N	Y
44X1907	5482	QLogic 8Gb Intelligent Pass-thru Module ^h	N	N	Y	Y	N	N	N	Y	N	Y
88Y6406	A24C	QLogic 20-Port 4/8Gb SAN Switch Module ^h	N	N	Y	Y	N	N	N	Y	N	Y
88Y6410	A24D	QLogic 4/8Gb Intelligent Pass-thru Module ^h	N	N	Y	Y	N	N	N	Y	N	Y
46M6172	4799	QLogic Virtual Fabric Extension Module	N	N	Y	N	Y	N	N	N	N	N
SAS modules												
39Y9195	2980	SAS Connectivity Module	N	N	Y	Y	N	N	N	N	N	N
43W3584	3734	SAS RAID Controller Module ^k	N	N	Y	Y	N	N	N	N	N	N
InfiniBand modules												
46M6005	0057	Voltaire 40 Gb InfiniBand Switch Module ^d	N	N	N	N	N	N	Y		Y	

Part number	Feature code	I/O Module	I/O bay number									
			1	2	3	4	5	6	7	8	9	10
Pass-through and interconnect modules												
44X1907	5482	QLogic 8Gb Intelligent Pass-thru Module ^h	N	N	Y	Y	N	N	N	Y	N	Y
88Y6410	A24D	QLogic 4/8Gb Intelligent Pass-thru Module ^h	N	N	Y	Y	N	N	N	Y	N	Y
44W4483	5452	Intelligent Copper Pass-thru Module ^b	Y	Y	Y	Y	N	N	Y	Y	Y	Y
39Y9314	1465	Multi-Switch Interconnect Module ^{g,l}	N	N	N	N	N	N	Y		Y	
44R5913	5491	Multi-Switch Interconnect Module for HT ^{g,m}	N	N	N	N	N	N	Y		Y	
46M6181	1641	10Gb Ethernet Pass-thru Module	N	N	N	N	N	N	Y	Y	Y	Y

a. Replaces 43W4395.

b. Requires MSIM (39Y9314, feature code 1465) or MSIM-HT (44R5913, feature code 5491) to be installed in high-speed I/O bays 7 and 8 or 9 and 10 or both when used in bays 7 - 10.

c. Replaces 41Y8523.

d. This I/O module is not supported in the BladeCenter S chassis.

e. Replaces 41Y8522.

f. Replaces 46M6071.

g. This module occupies two adjacent high-speed bays.

h. Requires MSIM (39Y9314, feature code 1465) or MSIM-HT (44R5913, feature code 5491) to be installed in high-speed I/O bays 7 and 8 or 9 and 10 or both when used in bay 8 or 10 or both.

i. Replaces 39Y9280.

j. Replaces 39Y9284.

k. This module is supported only in the BladeCenter S chassis.

l. This module is supported only with BladeCenter H chassis.

m. This module is supported only with BladeCenter HT chassis.

I/O module and expansion card support: See Table 2-83 on page 178 (MSIM) and Table 2-86 on page 182 (MSIM-HT) for a list of I/O modules and expansion cards supported by MSIM and MSIM-HT.

2.6 I/O modules overview

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 in several ways. The I/O topology subsection of each BladeCenter chassis section gives an overview of the internal I/O architecture and the supported combinations of I/O modules and expansion cards for it.

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

- ▶ ESMs (which include both standard and high-speed ESMs) are covered in 2.7, “Ethernet Switch Modules” on page 84.
- ▶ InfiniBand switch modules (which include high-speed InfiniBand switch modules) are covered in 2.8, “InfiniBand switch module” on page 150.
- ▶ Fibre Channel switch modules (which include standard FC switch modules) are covered in 2.9, “Fibre Channel switch modules” on page 153.

- ▶ SAS I/O modules (which include SAS Connectivity Module and SAS RAID Controller Module) are covered in 2.10, “SAS I/O modules” on page 168.
- ▶ Pass-through and interconnect modules (which include copper and optical pass-through modules, 4X InfiniBand Pass-thru Module, 10 Gb Ethernet Pass-thru Module, and multi-switch interconnect modules) are covered in 2.11, “Pass-through and interconnect modules” on page 173.

In addition, there are two other module types:

- ▶ Management modules provide system management functions and KVM multiplexing for all blade servers in the BladeCenter chassis. Every BladeCenter chassis has one or two of these modules. For more information, see 2.12, “Advanced Management Module” on page 185.
- ▶ Direct serial modules and cables provide hardwired and dedicated terminal console access to blade servers for management purposes. BladeCenter H, HT, and S offer this feature. See 2.13, “Direct serial connections” on page 190 for more details.

2.7 Ethernet Switch Modules

Various types of ESMs from several vendors are available for BladeCenter, and they support various network layers and services, as shown in Table 2-33.

Table 2-33 Ethernet Switch Modules

Part number	Feature code	Option description	Number x type of external ports	Network layers	Page
46C9272 ^a	A3FE	Cisco Catalyst Switch Module 3012	4 x Gigabit Ethernet	Layer 2/3	88
00Y3254 ^b	A3FD	Cisco Catalyst Switch Module 3110G	4 x Gigabit Ethernet, 2 x StackWise Plus	Layer 2/3	96
00Y3250 ^c	A3FC	Cisco Catalyst Switch Module 3110X	1 x 10 Gb Ethernet, 2 x StackWise Plus	Layer 2/3	104
46C9270 ^d	A3FF	Cisco Nexus 4001I Switch Module	6 x 10 Gb Ethernet 1 x Gigabit Ethernet	Layer 2	115
39Y9324	1484	IBM Server Connectivity Module	6 x Gigabit Ethernet	Layer 2	121
32R1860	1495	IBM L2/3 Copper Gigabit Ethernet Switch Module	6 x Gigabit Ethernet	Layer 2/3	124
32R1861	1496	IBM L2/3 Fiber Gigabit Ethernet Switch Module	6 x Gigabit Ethernet	Layer 2/3	127
32R1859	1494	IBM Layer 2-7 Gigabit Ethernet Switch Module	4 x Gigabit Ethernet	Layer 2/7	131
44W4404	1590	IBM 1/10 Gb Uplink Ethernet Switch Module	3 x 10 Gb Ethernet, 6 x Gigabit Ethernet	Layer 2/3	134
46C7191	1639	IBM Virtual Fabric 10 Gb Switch Module	10 x 10 Gb Ethernet, 1 x 1 Gb Ethernet	Layer 2/3	139
69Y1909	7656	Brocade Converged 10 GbE Switch Module	8 x 10 Gb Ethernet, 8 x 8 Gb FC	Layer 2/3	145

a. Replaces 43W4395.

b. Replaces 41Y8523.

c. Replaces 41Y8522.

d. Replaces 46M6071.

Layer 2 support means that the Ethernet switch provides basic connectivity. It can deal with frame headers that contain physical addresses, such as MAC addresses, 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 can deal with packet headers that contain logical addresses, such as IP addresses, 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, the switch 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 capability enables implementation of advanced features, such as server load balancing or application security.

To meet the requirements of a client's network environment, the ESMs must support certain network standards, protocols, and features that are required by the network design. These features can be grouped into several categories that include VLAN support, performance, high-availability and redundancy, management, security, quality of service (QoS), and routing.

Table 2-34 compares features and functions supported by various switches.

Table 2-34 Basic features supported by Ethernet Switch Modules

Feature	Cisco 3012	Cisco 3110G	Cisco 3110X	Cisco Nexus 4001I	IBM Server Connectivity Module	IBM L2/3 GbE	IBM L2/7 GbE	IBM 1/10 Gb Uplink	IBM Virtual Fabric 10 GbE Switch	Brocade Converged 10 GbE Switch
Ethernet ports										
Number of external ports	4	4	1	7	6	6	4	9	11	9
Speed of external ports, Mbps (G=Gbps)	10/ 100/ 1 G	10/ 100/ 1 G	10 Gb	6x1G/ 10G, 1x1G	10/ 100/ 1G	10/ 100/ 1G	10/ 100/ 1G	3x10G 6x10/ 100/ 1G	10x1G /10G 1x 1G	8x10G 1x 1G
Dedicated stacking ports	None	2 x Stack Wise Plus, 32 Gb	2 x Stack Wise Plus, 32 Gb	None	None	None	None	None	None	None
Number of internal ports	14	14	14	14	14	14	14	14	14	14
Speed of internal ports	1000	1000	1000	1000/ 10000	1000	1000	1000	1000	1000/ 10000	10000
Fibre Channel ports (for converged switches if applicable)										
Number of external ports	0	0	0	0	0	0	0	0	0	8
Speed of external ports, Mbps (G=Gbps)	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	8G

Feature	Cisco 3012	Cisco 3110G	Cisco 3110X	Cisco Nexus 4001I	IBM Server Connectivity Module	IBM L2/3 GbE	IBM L2/7 GbE	IBM 1/10 Gb Uplink	IBM Virtual Fabric 10 GbE Switch	Brocade Converged 10 GbE Switch
VLAN support										
Number of VLANs supported	1005	1005	1005	512	256	1024	1024	1024	1024	4096
802.1Q VLAN Tagging	Yes	Yes	Yes	Yes	Yes ^a	Yes	Yes	Yes	Yes	Yes
Cisco ISL Trunking	Yes	Yes	Yes	No	No	No	No	No	No	No
Performance										
Link Aggregation	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Jumbo frames	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
IGMP snooping	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Server load balancing	No	No	No	No	No	No	Yes	No	No	No
Content switching	No	No	No	No	No	No	Yes	No	No	No
High availability and redundancy										
Spanning tree protocol	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
Trunk failover (link state tracking)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
VRRP	No	No	No	No	No	Yes	Yes	Yes	Yes	No
HSRP	Yes	Yes	Yes	No	No	No	No	No	No	No
Management										
Serial port	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
Telnet/SSH	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SNMP protocol	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cisco CDP protocol	Yes	Yes	Yes	Yes	No	No	No	No	No	No
Cisco VTP protocol	Yes	Yes	Yes	No ^b	No	No	No	No	No	No
Serial over LAN support	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Switch stacking	No	Yes	Yes	No	No	No	No	Yes	Yes	No
Protected mode	Yes	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Yes

Feature	Cisco 3012	Cisco 3110G	Cisco 3110X	Cisco Nexus 4001I	IBM Server Connectivity Module	IBM L2/3 GbE	IBM L2/7 GbE	IBM 1/10 Gb Uplink	IBM Virtual Fabric 10 GbE Switch	Brocade Converged 10 GbE Switch
Security										
Port security (802.1x)	Yes	Yes	Yes	No	No	Yes	No	Yes	Yes	Yes
Access control lists (MAC-based)	Yes	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Yes
Access control lists (IP-based)	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
TACACS+	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
RADIUS	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
NAT	No	No	No	No	No	No	Yes	No	No	No
Quality of service (QoS)										
802.1p CoS	Yes	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Yes
IP ToS/DSCP	Yes	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Yes
Routing protocols										
RIP	Yes	Yes	Yes	No	No	Yes	Yes	Yes	Yes	No
OSPF	No	Yes ^C	Yes ^C	No	No	Yes	Yes	Yes	Yes	No
BGP	No	Yes ^C	Yes ^C	No	No	Yes	Yes	Yes	Yes	No
EIGRP	No	Yes ^C	Yes ^C	No	No	No	No	No	No	No
IPv6 support										
Host management (IPv6)	Yes	Yes	Yes	Yes	No	No	No	Yes	Yes	No
Static routes (IPv6)	No	Yes ^C	Yes ^C	No	No	No	No	Yes	Yes	No
RIP (IPv6)	No	Yes ^C	Yes ^C	No	No	No	No	No	No	No
OSPF v3 (IPv6)	No	Yes ^C	Yes ^C	No	No	No	No	Yes	Yes	No
BGP (IPv6)	No	No	No	No	No	No	No	No	No	No
EIGRP (IPv6)	No	Yes ^C	Yes ^C	No	No	No	No	No	No	No
Access control lists (IPv6)	No	Yes ^C	Yes ^C	No	No	No	No	No	No	No

Feature	Cisco 3012	Cisco 3110G	Cisco 3110X	Cisco Nexus 4001I	IBM Server Connectivity Module	IBM L2/3 GbE	IBM L2/7 GbE	IBM 1/10 Gb Uplink	IBM Virtual Fabric 10 GbE Switch	Brocade Converged 10 GbE Switch
Converged networking										
CEE/FCoE support	No	No	No	Yes ^d	No	No	No	No	Yes	Yes
vNIC Virtual Fabric mode ^e	No	No	No	No	No	No	No	No	Yes	No

a. VLAN tagging is only supported for external ports.

b. VTP mode is off and cannot be changed (VTP BPDUs are dropped on all interfaces).

c. An IP Services license is required for the Cisco Catalyst 3110 family of switches to support this feature.

d. A Software Upgrade License is required for Cisco Nexus 4001I switch to support this feature.

e. For more information, see 4.9, "BladeCenter Virtual Fabric vNIC solution" on page 439.

Switch module features: 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.

See 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.

2.7.1 Cisco Catalyst Switch Module 3012

The Cisco Catalyst Switch Module 3012 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 designed to help meet the rigors of a blade server-based application infrastructure. The switches are designed to deliver scalable, high-performance, and highly resilient connectivity while supporting ongoing IT initiatives around reducing server infrastructure complexity and TCO by seamlessly integrating into existing Cisco management networks.

Figure 2-25 shows the Cisco Catalyst Switch Module 3012.



Figure 2-25 Cisco Catalyst Switch Module 3012

Features support: Unlike the Cisco Catalyst 3110 family, which supports license upgrades, the Cisco Catalyst Switch Module 3012 supports only the IP Base software feature set.

Table 2-35 shows the part number to use to order the module.

Table 2-35 Cisco Catalyst Switch Module 3012 part number and feature code for ordering

Description	Part number	Feature code
Cisco Catalyst Switch Module 3012 for IBM BladeCenter	46C9272 ^a	A3FE

a. Replaces 43W4395.

Ordering the switch: 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).

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

- Ports
 - USB-style serial port: This Cisco console port offers an out-of-band management path. A USB to DB9 cable is used to connect the switch module to the PC. This cable is included with the switch.
 - Four external RJ-45 1000BASE-T connectors: These connectors are used to make 10/100/1000 Mbps connections to a backbone, end stations, and servers.

- Fourteen internal full-duplex Gigabit ports: One each is 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 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
 - Dynamic address learning for enhanced security
 - Support for 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 AMM 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 no cost GUI-based application tool to configure most features of this switch. For more information and to download CNA, go to the following address:
<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 providing 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 moving 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 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 processor 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 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 regardless of the authorized or unauthorized state of the port
 - IEEE 802.1X port-based authentication with guest VLAN to provide 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. This classification allows 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 various 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 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

Cisco switches: Unlike the Cisco Catalyst 3110 family of switches, 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.

Cisco switches: Unlike the Cisco Catalyst 3110 family of switches, the Cisco Catalyst 3012 module supports only 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 timeout 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 diagnostic tests to test the hardware functionality of the supervisor engine, modules, and switch while the switch is connected to a live network.
- Onboard 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, and 5 (100 m (328 ft.) maximum)
 - 100-ohm STP (100 m maximum)
 - 100BASE-TX
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568 100-ohm STP (100 m maximum)
 - 1000BASE-T
 - UTP Category 6.
 - UTP Category 5e (100 m maximum)
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568B 100-ohm STP (100 m 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

For more information, see the following documentation:

- ▶ *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 the following address:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5075938>

2.7.2 Cisco Catalyst Switch Module 3110G

The Cisco Catalyst Switch Module 3110G 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 designed to help meet the rigors of blade server-based application infrastructure. The switches are designed to deliver scalable, high performance, and 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 2-26.



Figure 2-26 Cisco Catalyst Switch Module 3110G

Table 2-36 shows the part number to use to order the module.

Table 2-36 Cisco Catalyst Switch Module 3110G part number and feature code for ordering

Description	Part number	Feature code
Cisco Catalyst Switch Module 3110G for IBM BladeCenter	00Y3254 ^a	A3FD

a. Replaces 41Y8523.

Ordering the module: 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).

The Cisco Catalyst Switch Module 3110G includes, as standard, the IP Base feature set software. Additional features, such as IPv6 forwarding and routing support, and advanced routing protocol support (EIGRP, OSPF, BGP, PIM, and so on) for standard IP (IPv4), require an additional license, as listed in Table 2-37.

Table 2-37 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

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

Licensing: The 3110-AISK9-LIC-I license is withdrawn from market, and its functionality is included in the IP Services license (IBM part number 43W4434 or Cisco part number 3110-IPS-LIC-I).

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

► Ports

- USB-style serial port: This Cisco console port offers an out-of-band management path. A USB to DB9 cable is used to connect the switch module to PC. This cable is included 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 includes 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 m cable)
 - CAB-STK-E-1M= (1 m cable)
 - CAB-STK-E-3M= (3 m cable)

For more information about VBS, see 4.7, “Virtual Blade Switch technology” on page 432.

- Fourteen internal full-duplex Gigabit ports: One each is 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 to enhance fault tolerance and to provide 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 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.
 - 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 and 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 AMM, 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 no cost GUI-based application tool to configure most features of this switch. For more information and download of CNA, go to the following address:
<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 stacks.
 - 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 providing 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 moving 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 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 processor 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 regardless 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 and class of service
 - 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. This setup allows 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 and 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 and routed backbones:
 - RIP Versions 1 and 2.
 - OSPF (an IP services feature set is required).
 - Enhanced IGRP (EIGRP) (an IP services feature set is required).
 - Border Gateway Protocol (BGP) Version 4 (an 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 (an IP services feature set is required).
 - VPNs (an IP services feature set is required).
 - Fallback bridging for forwarding non-IP traffic between two or more VLANs (an 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 (an IP services feature set is required).
 - Multicast Source Discovery Protocol (MSDP) for connecting multiple PIM-SM domains (an 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, and 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 (IP services feature set is required).
 - Support for EIGRP IPv6, which uses IPv6 transport, communicates with IPv6 peers, and advertises IPv6 routes (an IP services feature set is required).
 - Support for IPv6 Access Control Lists (ACLs) (an IP services feature set is required).

- IP unicast reverse path forwarding (unicast RPF) for confirming source packet IP addresses.
- Nonstop forwarding (NSF) awareness enables 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 learned or removed.
 - Syslog facility for logging system messages about authentication or authorization errors, resource issues, and timeout 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 diagnostic tests to test the hardware functionality of the supervisor engine, modules, and switch while the switch is connected to a live network.
 - Onboard 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).
- Network cables:
 - 10BASE-T:
 - UTP Category 3, 4, 5 (100 m (328 ft.) maximum)
 - 100-ohm STP (100 m maximum)
 - 100BASE-TX:
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568 100-ohm STP (100 m maximum)
 - 1000BASE-T:
 - UTP Category 6
 - UTP Category 5e (100 m maximum)
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568B 100-ohm STP (100 m 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

For more information, see the following documentation:

- ▶ *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 the following address:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5075938>

2.7.3 Cisco Catalyst Switch Module 3110X

The Cisco Catalyst Switch Module 3110X is an I/O module that is installed into a BladeCenter unit. This switch is a full wire-rated and 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 designed to help meet the rigors of blade server-based application infrastructure. The switches are designed to deliver scalable, high performance, and highly resilient connectivity while supporting ongoing IT initiatives around reducing server infrastructure complexity and TCO by seamlessly integrating into existing Cisco management networks.

Figure 2-27 shows the Cisco Catalyst Switch Module 3110X.



Figure 2-27 Cisco Catalyst Switch Module 3110X with optional X2 transceiver module installed

Table 2-38 shows the part number to use to order the module.

Table 2-38 Cisco Catalyst Switch Module 3110X part number and feature code for ordering

Description	Part number	Feature code
Cisco Catalyst Switch Module 3110X for IBM BladeCenter	00Y3250 ^a	A3FC

a. Replaces 41Y8522.

Ordering information: 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).

The Cisco Catalyst Switch Module 3110X comes standard with IP Base feature set software. Additional features, such as IPv6 forwarding and routing support, and advanced routing protocols support (EIGRP, OSPF, BGP, PIM, and so on) for standard IP (IPv4), require additional licenses, as listed in Table 2-39.

Table 2-39 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

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

Licensing: The 3110-AISK9-LIC-I license is withdrawn from market, and its functionality is included in the IP Services license (IBM part number 43W4434 or Cisco part number 3110-IPS-LIC-I).

The Cisco Catalyst Switch Module 3110X has one X2 slot. Optionally, this slot can be converted to one SFP+ or two SFP slots by respective X2 to SFP conversion modules to allow migration from X2 form-factor to SFP form-factor without changing the switch module. Conversion modules come without SFP+ transceivers; you can choose either SFP+ transceiver or SFP+ direct-attach cable (DAC) connectivity options, and you need to purchase them separately.

Two types of X2 to SFP+ converters are available:

- OneX Converter has one SFP+ slot to plug 10 Gb Ethernet SFP+ transceivers or cables.
- The TwinGig Converter has two SFP slots where you can plug in Gigabit Ethernet SFP transceivers.

Converter required: If you plan to use SFP or SFP+ transceivers or cables, a converter *must* be ordered separately (one converter per one 3110X).

Table 2-40 lists the part numbers to use to order the X2, SFP+ transceivers, and DAC cables from IBM or from Cisco Systems and authorized Cisco Systems resellers.

Table 2-40 X2 and SFP+ transceivers and copper cables for Cisco 3110X Switch Module

Description	IBM part number	IBM feature code	Cisco Systems part number
X2 Transceivers			
Cisco 10GBASE-CX4 X2 Transceiver (CX4, copper, InfiniBand 4X)	None	None	X2-10GB-CX4(=)
Cisco 10GBASE-SR X2 Transceiver (MMF, 850 nm, SC)	None	None	X2-10GB-SR(=)
Cisco 10GBASE-LX4 X2 Transceiver (MMF, 1310 nm, SC)	None	None	X2-10GB-LX4(=)
Cisco 10GBASE-LRM X2 Transceiver (MMF, 1310 nm, SC)	None	None	X2-10GB-LRM(=)
Cisco 10GBASE-LR X2 Transceiver (SMF, 1310 nm, SC)	None	None	X2-10GB-LR(=)
X2 to SFP+ Converters			
Cisco OneX Converter Module	88Y6066	A1A9	CVR-X2-SFP10G(=)
Cisco TwinGig Converter Module	None	None	CVR-X2-SFP(=)
10 Gb Ethernet SFP+			
Cisco 10GBASE-SR SFP+ Transceiver (MMF, 850 nm, LC)	88Y6054	A1A6	SFP-10G-SR(=)
DAC Cables			
Cisco 1 m 10G SFP+ Twinax cable assembly, passive	None	None	SFP-H10GB-CU1M(=)
Cisco 3 m 10G SFP+ Twinax cable assembly, passive	None	None	SFP-H10GB-CU3M(=)
Cisco 5 m 10G SFP+ Twinax cable assembly, passive	None	None	SFP-H10GB-CU5M(=)
Gigabit Ethernet SFP+			
Cisco 1000BASE-T SFP Transceiver (CAT 5 copper wire, RJ-45)	88Y6058	A1A7	GLC-T(=)
Cisco 1000BASE-SX SFP Transceiver (MMF, 850 nm, LC)	88Y6062	A1A8	GLC-SX-MM(=)
Cisco 1000BASE-LX/LH SFP Transceiver (MMF/SMF, 1300 nm, LC)	None	None	GLC-LH-SM(=)

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

- Ports:
 - USB-style serial port: This Cisco console port offers an out-of-band management path. A USB to DB9 cable is used to connect the switch module to PC. This cable is included with the switch.
 - One external 10 Gb Ethernet Module slot for making 10000 Mbps uplinks to 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.

- Two external high-speed StackWise Plus ports for switch module stacking to support Virtual Blade Switch (VBS) technology. Each 3110X switch module includes 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 m cable)
- CAB-STK-E-1M= (1 m cable)
- CAB-STK-E-3M= (3 m cable)

For more information about VBS, see 4.7, “Virtual Blade Switch technology” on page 432.

- Fourteen internal full-duplex Gigabit Ethernet ports (1000Base-X). Each one is connected to one blade server in the BladeCenter unit.
- One internal full-duplex 100 Mbps port connected to the management module.
- Performance features:
 - Fixed 10 Gbps speed on an external 10 Gb Ethernet port for maximum uplink bandwidth.
 - Up to 64 Gbps of throughput in a switch stack.
 - 10 Gb 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 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.
 - Dynamic address learning for enhanced security.
 - Supports multiple EtherChannel load balance algorithms (SMAC or DMAC, SIP or DIP, or 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 a switch module's console port.
- Supported by CiscoWorks management software.
- Protected Mode feature to isolate switch management from AMM, 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 no cost GUI-based application tool to configure most features of this switch. For more information and to download CNA, go to the following address:
<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 a 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 moving 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 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 processor 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 regardless of the authorized or unauthorized state of the port.
 - IEEE 802.1X port-based authentication with guest VLAN to provide limited services to non-802.1X-compliant users.
 - IEEE 802.1X accounting to track network usage.
- Quality of service and class of service:
 - 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. This configuration allows 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, and 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 (an IP services feature set is required).
 - Enhanced IGRP (EIGRP) (an IP services feature set is required).
 - Border Gateway Protocol (BGP) Version 4 (an 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 (an IP services feature set is required).
 - VPNs (an IP services feature set is required).
 - Fallback bridging for forwarding non-IP traffic between two or more VLANs (an 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 (an IP services feature set is required).
 - Multicast Source Discovery Protocol (MSDP) for connecting multiple PIM-SM domains (an 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, and 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, and RIP and OSPF) for forwarding IPv6 traffic through configured interfaces (an IP services feature set is required).
 - Support for EIGRP IPv6, which uses IPv6 transport, communicates with IPv6 peers, and advertises IPv6 routes (an IP services feature set is required).
 - Support for IPv6 Access Control Lists (ACLs) (an IP services feature set is required).
- IP unicast reverse path forwarding (unicast RPF) for confirming source packet IP addresses.
- Nonstop forwarding (NSF) awareness enables 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 learned or removed.
 - Syslog facility for logging system messages about authentication or authorization errors, resource issues, and timeout events.
 - Layer 2 trace route to identify the physical path that a packet takes from a source device to a destination device.
 - Online diagnostic tests to test the hardware functionality of the supervisor engine, modules, and switch while the switch is connected to a live network.
 - Onboard 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).

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

► Network cables:

– 10GBASE-SR

Table 2-41 lists the 10GBASE-SR cabling specifications.

Table 2-41 10GBASE-SR cabling specifications

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

– 10GBASE-LRM

Table 2-42 lists the 10GBASE-LRM cabling specifications.

Table 2-42 10GBASE-LRM cabling specifications

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

For more information, see the following documentation:

- ▶ *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 the following address:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5075938>

2.7.4 Cisco Nexus 4001I Switch Module

The Cisco Nexus 4001I Switch Module is a blade switch solution for the BladeCenter H and HT chassis, providing the server I/O solution required for high-performance, scale-out, virtualized, and non-virtualized x86 computing architectures. It is a line rate, extremely low-latency, non-blocking, Layer 2, 10-Gigabit Ethernet blade switch that is fully compliant with Fibre Channel over Ethernet (FCoE) and IEEE Data Center Bridging standards.

The Cisco Nexus 4001I enables a standards-based, high-performance Unified Fabric running over 10 Gigabit Ethernet in the blade server environment. This Unified Fabric enables consolidation of LAN traffic, storage traffic (IP-based such as iSCSI, NAS, and Fibre Channel SAN), and high-performance computing (HPC) traffic over a single 10 Gigabit Ethernet server network.

This offering works with BladeCenter Open Fabric Manager, providing all the benefits of I/O virtualization at 10 Gbps speed. Figure 2-28 shows the switch module.



Figure 2-28 Cisco Nexus 4001I Switch Module

Table 2-43 shows the part number to order these modules.

Table 2-43 Cisco Nexus 4001I Switch Module part numbers and feature codes for ordering

Description	Part number	Feature code
Cisco Nexus 4001I Switch Module for IBM BladeCenter	46C9270 ^a	A3FF

a. Replaces 46M6071.

Ordering information: 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 N4K-4001I-XPX).

The module part numbers include the following items:

- ▶ One Cisco Nexus 4001I Switch Module
- ▶ Cisco Console Cable RJ45-to-DB9
- ▶ One filler panel
- ▶ Important Notices document
- ▶ Documentation CD-ROM

Software Upgrade License for Cisco Nexus 4001I

The Cisco Nexus 4001I Switch Module is designed to support both 10 Gb Ethernet and Fibre Channel over Ethernet. The Software Upgrade License for Cisco Nexus 4001I, part number 49Y9983, enables the switch to work in FCoE mode. When connected to a converged adapter in the server, this switch can route CEE packets to an upstream FCoE switch, which can then route the packets to the LAN or SAN.

Table 2-44 shows the part numbers for the Software Upgrade License for Cisco Nexus 4001I.

Table 2-44 Software Upgrade License part numbers and feature codes for ordering

Description	Part number	Feature code	Cisco Systems part number
Software Upgrade License for Cisco Nexus 4001I	49Y9983	1744	N4K-4001I-SSK9

SFP+ Transceivers and Copper Cables

The Cisco Nexus 4001I Switch Module does not include either SFP+ Transceivers or SFP+ Copper Cables. These parts can be ordered either from IBM (see Table 2-43 on page 115) or directly from Cisco Systems or authorized Cisco Systems resellers, as listed in Table 2-45.

Table 2-45 SFP+ transceivers and copper cables for Cisco Nexus 4001I Switch Module

Description	IBM part number	IBM feature code	Cisco Systems part number
10 Gb Ethernet SFP+			
Cisco 10GBASE-SR SFP+ Transceiver (MMF, 850 nm, LC)	88Y6054	A1A6	SFP-10G-SR(=)
Cisco 10GBASE-LR SFP+ Transceiver (SMF, 1310 nm, LC)	None	None	SFP-10G-LR(=)
DAC Cables			
Cisco 1 m 10G SFP+ Twinax cable assembly, passive	None	None	SFP-H10GB-CU1M(=)
Cisco 3 m 10G SFP+ Twinax cable assembly, passive	None	None	SFP-H10GB-CU3M(=)
Cisco 5 m 10G SFP+ Twinax cable assembly, passive	None	None	SFP-H10GB-CU5M(=)
Gigabit Ethernet SFP+			
Cisco 1000BASE-T SFP Transceiver (CAT 5 copper wire, RJ-45)	88Y6058	A1A7	GLC-T(=)
Cisco 1000BASE-SX SFP Transceiver (MMF, 850 nm, LC)	88Y6062	A1A8	GLC-SX-MM(=)
Cisco 1000BASE-LX/LH SFP Transceiver (MMF/SMF, 1300 nm, LC)	None	None	GLC-LH-SM(=)

The following features and specifications are supported for the The Cisco Nexus 4001I Switch Module:

- ▶ Form-factor:
 - Single-height high-speed switch module.
- ▶ External ports:
 - Six 10 Gb SFP+ ports operating at wire speed. Also designed to support 1 Gb SFP if required, with the flexibility of mixing 1 Gb/10 Gb. Table 2-45 on page 116 lists the supported transceivers and cables.
 - One 10/100/1000 Mb Ethernet copper RJ-45 used for management.
 - An RS-232 RJ-45 connector for serial port that provides an additional means to configure the: switch module. The console cable is supplied with the switch module.
- ▶ Internal ports
 - Fourteen internal auto-negotiating ports: 1 Gb or 10 Gb to the server blades
 - Two internal full-duplex 100 Mbps ports connected to the management modules
- ▶ Scalability and performance:
 - Autosensing 1 Gb/10 Gb internal and external Ethernet ports for bandwidth optimization
 - Non-blocking architecture with wire-speed forwarding of traffic and full line rate performance of 400 Gbps full duplex
 - Forwarding rate of 300 million packets per second (mpps)
 - Low, predictable, and consistent latency of 1.5 μ s regardless of packet size, traffic pattern, or enabled features on a 10 Gigabit Ethernet interface
 - Media access control (MAC) address learning: Automatic update, supports up to 8 KB MAC addresses
 - EtherChannels and LACP (IEEE 802.3ad) link aggregation, up to 60 Gb of total uplink bandwidth per switch, up to seven trunk groups, and up to six ports per group
 - Support for jumbo frames (up to 9216 bytes)
 - Traffic suppression (unicast, multicast, and broadcast)
 - IGMP snooping to limit flooding of IP multicast traffic (IGMP V2 and V3)
 - Configurable traffic distribution schemes over EtherChannel links based on source/destination IP addresses, MAC addresses, or ports
 - Spanning Tree edge ports (formerly PortFast) for rapid STP convergence
- ▶ Availability and 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.
 - 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 providing 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 moving root and designated ports to the converging state.
- 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.
- ▶ VLAN support:
 - Up to 512 VLANs supported per switch; VLAN numbers ranging 1 - 4000
 - 802.1Q VLAN tagging support on all ports
 - Private VLANs
- ▶ Security:
 - VLAN-based, MAC-based, and IP-based access control lists (ACLs).
 - Role-based access control (RBAC) to restrict the authorization of the user to perform switch management functions
 - 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
- ▶ Quality of service:
 - Support for IEEE 802.1p CoS, IP ToS/DSCP, Protocol, IP Real Time Protocol, and ACL-based (MAC/IP source and destination addresses, VLANs) traffic classification and processing
 - Trust boundaries on incoming interfaces to automatically classify incoming packets into system classes based on packet's CoS value
 - Traffic shaping, guaranteeing bandwidth, or prioritizing based on defined policies
 - Up to eight egress queues per interface (one queue per CoS) for processing qualified traffic (support for strict priority and weighted round-robin (WRR) CoS policies)
- ▶ Fibre Channel over Ethernet:
 - Support for T11-compliant FCoE on all 10 Gb Ethernet interfaces.
 - FCoE Initialization Protocol (FIP) snooping to enable the switch to operate as a loss-less Ethernet bridge transparently forwarding FCoE packets.
 - 802.1Q VLAN tagging for FCoE frames.
 - Priority-based flow control (IEEE 802.1Qbb) simplifies management of multiple traffic flows over a single network link and creates lossless behavior for Ethernet by allowing class-of-service (CoS)-based flow control.
 - Enhanced Transmission Selection (IEEE 802.1Qaz) enables consistent management of QoS at the network level by providing consistent scheduling of different traffic types (IP, storage, and so on).
 - Data Center Bridging Exchange (DCBX) Protocol (IEEE 802.1AB) simplifies network deployment and reduces configuration errors by providing autonegotiation of IEEE 802.1 DCB features between the network interface card (NIC) and the switch and between switches.
- ▶ IPv6 support:

IPv6 host support (IPv6 unicast addressing, IPv6 traffic processing, and IPv6 applications support, including syslog server, RADIUS, TACACS+, NTP, telnet, and ssh). IPv6 traffic forwarding is not supported.

► Manageability:

- Command-line interface: You can configure switches using the CLI from an SSH V2 session, a Telnet session, or the console port. SSH provides a secure connection to the device.
- XML Management Interface over SSH: You can configure switches using the XML management interface, which is a programming interface based on the NETCONF protocol that complements the CLI functionality. For more information, see the *Cisco NX-OS XML Interface User Guide*:
http://www.cisco.com/en/US/docs/switches/datacenter/sw/nx-os/xml/user/guide/nxos_xml_interface.html
- Cisco Data Center Manager support.
- SNMP V1, 2, and 3 support.
- Protected Mode feature to isolate switch management from the AMM for increased security of the switch.
- Cisco Discovery Protocol (CDP) Versions 1 and 2 to aid in troubleshooting and reporting on misconfiguration of ports connecting to other devices supporting CDP.

► Monitoring:

- Switch LEDs for external port status and switch module status indication
- RMON
- Change tracking and remote logging with syslog feature
- Online diagnostic tests
- Cisco Fabric Services
- Session Manager
- SPAN for local traffic monitoring

► Special functions:

Serial over LAN (SOL)

The switch module supports the following IEEE standards:

- IEEE 802.1D: Spanning Tree Protocol
- IEEE 802.1p: CoS Prioritization
- IEEE 802.1Q: VLAN Tagging
- IEEE 802.1s: Multiple VLAN Instances of Spanning Tree Protocol
- IEEE 802.1w: Rapid Reconfiguration of Spanning Tree Protocol
- IEEE 802.3ab: 1000Base-T (10/100/1000 Ethernet over copper)
- IEEE 802.3ad: Link Aggregation Control Protocol (LACP)
- IEEE 802.3ae 10 Gb Ethernet over fiber optics
 - 10GBASE-SR
 - 10GBASE-LR
- IEEE 802.3z Gigabit Ethernet over fiber optics
 - 1000BASE-SX
 - 1000BASE-LX
- IEEE 802.3ab 1000BASE-T over copper twisted pair

Network cables requirements are as follows:

► 10GBASE-SR

Table 2-46 on page 120 lists the 10GBASE-SR cabling specifications.

Table 2-46 10GBASE-SR cabling specifications

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

► 10GBASE-LR

Table 2-47 lists the 10GBASE-LR cabling specifications.

Table 2-47 10GBASE-LR cabling specifications

Wavelength	Cable type	Core size (microns)	Modal bandwidth (MHz/km)	Maximum cable length
1310 nm	SMF	G.652	Not applicable	6.2 miles (10 km)

► 1000BASE-SX

Table 2-48 lists the 1000BASE-SX cabling specifications.

Table 2-48 1000BASE-SX cabling specifications

Wavelength	Cable type	Core size (microns)	Modal bandwidth (MHz/km)	Maximum cable length
850 nm	MMF	62.5	160	722 ft. (220 m)
850 nm	MMF	62.5	200	902 ft. (275 m)
850 nm	MMF	50	400	1640 ft. (500 m)
850 nm	MMF	50	500	1804 ft. (550 m)

► 1000BASE-LX

Table 2-49 lists the 1000BASE-LX cabling specifications.

Table 2-49 1000BASE-LX cabling specifications

Wavelength	Cable type	Core size (microns)	Modal bandwidth (MHz/km)	Maximum cable length
1310 nm	MMF	62.5	500	1804 ft. (550 m)
1310 nm	MMF	50.0	400	1804 ft. (550 m)
1310 nm	MMF	50.0	500	1804 ft. (550 m)
1310 nm	SMF	G.652	Not applicable	6.2 miles (10 km)

► 1000BASE-T

- UTP Category 6 (100 m maximum)
- UTP Category 5e (100 m maximum)
- UTP Category 5 (100 m maximum)
- EIA/TIA-568B 100-ohm STP (100 m maximum)

For more information, see the following documentation:

- ▶ *Configuration Guide for Cisco Nexus 4001I Switch Module*
- ▶ *Hardware Installation Guide for Cisco Nexus 4001I Switch Module*
- ▶ *Getting Started Guide for Cisco Nexus 4001I Switch Module*
- ▶ *Command Reference Guide for Cisco Nexus 4001I Switch Module*

These publications can be downloaded from the following address:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5082494>

2.7.5 IBM Server Connectivity Module

The IBM Server Connectivity Module for IBM BladeCenter is a switch that provides basic Layer 2 functionality. This device can be configured by a non-networking system administrator through a graphical user interface (GUI) or a 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 is 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 2-29).



Figure 2-29 IBM Server Connectivity Module

Table 2-50 shows the part number to use to order the module.

Table 2-50 IBM Server Connectivity Module part number and feature code for ordering

Description	Part number	Feature code
IBM Server Connectivity Module for IBM BladeCenter	39Y9324	1484

The following features and specifications are supported for the IBM Server Connectivity Module:

- ▶ Internal ports:
 - Fourteen 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 KB 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.
- ▶ A “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 m or 328.1 ft. maximum)
 - 100-ohm STP (100 m or 328.1 ft. maximum)
 - 100BASE-TX:
 - UTP Category 5 (100 m or 328.1 ft. maximum)
 - EIA/TIA-568 100-ohm STP (100 m or 328.1 ft. maximum)
 - 1000BASE-T:
 - UTP Category 6
 - UTP Category 5e (100 m or 328.1 ft. maximum)
 - UTP Category 5 (100 m or 328.1 ft. maximum)
 - EIA/TIA-568B 100-ohm STP (100 m 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

For more information, see the following documentation:

- ▶ *IBM Server Connectivity Module Installation Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-64155>
- ▶ *IBM Server Connectivity Module User's Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-64190>

Important: If you plan to use this switch with HC10 workstation blades, be sure to disable IGMP snooping. See IBM RETAIN® tip H192042 for details at the following address:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5073011>

2.7.6 IBM Layer 2/3 Copper GbE Ethernet Switch Module

The IBM Layer 2/3 Copper Gigabit Ethernet Switch Module is a switch option that enables administrators to consolidate full Layer 2 and Layer 3 LAN switching and routing capabilities into a BladeCenter chassis. Consolidation flattens the topology of the data center infrastructure and reduces the number of discrete devices, management consoles, and different management systems.

Figure 2-30 shows the switch module.



Figure 2-30 IBM L2/3 Copper GbE Switch Module

Table 2-51 lists the part number to use to order the module.

Table 2-51 IBM Layer 2/3 Copper GbE Switch Module part number and feature code for ordering

Description	Part number	Feature code
IBM Layer 2/3 Copper GbE Switch Module for IBM BladeCenter	32R1860	1495

The following features and specifications are supported for the IBM L2/3 GbE Switch Module:

- ▶ Internal ports:
 - Fourteen 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 10/1000/10000 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. The VLAN numbers range is 1 - 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 remarking 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 are supported.
 - Support for DHCP Relay.
- ▶ Manageability:
 - SNMP (V1 and V3)
 - HTTP browser GUI
 - Telnet interface for CLI
 - SSH
 - Serial interface for CLI
 - Scriptable CLI
 - Firmware image update (TFTP and FTP)
 - Network Time Protocol (NTP) for switch clock synchronization
 - IBM System Networking Switch Center 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 diagnostic tests
- ▶ Special functions:
 - Serial over LAN (SOL)
- ▶ Network cables:
 - 10BASE-T:
 - UTP Category 3, 4, 5 (100 m (328 ft.) maximum)
 - 100-ohm STP (100 m maximum)
 - 100BASE-TX:
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568 100-ohm STP (100 m maximum)
 - 1000BASE-T:
 - UTP Category 6
 - UTP Category 5e (100 m maximum)
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568B 100-ohm STP (100 m maximum)
- ▶ RS-232 serial cable: 3 m console cable DB-9 to USB connector (nonstandard use of USB connector) includes 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

For more information, see the following documentation:

- ▶ *IBM Layer 2-3 GbE Switch Module Installation Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-58334>
- ▶ *IBM Layer 2-3 GbE Switch Module Application Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-58325>

2.7.7 IBM Layer 2/3 Fiber GbE Ethernet Switch Module

The IBM Layer 2/3 Fiber Gigabit Ethernet Switch Module is a switch option that enables administrators to consolidate full Layer 2/3 LAN switching and routing capabilities into a BladeCenter chassis. Consolidation flattens the topology of the data center infrastructure and reduces the number of discrete devices, management consoles, and different management systems.

Figure 2-31 shows the switch module.



Figure 2-31 IBM Layer 2/3 Fiber GbE Switch Module

SFP modules: Six SFP modules for SX connections are standard with the IBM Layer 2/3 Fiber GbE Switch Module.

Table 2-52 lists the part number to use to order the module.

Table 2-52 IBM Layer 2/3 Fiber GbE Switch Module part number and feature code for ordering

Description	Part number	Feature code
IBM Layer 2/3 Fiber GbE Switch Module for IBM BladeCenter	32R1861	1496

The following features and specifications are supported for the IBM Layer 2/3 Fiber GbE Switch Module:

- ▶ Internal ports:
 - Fourteen 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 transceiver-based LC fiber 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
- ▶ 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 1 - 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:
 - 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 remarking 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
 - IBM System Networking Switch Center 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 diagnostic tests

- ▶ Special functions:
 - Serial over LAN (SOL)
- ▶ Network cables:
 - 1000BASE-SX: 850 nm wavelength, multimode fiber, 50 μ or 62.5 μ (550 m maximum), with LC duplex connector
 - RS-232 serial cable: 3 m console cable DB-9 to USB connector (nonstandard use of USB connector) that includes 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

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

For more information, see the following documentation:

- ▶ *IBM Layer 2/3 GbE Switch Module Installation Guide*
http://publib.boulder.ibm.com/infocenter/bladectr/documentation/topic/com.ibm.bladecenter.io_32R1860.doc/00d9265.pdf
- ▶ *IBM Layer 2/3 GbE Switch Module Application Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ldocid=MIGR-58325>

2.7.8 IBM Layer 2-7 Gigabit Ethernet Switch Module

The IBM Layer 2-7 Gigabit Ethernet Switch Module serves as a switching and routing fabric for the IBM BladeCenter server chassis. It also introduces Layer 4-7 functionality for application-based and server-based load balancing, advanced filtering, content-aware intelligence, embedded security services, and persistence support.



Figure 2-32 IBM Layer 2-7 GbE Switch Module

Table 2-53 lists the part number to use to order the module.

Table 2-53 IBM Layer 2-7 GbE Switch Module part number and feature code for ordering

Description	Part number	Feature code
IBM Layer 2-7 GbE Switch Module for IBM BladeCenter	32R1859	1494

The following features and specifications are supported for the IBM Layer 2-7 GbE Switch Module:

- Ports
 - Four external RJ-45 1000BASE-T connectors for making 10/100/1000 Mbps connections to a backbone, end stations, and servers.
 - Fourteen 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 1 - 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 CLI
 - Firmware image update (TFTP and FTP)
 - Network Time Protocol (NTP) for switch clock synchronization
 - IBM System Networking Switch Center 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 diagnostic tests
- ▶ Special functions:
 - Serial over LAN (SOL).
- ▶ Network cables:
 - 10BASE-T:
 - UTP Category 3, 4, 5 (100 m (328 ft.) maximum)
 - 100-ohm STP (100 m maximum)
 - 100BASE-TX:
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568 100-ohm STP (100 m maximum)
 - 1000BASE-T:
 - UTP Category 6
 - UTP Category 5e (100 m maximum)
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568B 100-ohm STP (100 m maximum)
- ▶ RS-232 serial cable: A 3 m console cable DB-9-to-USB connector (nonstandard use of USB connector) that includes the GbE switch module.

Cabling: 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

For more information, see the following documentation:

- ▶ *IBM Layer 2-7 GbE Switch Module Installation Guide:*

http://publib.boulder.ibm.com/infocenter/bladectr/documentation/topic/com.ibm.bladecenter.io_32R1859.doc/00d9266.pdf

- ▶ *IBM Layer 2-7 GbE Switch Module Application Guide:*

<http://ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-53098>

2.7.9 IBM 1/10 Gb Uplink Ethernet Switch Module

The IBM 1/10 Gb Uplink Ethernet Switch Module for IBM BladeCenter 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 reduces the number of discrete devices, management consoles, and management systems. In addition, the next-generation switch module hardware supports IPv6 Layer 3 frame forwarding protocols.

This module delivers port flexibility, efficient traffic management, increased uplink bandwidth, and strong Ethernet switching price/performance.

Figure 2-33 shows this switch module.



Figure 2-33 IBM 1/10 Gb Uplink Ethernet Switch Module

Table 2-54 lists the part number to use to order the module.

Table 2-54 IBM 1/10 Gb Uplink Ethernet Switch Module part number and feature code for ordering

Description	Part number	Feature code
IBM 1/10 Gb Uplink Ethernet Switch Module	44W4404	1590

To communicate outside of the chassis, you must have either SFP+ transceivers or SFP+ direct-attach cables (DAC) connected. Direct-attach cables have SFP+ transceivers on both ends. You have the flexibility to expand your bandwidth as needed, using anywhere from one connection up to three connections per switch.

Table 2-55 lists the part numbers to use to order the SFP+ transceivers, FC cables, and DAC cables from IBM.

Table 2-55 IBM part numbers for ordering SFP+ transceivers, FC cables, and DAC cables

Description	Part number	Feature code
10 Gb SFP+		
IBM 10GBase-SR 10 GbE 850 nm Fiber SFP+ Transceiver	44W4408	4942
IBM SFP+ Transceiver	46C3447	5053
DAC cables		
0.5 m Molex Direct Attach Copper SFP+ Cable	59Y1932	3735
1 m Molex Direct Attach Copper SFP+ Cable	59Y1936	3736
3 m Molex Direct Attach Copper SFP+ Cable	59Y1940	3737
7 m Molex Direct Attach Copper SFP+ Cable	59Y1944	3738
FC cables		
3 m Intel Connects Optical Cable	46D0153	3852
10 m Intel Connects Optical Cable	46D0156	3853
30 m Intel Connects Optical Cable	46D0159	3854

The following features and specifications are supported for the IBM 1/10 Gb Uplink Ethernet Switch Module:

- ▶ Internal ports:
 - Fourteen 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 or DAC cables. SFP+ modules and DAC cables are optional and must be ordered from IBM (Table 2-55 for part numbers/feature codes).
 - 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.
- Switch stacking (up to eight switches in one stack).
- Support for Active Multipath Protocol (AMP), as an Access Switch allows this switch to connect to two Aggregator Switches, forming a loop topology without port blocking for load balancing purposes.
- ▶ 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, such as 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. The VLAN numbers range is 1 - 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:
 - Support for IEEE 802.1p, IP ToS/DSCP, and ACL-based (MAC/IP source and destination addresses, and VLANs) traffic classification and processing.
 - Traffic shaping and remarking 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.
- IPv6 host management.
- IPv6 forwarding based on static routes and OSPF V3.
- ▶ Manageability:
 - SNMP (V1 and V3)
 - HTTP browser graphical user interface (GUI)
 - Telnet interface for CLI
 - SSH
 - Serial interface for CLI
 - Scriptable 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 diagnostic tests
- ▶ Special function:
 - Serial over LAN (SOL)
- ▶ Virtualization feature:
 - IBM VMready®
- ▶ 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 m (328 ft.) maximum)
 - 100-ohm STP (100 m maximum)
 - 100BASE-TX:
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568 100-ohm STP (100 m maximum)
 - 1000BASE-T:
 - UTP Category 6
 - UTP Category 5e (100 m maximum)
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568B 100-ohm STP (100 m maximum)
- ▶ RS-232 serial cable: 3 m console cable DB-9-to-USB connector (nonstandard use of USB connector) that includes the GbE switch module.

Switch features: When stacking mode is enabled on a switch, the following features are not available:

- ▶ Active Multi-Path Protocol
- ▶ Protocol-based VLANs
- ▶ Routing protocols
- ▶ Virtual Router Redundancy Protocol
- ▶ IP v6

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

For more information, see the following documentation:

- ▶ *IBM 1/10 Gb Uplink Ethernet Switch Module Installation Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5076217>
- ▶ *IBM 1/10 Gb Uplink Ethernet Switch Module Application Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5076214>
- ▶ *IBM 1/10 Gb Uplink Ethernet Switch Module Command Reference:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5076525>

2.7.10 IBM Virtual Fabric 10 Gb Switch Module

The IBM Virtual Fabric 10 Gb Switch Module for IBM BladeCenter offers the most bandwidth of any blade switch and represents the perfect migration platform for clients who are still at 1 Gb outside the chassis by seamlessly integrating into the existing 1 Gb infrastructure. This switch is the first 10 Gb switch for IBM BladeCenter that supports converged networking (that is, it is able to transmit Converged Enhanced Ethernet (CEE) to a Fibre Channel over Ethernet (FCoE) capable top-of-rack switch). This new feature is available with firmware release 6.1.

In addition, this switch is a key part of the IBM Virtual Fabric offering, which allows clients to form eight virtual network interface controllers (NICs) from one physical NIC and to manage them in virtual groups. This switch can be managed through a CLI or a graphical interface of the switch and in the future with BladeCenter Open Fabric Manager, providing all the benefits of I/O Virtualization at 10 Gb speeds.

If you have a chassis with multiple servers, some operating at 1 Gb or at 10 Gb, and some transmitting converged packets, this single switch can handle all these workloads and connect to a 1 Gb or 10 Gb infrastructure or both.

With the extreme flexibility of the IBM switch, clients can take advantage of the technologies they require for multiple environments. For 1 Gb uplinks, they can take advantage of SFP transceivers. For 10 Gb uplinks, they have a choice of either SFP+ transceivers (SR or LR) for longer distances or more cost-effective and lower-power-consuming options, such as SFP+ direct-attached cables (DAC). DACs can be 1 - 7 m in length and are ideal for connecting chassis together, connecting to a top-of-rack switch, or even connecting to an adjacent rack.

Tip: Previously, this switch was known as the IBM 10-port 10 Gb Ethernet Switch Module. In November 2009, when firmware Version 6.1 was released, the name of the switch was changed to the IBM Virtual Fabric 10 Gb Switch Module.

Figure 2-34 shows the switch module.



Figure 2-34 IBM Virtual Fabric 10 Gb Switch Module

Table 2-56 lists the part number to use to order the module.

Table 2-56 IBM Virtual Fabric 10 Gb Switch Module part number and feature code for ordering

Description	Part number	Feature code
IBM Virtual Fabric 10 Gb Switch Module for IBM BladeCenter	46C7191	1639

The part numbers include the following items:

- ▶ One IBM Virtual Fabric 10 Gb Switch Module
- ▶ One 3 m mini-USB-to-DB9 serial console cable
- ▶ One filler module
- ▶ *IBM Virtual Fabric 10 Gb Switch Module Installation Guide*
- ▶ IBM user license agreement
- ▶ Important Notices document
- ▶ Documentation CD-ROM

Important: SFP+ (small form-factor pluggable plus) transceivers are not included and must be purchased separately.

To communicate outside of the chassis, you must have either SFP+ transceivers or SFP+ direct-attach cables (DAC) connected. Direct-attach cables have SFP+ transceivers on both ends. You have the flexibility to expand your bandwidth as needed, using anywhere from one connection up to 10 connections per switch.

Table 2-57 lists the part numbers to use to order the SFP+ transceivers, FC cables, and DAC cables from IBM.

Table 2-57 IBM part numbers for ordering SFP+ transceivers, FC cables, and DAC cables

Description	Part number	Feature code
10 Gb SFP+		
IBM 10GBase-SR 10 GbE 850 nm Fiber SFP+ Transceiver	44W4408	4942
IBM SFP+ SR Transceiver	46C3447	5053
IBM SFP+ LR Transceiver	90Y9412	A1PM
IBM SFP+ ER Transceiver	90Y9415	A1PP
1 Gb SFP+		
IBM SFP RJ45 Transceiver	81Y1618	3268
IBM SFP SX Transceiver	81Y1622	3269
IBM SFP LX Transceiver	90Y9424	N/A
DAC cables		
0.5 m Molex Direct Attach Copper SFP+ Cable	59Y1932	3735
1 m Molex Direct Attach Copper SFP+ Cable	59Y1936	3736
3 m Molex Direct Attach Copper SFP+ Cable	59Y1940	3737
7 m Molex Direct Attach Copper SFP+ Cable	59Y1944	3738
FC cables		
3 m Intel Connects Optical Cable	46D0153	3852
10 m Intel Connects Optical Cable	46D0156	3853
30 m Intel Connects Optical Cable	46D0159	3854

The following features and specifications are supported for the IBM Virtual Fabric 10 Gb Switch Module:

- ▶ Form-factor:
 - Single-height high-speed switch module.
- ▶ Internal ports:
 - Fourteen internal auto-negotiating ports: 1 Gb or 10 Gb to the server blades
 - Two internal full-duplex 100 Mbps ports connected to the management module
- ▶ External ports:
 - Up to ten 10 Gb SFP+ ports (also designed to support 1 Gb SFP if required, with the flexibility of mixing 1 Gb/10 Gb)
 - One 10/100/1000 Mb copper RJ-45 adapter used for management or data
 - An RS-232 mini-USB connector for serial port that provides an additional means to install software and configure the switch module

- ▶ Scalability and performance:
 - Autosensing 1 Gb/10 Gb internal and external Ethernet ports for bandwidth optimization.
 - Non-blocking architecture with wire-speed forwarding of traffic and full line rate performance of 480 Gbps full duplex.
 - Media access control (MAC) address learning: Automatic update that supports up to 32 KB MAC addresses.
 - Up to 128 IP interfaces per switch.
 - Static, EtherChannel, and LACP (IEEE 802.3ad) link aggregation, up to 100 Gb of total bandwidth per switch, up to 18 trunk groups, and up to eight ports per group.
 - Support for jumbo frames (up to 12288 bytes).
 - Broadcast/multicast storm control.
 - IGMP snooping for limit flooding of IP multicast traffic (IGMP V1, V2, and V3).
 - Configurable traffic distribution schemes over trunk links based on source/destination IP addresses, MAC addresses, or both.
 - Fast port forwarding and fast uplink convergence for rapid STP convergence.
 - Switch stacking (up to eight switches in single stack).
 - Support for Active Multipath Protocol (AMP), as an Access Switch allows this switch to connect to two Aggregator Switches, forming a loop topology without port blocking for load balancing purposes.
- ▶ Availability and redundancy:
 - VRRP for Layer 3 router redundancy.
 - IEEE 802.1D STP for providing Layer 2 redundancy with PVRST+.
 - 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 supported per switch. The VLAN numbers range is 1 - 4095 (4095 is used for the management module's connection only).
 - 802.1Q VLAN tagging support on all ports.
 - Protocol-based 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+, an dLDAP

- ▶ Quality of service:
 - Up to eight queues per port.
 - 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 remarking 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 supported)
 - VRRP for router redundancy
 - Support for up to 128 static routes
 - Routing protocol support (Router Information Protocol (RIP) v1, RIP v2, OSPF v1, v2, and v3, and BGP-4), with up to 1024 entries in the routing table
 - Support for DHCP Relay
 - IPv6 host management
 - IPv6 forwarding based on static routes
- ▶ Manageability:
 - Simple Network Management Protocol (SNMP: V1, V2, and V3)
 - HTTP/HTTPS browser GUI
 - Industry standard CLI and BLADEOS/AlteonOS CLI
 - Telnet interface for CLI
 - SSH v1/v2
 - Serial interface for CLI
 - Scriptable CLI
 - Firmware image update (TFTP and FTP)
 - Network Time Protocol (NTP) for switch clock synchronization
 - IBM System Networking Switch Center 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 the syslog feature.
 - POST diagnostic tests.
- ▶ Special function:
 - Serial over LAN (SOL)
- ▶ Virtualization features:
 - VMready
 - Virtual Fabric Adapter vNIC support:
 - Ethernet, iSCSI, or FCoE traffic is supported on vNICs

- 802.1Qbg Edge Virtual Bridging (EVB) is an emerging IEEE standard for allowing networks to become virtual machine-aware (VM-aware):
 - Virtual Ethernet Bridging (VEB) and Virtual Ethernet Port Aggregator (VEPA) are mechanisms for switching between VMs on the same hypervisor.
 - Edge Control Protocol (ECP) is a transport protocol that operates between two peers over an IEEE 802 LAN providing reliable, in-order delivery of upper layer protocol data units.
 - Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP) allows centralized configuration of network policies that will persist with the VM, independent of its location.
 - EVB Type-Length-Value (TLV) is used to discover and configure VEPA, ECP, and VDP.
- Converged Enhanced Ethernet and FCoE features:
 - Fibre Channel over Ethernet (FCoE) allows Fibre Channel traffic to be transported over Ethernet links.
 - FCoE Initialization Protocol (FIP) snooping enforces point-to-point links for FCoE traffic outside the regular Fibre Channel topology.
 - Priority-Based Flow Control (PFC) (IEEE 802.1Qbb) extends 802.3x standard flow control to allow the switch to pause traffic based on the 802.1p priority value in each packet's VLAN tag.
 - Enhanced Transmission Selection (ETS) (IEEE 802.1Qaz) provides a method for allocating link bandwidth based on the 802.1p priority value in each packet's VLAN tag.
 - Data Center Bridging Capability Exchange Protocol (DCBX) (IEEE 802.1AB) allows neighboring network devices to exchange information about their capabilities.
 - Supports the QLogic Virtual Fabric Extension Module for IBM BladeCenter, which provides FCoE gateway functionality inside the BladeCenter Chassis. See 2.9.5, “QLogic Virtual Fabric Extension Module” on page 165.

Stacking mode: When stacking mode is enabled on a switch, the following features are not available:

- Active Multi-Path Protocol
- Protocol-based VLANs
- Routing protocols
- Virtual Router Redundancy Protocol
- Converged Enhanced Ethernet
- Fibre Channel over Ethernet
- IP v6
- vNIC Virtual Fabric mode

VMready is a unique solution that enables the network to be virtual machine-aware. The network can be configured and managed for virtual ports (v-ports) rather than just for physical ports. With VMready, as VMs migrate across physical hosts, so do their network attributes. Virtual machines can be added, moved, and removed while retaining the same ACLs, QoS, and VLAN attributes. VMready allows for a define-once-use-many configuration that evolves as the server and network topologies evolve. VMready works with all virtualization products, including VMware, Hyper-V, Xen, and KVM, without modification of virtual machine hypervisors or guest operating systems.

VMready is able to automatically discover virtual machines on hypervisor-based hosts connected to internal switch ports (up to 2048 VMs can be recognized) and to pre-provision

network connections and group membership for Virtual Entities. In addition, VMready together with IBM NMotion® allows seamless migration/failover of VMs to different hypervisor hosts, preserving network connectivity configurations.

Virtual Switch Groups (VSGs) allows segmentation of the switch into smaller logical switches that operate independently from each other (up to 32 VSGs are supported). VSGs can be formed by using port groups or VM groups. Port groups may contain both external and internal ports on the switch. VM groups may contain both VMs and external ports on the switch. Internal switch ports cannot be assigned to VM group.

The switch module supports the following IEEE standards:

- ▶ IEEE 802.1D STP with PVRST+
- ▶ IEEE 802.1s MSTP
- ▶ IEEE 802.1w 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 802.3ab 1000BASE-T Gigabit Ethernet
- ▶ IEEE 802.3ae 10GBASE-SR 10 Gb Ethernet fiber optics short range
- ▶ IEEE 802.3ae 10GBASE-LR 10 Gb Ethernet fiber optics long range
- ▶ IEEE 802.3z 1000BASE-SX Gigabit Ethernet

The following network cables are supported for the IBM Virtual Fabric 10 Gb Switch Module:

- ▶ 10GBASE-SR for 10 Gb ports: 850 nm wavelength, multimode fiber, 50 µ or 62.5 µ (300 m maximum), with LC duplex connector
- ▶ 1000BASE-T for RJ-45 port:
 - UTP Category 6 (100 m maximum)
 - UTP Category 5e (100 m maximum)
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568B 100-ohm STP (100 m maximum)

For more information, see the following documents:

- ▶ *IBM Virtual Fabric 10 Gb Switch Module Installation Guide*
- ▶ *Application Guide*
- ▶ *Command Reference*
- ▶ *isCLI Reference*
- ▶ *BBI (Browser-based Interface) Quick Guide*
- ▶ *Release Notes*
- ▶ *End User License Agreement*

These documents can be downloaded from the following address:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5080917>

2.7.11 Brocade Converged 10 GbE Switch Module

The Brocade Converged 10 GbE Switch Module and Brocade 2-Port 10 Gb Converged Network Adapter are part of a leading Converged Ethernet solution for IBM BladeCenter that offers Fibre Channel investment protection, maximum bandwidth and performance, and simplicity in a converged environment. This module is the industry's first truly converged FCoE blade-based solution with native Fibre Channel and Ethernet ports built into one switch

module. The Converged Ethernet switching provides up-front reduction in SAN/LAN cost and complexity and increased rack utilization. It requires less cabling, cooling, and power, which means up-front reduction in network investment.

The Brocade Converged 10 GbE Switch Module also features Dynamic Ports on Demand (DPoD) capability through the Port Upgrade Key. This capability allows you to enable any combination of Fibre Channel and Ethernet ports based on your infrastructure requirements, and to experience Converged Ethernet benefits without significant investment.

With the base model (Converged 10 GbE Switch Module), you can enable 16 of the 30 ports on the switch (8 x 10 Gb CEE faceplate ports, 8 x 8 Gb FC faceplate ports, and 14 x 10 Gb CEE backplane ports). If you purchase the Port Upgrade Key, you can enable all 30 ports on the Switch Module for a fully realized Converged Ethernet solution.

Figure 2-35 shows the switch module.



Figure 2-35 Brocade Converged 10 GbE Switch Module

Table 2-58 lists the part numbers to use to order the module.

Table 2-58 Brocade Converged 10 GbE Switch Module and related options part numbers and feature codes for ordering

Description	Part number	Feature code
Brocade Converged 10 GbE Switch Module for IBM BladeCenter	69Y1909	7656
Brocade 10 Gb SFP+ SR Optical Transceiver (two standard)	49Y4216	0069
Brocade 8 Gb SFP+ SW Optical Transceiver	44X1962	5084
Brocade 14-port Converged 10 GbE Switch Port Upgrade for IBM BladeCenter	44E5686 ^a	
Brocade ISL Trunking	26K5607	4247
Brocade Fabric Watch	45W0504	4398
Brocade Advanced Performance Monitoring	45W0508	4248

a. Replaces 69Y1917.

The module part numbers include the following items:

- ▶ One Brocade Converged 10 GbE Switch Module for IBM BladeCenter (includes two 10 GbE SFP+ modules standard)
- ▶ *IBM Important Notices* document
- ▶ *IBM Warranty and Support Guide*
- ▶ One mini-USB console cable with serial connectors
- ▶ BCHT interposer gasket kit (used if the switch module is inserted in the BCHT chassis)

- DB9 to RJ45 adapter
- The IBM Documentation CD, which contains the following documents:
 - *Brocade Converged 10GbE Switch Module Installation and Users Guide*
 - *IBM Safety Information* document (multilingual)
 - *IBM Environmental Notices and User's Guide*
 - Brocade EULA
 - Brocade CEE Admin guide
 - Brocade CEE CLI guide
 - Brocade MIBS manual

Important: The switch includes, as standard, two small form-factor pluggable plus (SFP+) modules. Additional SFP+ transceivers can be ordered separately if required. See Table 2-58 for a list of supported transceivers and their ordering part numbers.

The base model of the Brocade Converged 10 GbE Switch Module, 69Y1909, includes 16 licensed ports. This module supports the DPoD feature, where ports are licensed as they come online independent of the port type. In the base model, two external CEE ports (ports 15 and 16) are pre-licensed and are preinstalled with two 10 GbE optical transceivers. The remaining 14 licenses are available to be assigned to other internal or external ports. The 16 licensed ports can be any combination of external CEE ports, external FC ports, or internal CEE ports. After all the licenses are assigned, you can manually move licenses from one port to another.

Additional ports can be activated by purchasing the Brocade Converged 10 GbE Switch Port Upgrade, part number 44E5686. This upgrade option activates the license for the remaining 14 ports. All internal and external ports are then licensed and can be used.

Important: The Brocade Converged 10 GbE Switch Port Upgrade includes, as standard, two Brocade 8 Gb SFP+ SW Optical Transceivers.

The Brocade Converged 10 GbE Switch Module with a supported converged network adapter (CNA) provides these benefits:

- Highly integrated BladeCenter switch module

The Brocade Converged 10 GbE Switch Module offers one of the industry's best integrated I/O solutions. The compact design incorporates Ethernet and Fibre Channel switching and provides a total of 30 ports: eight external 10 Gb Ethernet CEE ports for LAN connectivity and eight external 8 Gb Fibre Channel ports for storage and SAN connectivity. With the high integration of this module, you can achieve all of your networking and storage I/O needs with a single module.
- Unique flexible ports and server deployment

The DPoD feature of the Brocade Converged 10 GbE Switch Module allows you to activate any combination of internal or external ports. You can balance server and I/O port assignments to accommodate the workloads of applications and satisfy business needs. DPoD gives you the ability to enable only 10 GbE ports for networking needs, only Fibre Channel ports for storage support, or any combination of Ethernet and Fibre Channel ports.
- Support for DCFM

IBM System Storage® Data Center Fabric Manager (DCFM) is available for use with the switch. DCFM is fully integrated with IBM Systems Director for end-to-end data center management. A 75-day trial version can be downloaded from the following address:

<http://www.ibm.com/systems/storage/san/b-type/dcfm/>

- ▶ Highest bandwidth and low latency

The total available bandwidth of 144 Gbps (80 Gbps Ethernet and 65 Gbps Fibre Channel) accessible through external ports means that the switch module can support high-demand applications and storage solutions. In addition to the high bandwidth available on external ports, the switch module supports an additional 140 Gbps on 14 internal 10 GbE ports. The switch module uses cut-through and non-blocking to deliver high performance and low latency for demanding virtual applications and high-speed environments.

- ▶ Low total cost of ownership

The highly integrated design delivers the lowest total cost of ownership (TCO). The Brocade Converged 10 GbE Switch Module is two switches in one (Ethernet and Fibre Channel), which reduces cost of acquisition. The switch module is built around a pay-as-you-go model that gives you a lower price entry point and allows you to add ports as your business needs dictate. The switch module includes two 10 GbE SFP+ transceivers to further lower initial investments. As it is designed for seamless integration into existing networking and storage environments, this module truly protects existing investments. Overall, with this integrated switch module, your server infrastructure uses less hardware components with lower costs and higher reliability, which further drives a lower TCO.

The converged switch modules have the following features:

- ▶ Eight external 10 Gb Converged Enhanced Ethernet ports
- ▶ Eight external auto-negotiated Fibre Channel ports (2 Gbps, 4 Gbps, or 8 Gbps)
- ▶ Fourteen internal auto-negotiated 10 GbE ports (1 Gbps or 10 Gbps)
- ▶ Two internal full-duplex 100 Mbps Ethernet interfaces for management purposes
- ▶ One external RS232 console port with a mini-USB interface for serial console management
- ▶ One external 10/100/1000 Mb RJ45 Ethernet copper port for debugging and field support
- ▶ CEE features:
 - Priority-based Flow Control (PFC): IEEE 802.1Qbb
 - Enhanced Transmission Selection (ETS): IEEE 802.1Qaz
 - Data Center Bridging eXchange (DCBX)
- ▶ Layer 2 features:
 - Layer 2 Virtual Local Area Networks (VLANs): 4096
 - VLAN Encapsulation 802.1Q
 - Rapid Spanning Tree Protocol (RSTP)
 - Multiple Spanning Tree MSTP (802.1s): 16 instances
 - Link Aggregation Control Protocol (LACP) IEEE- 802.3ad
 - Brocade enhanced frame-based trunking
 - Advanced PortChannel hashing based on Layer 2, 3, and 4 information
 - Pause Frames (802.3x)
 - Storm Control (unicast, multicast, and broadcast)
 - Address Resolution Protocol (ARP) RFC 826
- ▶ Layer 2 security:
 - Ingress Access Control Lists (ACLs)
 - Standard and extended Layer 2 ACLs
 - VLAN-based ACLs (VACLs)
 - Port-based ACLs (PACLs)
 - ACL statistics
 - Port-based Network Access Control: IEEE 802.1X

- ▶ Layer 2 Quality of service (QoS):
 - Eight priority levels for QoS
 - IEEE 802.1p Class of Service (CoS)
 - Eight hardware queues per port
 - Per-port QoS configuration
 - CoS trust: IEEE 802.1p
 - Per-port Virtual Output Queuing
 - CoS-based egress queuing
 - Egress strict priority queuing
 - Egress port-based scheduling: Weighted Round-Robin (WRR)
- ▶ Layer 3 feature:
 - Static IP routes
- ▶ Fibre Channel features:
 - External Fibre Channel ports that can operate as F_ports (fabric ports), FL_ports (fabric loop ports), or E_ports (expansion ports)
 - FC fabric services
 - Simple Name Server (SNS)
 - Registered State Change Notification (RSCN)
 - Dynamic Path Selection (DPS)
 - Enhanced Group Management (EGM)
 - ISL Trunking (optional)
 - Fabric Watch (optional)
 - Advanced Performance Monitor (optional)
- ▶ Power-on self-test diagnostics and status reporting

The following software features are included with the switch modules:

- ▶ Brocade Web Tools
- ▶ 16-port licensing

The switch supports the following fabric management (all management connections go through the management module except direct serial connection, which goes through the mini-USB port):

- ▶ Web interface through Web Tools
- ▶ Command-line interface (CLI) through the Telnet program
- ▶ A terminal emulation program connection to the mini-USB port interface
- ▶ IBM System Storage Data Center Fabric Manager (DCFM) application
- ▶ The switch's SNMP agent

The switch module supports the following IEEE standards:

- ▶ IEEE 802.1D STP
- ▶ IEEE 802.1s MSTP
- ▶ IEEE 802.1w 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 802.3ad Link Aggregation Control Protocol
- ▶ IEEE 802.3x Full-duplex Flow Control
- ▶ IEEE 802.3ab 1000BASE-T Gigabit Ethernet
- ▶ IEEE 802.3ae 10GBASE-SR 10 Gb Ethernet fiber optics short range

The following network cables are supported for the Brocade Converged 10 GbE Switch Module:

- ▶ 10GBASE-SR for 10 Gb ports: 850 nm wavelength, multimode fiber, 50 μ or 62.5 μ (300 m maximum), with LC duplex connector
- ▶ 8 Gb FC ports when operating at 8 Gbps speed
850 nm wavelength, multimode fiber, 50 μ (150 m maximum) or 62.5 μ (21 m maximum), with LC duplex connector
- ▶ 1000BASE-T for RJ-45 port:
 - UTP Category 6 (100 m maximum)
 - UTP Category 5e (100 m maximum)
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568B 100-ohm STP (100 m maximum)

The switch module has the following physical specifications:

- ▶ Dimensions:
 - Approximate height: 4.10 cm (1.61 in.)
 - Approximate width: 29.38 cm (11.57 in.)
 - Approximate depth: 25.70 cm (10.12 in.)
 - Approximate weight: 2.13 kg (4.7 lb)
- ▶ Temperature:
 - Operating: 0° - 40°C (32° - 104°F) at 0 - 3,048 m (0 - 10,000 ft.)
 - Non-operating: -20° - 70°C (-4 - 158°F) at 0 - 10,688 km (0 - 35,000 ft.)
- ▶ Humidity:
 - Operating: 10 - 90% (non-condensing) at 29°C
 - Non-operating: 5 - 95% (non-condensing) at 38°C

For more information, see the following documents:

- ▶ *Brocade Converged 10GbE Switch Module Installation and User's Guide*
- ▶ *Brocade Converged Enhanced Ethernet Administrator's Guide*
- ▶ *Brocade Fabric OS Administrator's Guide*

These documents can be downloaded from the following address:

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

For more information, see the *Brocade Converged 10GbE Switch Module for IBM BladeCenter at-a-glance* guide:

<http://www.redbooks.ibm.com/abstracts/tips0789.html>

2.8 InfiniBand switch module

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)
- ▶ Virtualizes and shares I/O across an entire BladeCenter or collection of BladeCenter chassis for cost savings and high availability

The part number and feature code for ordering the InfiniBand switch module for BladeCenter is listed in Table 2-59.

Table 2-59 *InfiniBand switch module for IBM BladeCenter*

Part Number	Feature Code	Description	Number of external ports	Type of external ports	Page
46M6005	0057	Voltaire 40 Gb InfiniBand Switch Module	16	4X QDR (40 Gbps)	151

2.8.1 Voltaire 40 Gb InfiniBand Switch Module

The Voltaire 40 Gb InfiniBand Switch Module for BladeCenter provides InfiniBand QDR connectivity between the server platform's high performance blade servers and external InfiniBand fabrics in non-blocking designs, all on a single device. This high speed module also accommodates performance-optimized fabric designs using a single BladeCenter chassis or stacking multiple BladeCenter chassis without requiring an external InfiniBand switch.

The InfiniBand switch module offers 14 internal ports, one to each server, and 16 ports out of the chassis per switch. This solution offers a no-compromise, congestion-free solution to meet even the most performance-hungry applications.

The module's HyperScale architecture also provides a unique inter-switch link or mesh capability to form highly scalable, cost-effective, and low latency fabrics. Because this switch has 16 uplink ports, they can create a meshed architecture and still have unblocked access to data using the 14 uplink ports. This solution can scale from 14 to 126 nodes and offer latency of less than 200 ns, allowing applications to operate at maximum efficiency.

The Voltaire 40 Gb InfiniBand Switch Module is shown in Figure 2-36.



Figure 2-36 *Voltaire 40 Gb InfiniBand Switch Module*

Table 2-60 shows the part number to use to order the module and supported cables.

Table 2-60 *Voltaire 40 Gb InfiniBand Switch Module part number and feature code for ordering*

Description	Part number	Feature code
Voltaire 40 Gb InfiniBand Switch Module for IBM BladeCenter	46M6005	0057

The Voltaire 40 Gb InfiniBand Switch Module part number includes the following items:

- ▶ One Voltaire 40 Gb InfiniBand Switch Module
- ▶ Documentation CD-ROM
- ▶ *IBM Important Notices* flyer

The Voltaire 40 Gb InfiniBand Switch Module includes the following features and functions:

- ▶ Form-factor
 - Double-height high-speed switch module
- ▶ Internal ports:
 - Fourteen internal 4X QDR/DDR/SDR InfiniBand ports to the server blades
 - Two internal 4X InfiniBand ports to the bridge modules
- ▶ External ports
 - Up to 16 auto-sensing 4X QDR/DDR/SDR InfiniBand QSFP ports (40, 20, or 10 Gbps auto-negotiate)
- ▶ Characteristics:
 - Full QDR rate InfiniBand switching
 - Based on the Infiniscale-IV device
 - Up to 40 Gbps performance for clusters and grids (bidirectional)
 - Hot-swap support, with soft start and current limiting
 - Bay address and presence support
 - High temperature monitoring
 - Full BladeCenter H AMM interface through I2C
 - Fully non-blocking QDR blade
 - QSFP front panel connectors
 - Onboard processor for chassis management capabilities
 - Ultra-low latency: Under 100 ns blade to blade connectivity
 - Built-in high availability
 - Port status controllability

Table 2-61 lists Voltaire 40 Gb InfiniBand Switch module-related option.

Table 2-61 Voltaire 40 Gb InfiniBand Switch Module options

Description	Part number	Feature code
3 m Copper QDR InfiniBand QSFP Cable	49Y9980	3866

For more information, see *Voltaire 40 Gb InfiniBand Switch Module for IBM BladeCenter User Guide*

<http://www.voltaire.com/SupportUserManuals/BladeSolutionsUM>

2.9 Fibre Channel switch modules

This section describes SAN storage options available for IBM BladeCenter. For information related to storage infrastructure planning and integration, see Chapter 5, “Storage integration” on page 449.

Firmware releases: 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.

See 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 2-62 lists the current Fibre Channel switch modules.

Table 2-62 Fibre Channel switch modules

Part number	Feature Code	Description	Number of external ports	Port interface bandwidth	Page
42C1828	5764	Brocade Enterprise 20-port 8 Gb SAN Switch Module	6	8 Gbps	154
44X1920	5481	Brocade 20-port 8 Gb SAN Switch Module	6	8 Gbps	154
44X1921	5483	Brocade 10-port 8 Gb SAN Switch Module ^a	6	8 Gbps	154
44E5696 ^b	A3FH	Cisco Systems 20-port 4 Gb FC Switch Module	6	4 Gbps	158
44E5692 ^c	A3FG	Cisco Systems 10-port 4 Gb FC Switch Module ^a	6	4 Gbps	158
44X1905	5478	QLogic 20-Port 8 Gb SAN Switch Module	6	8 Gbps	160
88Y6406	A24C	QLogic 20-Port 4/8 Gb SAN Switch Module	6	8 Gbps	160
44X1907	5482	QLogic 8 Gb Intelligent Pass-thru Module	6	8 Gbps	163
88Y6410	A24D	QLogic 4/8 Gb Intelligent Pass-thru Module	6	8 Gbps	163
46M6172	4799	QLogic Virtual Fabric Extension Module	6	8 Gbps	165

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

b. Replaces 39Y9280.

c. Replaces 39Y9284.

2.9.1 Brocade Enterprise 20-port, 20-port, and 10-port 8 Gb SAN Switch Modules

The Brocade Enterprise 20-port, 20-port, and 10-port 8 Gb SAN Switch Modules for IBM BladeCenter deliver embedded Fibre Channel switching with the industry's fastest Fibre Channel throughput performance (8 Gbps). The 20-port and 10-port switches provide basic FC switch functionality, and the Enterprise 20-Port model is a fully enabled switch with a complete set of advanced SAN fabric service licensed features. These switches are also known as Brocade 5470 switches, and they are functionally equivalent to Brocade 300 switches unless otherwise stated in the documentation.

Figure 2-37 shows the Brocade 8 Gb SAN switches.



Figure 2-37 Brocade 10-port and 20-port SAN Switch Modules for IBM BladeCenter

Brocade 8 Gb SAN switches operate in full-fabric switch mode for native interoperability with IBM System Storage b-type SAN switches, and in Brocade Access Gateway mode for industry-standard interoperability with any fabric that supports N_Port ID Virtualization (NPIV) technology.

The Brocade 8 Gb SAN Switch Modules fit in the standard I/O bays 3 and 4 of the BladeCenter chassis. With the addition of the Multi-Switch Interconnect Module (MSIM) in the BladeCenter H, they can also fit in the high-speed I/O bays 8 and 10. The Brocade 8 Gb SAN Switch Modules are not supported with MSIM-HT in high-speed bays of the BladeCenter HT chassis.

Table 2-63 lists the functional differences between these switches.

Table 2-63 Brocade Enterprise 20-port, 20-port, and 10-port switch modules feature comparison

Feature	Enterprise 20-port switch	20-port switch	10-port switch
Number of active ports	20	20	10
10-port Upgrade License	Not applicable	Not applicable	Optional
Full Fabric	Included	Included	Included
Access Gateway	Included	Included	Included
Enhanced Group Management	Included	Included	Included
Advanced Zoning	Included	Included	Included
Fabric Tools	Included	Included	Included
ISL Trunking	Included	Optional	Optional
Adaptive Networking	Included	Optional	Optional
Advanced Performance Monitoring	Included	Optional	Optional
Fabric Watch	Included	Optional	Optional
Extended Fabrics	Included	Optional	Optional
Server Application Optimization	Included	Not available	Not available

Table 2-64 shows Brocade 8 Gb SAN Switch Modules ordering information.

Table 2-64 Brocade 10-port and 20-port 8 Gb FC Switch Modules part numbers for ordering

Description	Part number	Feature code
Brocade Enterprise 20-port 8 Gb SAN Switch Module	42C1828	5764
Brocade 20-port 8 Gb SAN Switch Module	44X1920	5481
Brocade 10-port 8 Gb SAN Switch Module	44X1921	5483
Brocade 10-port Upgrade ^a	32R1822	4397
Brocade ISL Trunking	26K5607	4247
Advanced Performance Monitoring	45W0508	4248
Extended Fabrics	45W0505	4249
Fabric Watch	45W0504	4398
Adaptive Networking	45W0510	None

a. For 10-port model only

The switch module part numbers include the following items:

- ▶ Brocade Enterprise 20-port, 20-port, or 10-port 8 Gb SAN Switch Module
- ▶ Gasket for installation in a BladeCenter H Multi-Switch Interconnect Module (MSIM)
- ▶ Documentation CD containing:
 - *Brocade 20-port and 10-port 8Gb SAN Switch Module Installation and User Guide*
 - End User License Agreement (EULA)
 - Safety Information
 - Brocade documentation
- ▶ The *IBM Important Notices* document

The switch does not include small form-factor pluggable plus (SFP+) modules; they must be ordered separately. Table 2-65 lists the transceiver that is supported.

Table 2-65 Supported SFP+ part numbers for ordering

Description	Part number	Feature code
Brocade 8 Gb SFP+ SW Optical Transceiver	44X1962	5084

The Brocade switch modules include the following hardware features:

- ▶ Six external auto-sensing Fibre Channel ports (2 Gbps, 4 Gbps, or 8 Gbps for 8 Gb SFP+ transceivers) that operate at a maximum speed of 8 Gbps.
- ▶ Fourteen internal Fibre Channel ports that operate as F_ports (fabric ports) at 2 Gbps, 4 Gbps, or 8 Gbps, depending on FC expansion cards installed into blades.
- ▶ A 10-port model includes 10 Dynamic Ports-On-Demand (DPoDs) that can be activated in any combination of internal and external ports.
- ▶ Two internal full-duplex 100 Mb Ethernet ports for management.
- ▶ High availability features, including hot-pluggable design, nondisruptive software upgrades, automatic path failover, and power-on self-test diagnostic tests and status reporting.
- ▶ 8 Gb switch fabric aggregated I/O bandwidth: 320 Gbps (20 ports at 8 Gbps line rate at full-duplex).
- ▶ External ports that can operate as F_ports (fabric ports), FL_ports (fabric loop ports), or E_ports (expansion ports) when in full FC switch mode. When in Brocade Access Gateway mode, the external ports operate as NPIV-enabled N_Ports.
- ▶ Fabric latency of 0.74 μ s ensures the fastest response and access time to data for critical applications and servers.
- ▶ Supports full 8 Gbps operations at distances up to 129 km and up to 1036 km at 1 Gbps.
- ▶ Support up to 239 switches, depending on the configuration.
- ▶ Frame-filtering technology that enables advanced zoning and monitoring capabilities.
- ▶ Optional small form-factor pluggable plus (SFP+) modules.

Brocade switch modules come standard with the following software features:

- ▶ Brocade Full Fabric mode: Enables high performance 8 Gb fabric switching.
- ▶ Brocade Access Gateway mode: Uses NPIV to connect to any fabric without adding switch domains to reduce management complexity.
- ▶ Dynamic Path Selection: Enables exchange-based load balancing across multiple Inter-Switch Links (ISLs) for superior performance.

- ▶ Brocade Advanced Zoning: Segments a SAN into virtual private SANs to increase security and availability.
- ▶ Brocade Enhanced Group Management: Enables centralized and simplified management of Brocade fabrics through IBM System Storage Data Center Fabric Manager (DCFM).

Brocade switch modules support the following fabric management features (all management connections go through the management module):

- ▶ IBM System Storage DCFM: Centralizes and simplifies management of enterprise-level SAN fabrics consisting of IBM System Storage b-type (using FOS) and Brocade fabrics m-type (using M-EOS) SAN products.
- ▶ IBM System Storage DCFM Professional Plus: For customers managing single fabrics up to 1000 ports running Brocade Fabric OS (FOS) products.
- ▶ Brocade Web Tools: Provides a web-based management interface designed to help simplify monitoring and configuration management.
- ▶ Command-line interface through the Telnet program.
- ▶ Switch SNMP agent, which enables receipt of configuration values, traffic information, and FC failure data through SNMP and the Ethernet interface.

The Enterprise 20-port model (42C1828) is a fully enabled switch with a complete set of licensed features that maximizes performance, ensures availability, and simplifies management for the most demanding applications and expanding virtualization environments. The factory-installed licenses include:

- ▶ ISL Trunking: Consolidates ISLs into fault-tolerant and load-balanced trunks with bandwidth of up to 48 Gbps for greater performance and simplified management.
- ▶ Adaptive Networking: Ensures high-priority connections that obtain the bandwidth necessary for optimum performance.
- ▶ Advanced Performance Monitoring: Provides end-to-end visibility into fabric performance to maximize performance tuning and the usage of fabric resources.
- ▶ Fabric Watch: Monitors and creates alerts based on the health of switch and fabric elements.
- ▶ Extended Fabrics: Extends Fibre Channel SANs beyond traditional distance limitations for replication and backup at full bandwidth.
- ▶ Server Application Optimization: Allows virtual machine quality of service, authentication, and monitoring from adapter through SAN.

All of these features, except for Server Application Optimization, are available as options for the 10-port and 20-port models of the switch. Table 2-64 on page 155 lists the ordering information for each of them. These features are factory installed on the Enterprise 20-port module (42C1828).

For additional information, see the *Brocade 8 Gb 10-port and 20-port SAN Switch Module Installation and User's Guide*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5082884>

2.9.2 Cisco Systems 4 Gb 20-port and 10-port Fibre Channel Switch Modules

Cisco Systems 4 Gb Fibre Channel Switch Modules provide high-performance end-to-end SAN solutions using 4 Gb Fibre Channel technology. These modules allow seamless integration of IBM BladeCenter servers into existing Cisco SANs, and have functions equivalent to the Cisco MDS 9124 switch. These switch modules are shown in Figure 2-38.



Figure 2-38 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 10 activated ports, and the 20-port switch has 20 activated ports.

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 7 of 14 internal ports can be enabled), or external devices (the first three ports can be enabled consequently). Up to 10 ports can be active at the same time. It is possible to make ports ineligible to obtain the license, or move licenses among ports.

Table 2-66 shows the ordering information for the Cisco Systems 4 Gb Fibre Channel Switch Modules.

Table 2-66 Cisco Systems 4 Gb Fibre Channel switch modules' part numbers for ordering

Description	Part number	Feature code
Cisco Systems 4 Gb 20-port Fibre Channel Switch Module	44E5696 ^a	A3FH
Cisco Systems 4 Gb 10-port Fibre Channel Switch Module	44E5692 ^b	A3FG

a. Replaces 39Y9280.

b. Replaces 39Y9284.

These modules have the following features and functions:

- ▶ Six external autosensing Fibre Channel ports that operate at 4, 2, or 1 Gbps
- ▶ Fourteen internal autosensing Fibre Channel ports that operate at 2 or 4 Gbps
- ▶ Two internal full-duplex 100 Mbps Ethernet interfaces
- ▶ Power-on diagnostics and status reporting
- ▶ External ports that 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 or 4 Gbps
- ▶ N-port ID Virtualization (NPIV) support
- ▶ VSAN support (up to 16 VSANs per switch)
- ▶ Advanced Traffic Management capabilities:
 - 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 are included with the switch modules:

- ▶ Cisco Fabric Manager
- ▶ Twenty-port licensing (20-port version only)
- ▶ Ten-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) methods:

- ▶ Web interface through Cisco Fabric Manager
- ▶ Cisco MDS9000 family CLI through the Telnet program
- ▶ Application programming interface (API)

SFP modules: The switch comes without SFP modules, and they must be ordered. Table 2-67 on page 160 lists the available part numbers for ordering that are supported.

Table 2-67 Supported SFP part numbers for ordering

Part number	Description
41Y8598	Cisco Systems 4 Gb Short-wave Length SFP Module

Table 2-68 lists the optional feature available for these switches.

Table 2-68 Optional software features part numbers for ordering

Part number	Description
39Y9290	Cisco Systems 10 port License Upgrade ^a

a. Available for the 10-port switch only.

For more information, see *Cisco Fibre Channel Switch Module setup information - IBM BladeCenter*.

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5070625>

2.9.3 QLogic 20-port 8 Gb and 4/8 Gb SAN Switch Modules

The QLogic 20-port 8 Gb SAN Switch Module and the QLogic 20-port 4/8 Gb SAN Switch Module are high-speed additions to the BladeCenter switch portfolio that help enable scalability in SAN size and complexity while maintaining ease of management. The 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. Figure 2-39 shows the 8 Gb switch. The 4/8 Gb switch has an identical port configuration.



Figure 2-39 QLogic 20-port 8 Gb SAN Switch Module

Table 2-69 shows the ordering information for these switch modules.

Table 2-69 QLogic 20-port 8 Gb and 4/8 Gb SAN Switch Module part numbers for ordering

Description	Part number	Feature code
QLogic 20-port 8 Gb SAN Switch Module for IBM BladeCenter	44X1905	5478
QLogic 20-port 4/8 Gb SAN Switch Module for IBM BladeCenter	88Y6406	A24C

The switch modules have 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.
- ▶ The QLogic 20-port 8 Gb SAN Switch Module has 14 internal Fibre Channel ports that operate as F_ports (fabric ports) at 2, 4, or 8 Gbps, depending on the FC expansion cards installed into the blades.
- ▶ The QLogic 20-port 4/8 Gb SAN Switch Module has 14 internal Fibre Channel ports that operate as F_ports (fabric ports) at 2 Gbps, 4 Gbps, or 8 Gbps, depending on the FC expansion cards installed into blades and the I/O module bays the switch is installed in:
 - When installed in an MSIM (BladeCenter H only), the internal ports operate at up to 8 Gbps.
 - When installed in a BladeCenter T chassis, the internal ports operate at 2 Gbps.
 - In all other cases, the internal ports operate at up to 4 Gbps.
- ▶ 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 for up to 239 switches, depending on the 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 QuickTools web interface software feature includes the switch module.

The switch supports the following management methods (all management connections go through the management module):

- ▶ Web interface through QuickTools.
- ▶ Command-line interface through the 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.

Important: The switch does not include SFP+ modules, and they must be ordered separately. Table 2-70 lists the part numbers for ordering that are supported.

Table 2-70 Supported SFPs and SFPs+ part numbers for ordering

Part number	Feature code	Description
44X1964	5075	IBM 8 Gb SFP+ SW Optical Transceiver
22R4897	2414	4 Gbps SW SFP Transceiver 4 pack
22R4902	2410	4 Gbps SW SFP Transceiver

For more information, see the following documentation:

- ▶ *The QLogic 8 Gb Intelligent Pass-thru Module and 20-port 8 Gb SAN Switch Module Installation and User's Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5079302>
- ▶ *The QLogic 4/8 Gb Intelligent Pass-thru Module and 20-port 4/8 Gb SAN Switch Module Installation and User's Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5087750>

2.9.4 QLogic 8 Gb and 4/8 Gb Intelligent Pass-thru Modules

The QLogic 8 Gb Intelligent Pass-thru Module and QLogic 4/8 Gb Intelligent Pass-thru Module for IBM BladeCenter provide seamless integration of IBM BladeCenter solutions into existing Fibre Channel storage networks using N_Port ID Virtualization (NPIV) technology. Figure 2-40 shows the 8 Gb module. The 4/8 Gb module has an identical port configuration.



Figure 2-40 QLogic 8 Gb Intelligent Pass-thru Module for IBM BladeCenter

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-through 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.

Connection to Fibre Channel switch: This module *must* connect to the external Fibre Channel switch (or switches) that supports NPIV. Direct connections from the module's external ports to storage devices or hosts must be avoided.

Table 2-71 shows the ordering information.

Table 2-71 QLogic 8 Gb and 4/8 Gb Intelligent Pass-thru Module part numbers for ordering

Description	Part number	Feature code
QLogic 8 Gb Intelligent Pass-thru Module for IBM BladeCenter	44X1907	5482
QLogic 4/8 Gb Intelligent Pass-thru Module for IBM BladeCenter	88Y6410	A24D

The I/O 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.
- ▶ The QLogic 8 Gb Intelligent Pass-thru Module has 14 internal Fibre Channel ports that operate as TH_Ports (Transparent Host Ports) at 2, 4, or 8 Gbps, depending on the FC expansion cards installed into the blades.
- ▶ The QLogic 4/8 Gb Intelligent Pass-thru Module has 14 internal Fibre Channel ports that operate as F_ports (fabric ports) at 2 Gbps, 4 Gbps, or 8 Gbps, depending on the FC expansion cards installed into blades and the I/O module bays in which the switch is installed:
 - When installed in an MSIM (BladeCenter H only), the internal ports operate up to 8 Gbps.
 - When installed in a BladeCenter T chassis, the internal ports operate at 2 Gbps.
 - In all other cases, the internal ports operate at up to 4 Gbps.
- ▶ 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 for up to 239 switches, depending on the 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 QuickTools web interface software feature includes the switch module.

The module supports the following management methods (all management connections go through the management module):

- Web interface through QuickTools.
- Command-line interface 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.

Upgrading the modules: The QLogic 8 Gb and 4/8 Gb Intelligent Pass-thru Modules cannot be upgraded to a full-fabric Fibre Channel switch module.

The switch comes without SFP+ modules, which must be ordered separately. Table 2-72 lists the part numbers for ordering that are supported.

Table 2-72 Supported SFPs and SFPs+ part numbers for ordering

Part number	Feature code	Description
44X1964	5075	IBM 8 Gb SFP+ SW Optical Transceiver
22R4897	2414	4 Gbps SW SFP Transceiver 4-Pack
22R4902	2410	4 Gbps SW SFP Transceiver

For more information, see the following resources:

- *The QLogic 8 Gb Intelligent Pass-thru Module and 20-port 8 Gb SAN Switch Module Installation and User's Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5079302>
- *The QLogic 4/8 Gb Intelligent Pass-thru Module and 20-port 4/8 Gb SAN Switch Module Installation and User's Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5087750>

2.9.5 QLogic Virtual Fabric Extension Module

The QLogic Virtual Fabric Extension Module is another example of how IBM is at the forefront of offering new technology to clients. IBM was among the first to deliver Fibre Channel over Ethernet across System x and BladeCenter and this new module takes this process a step further by offering clients I/O convergence inside the chassis.

Clients using the IBM Virtual Fabric 10 Gb Switch Module for their LAN traffic can now combine it with this module and a Converged Network Adapter, such as the QLogic 2-port 10 Gb CNA, to converge their LAN and SAN on a single network. The QLogic Virtual Fabric Extension Module offers six ports of 8 Gb Fibre Channel connectivity, without the need for separate Fibre Channel expansion cards in the BladeCenter servers.

Figure 2-41 shows the QLogic Virtual Fabric Extension Module.



Figure 2-41 QLogic Virtual Fabric Extension Module for IBM BladeCenter

This solution not only offers a simplified fabric, but also significant savings over a separate 10 Gb/8 Gb infrastructure. In addition to the cost and management savings, this solution provides the flexibility to upgrade an existing chassis to implement I/O convergence without replacing existing 10 Gb switches. Additionally, this module can operate either in full-fabric mode or transparent (NPIV) mode, depending on your FC storage infrastructure requirements.

The QLogic Virtual Fabric Extension Module acts as a bridge between native Fibre Channel SANs and 10 Gb Converged Enhanced Ethernet networks. This module is installed in bays 3 and 5 of a BladeCenter H chassis. Installing two bridges, one in each bay, provides redundancy and performance benefits, although you can install one QLogic Virtual Fabric Extension Module in bay 5 only. The QLogic Virtual Fabric Extension Module requires one or two IBM Virtual Fabric 10 Gb Switch Modules, 46C7191. Depending on the configuration used (as described in 5.11, “Converged networks” on page 484), two or four of the 10 external ports of the IBM switch module are rerouted internally to the extension module.

Table 2-73 shows QLogic Virtual Fabric Extension Module ordering information.

Table 2-73 QLogic Virtual Fabric Extension Module part number for ordering

Description	Part number	Feature code
QLogic Virtual Fabric Extension Module for IBM BladeCenter	46M6172	4799

The part number includes the following items:

- ▶ One QLogic Virtual Fabric Extension Module for IBM BladeCenter
- ▶ Support CD
- ▶ The *IBM Important Notices* document
- ▶ Warranty information

The QLogic Virtual Fabric Extension Module comes without SFP+ modules. They must be ordered additionally. Table 2-74 lists the part number that is supported.

Table 2-74 Supported SFP+ for the QLogic Virtual Fabric Extension Module

Description	Part number	Feature code
IBM 8 Gb SFP+ SW Optical Transceiver	44X1964	5075

The QLogic Virtual Fabric Extension Module for IBM BladeCenter has the following features:

- ▶ Standard I/O module form factor.
- ▶ Six external autosensing Fibre Channel ports that operate at 8 Gbps, 4 Gbps, or 2 Gbps.
- ▶ External ports can be configured as full fabric (GL, G, F, FL, E) Fibre Channel ports, or transparent fabric (TF) Fibre Channel ports.
- ▶ Up to 40 Gbps of internal bandwidth to the switch module (eight internal 10 Gb ports) can connect to high-speed switch modules. Up to four ports can be active at the same time, depending on the configuration. See 5.11, “Converged networks” on page 484.
- ▶ Two internal full-duplex 100 Mbps Ethernet interfaces for management.
- ▶ 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.
- ▶ Support for transparent mode (NPIV).
- ▶ 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 for up to 239 Fabric Extension Modules, depending on the configuration.
- ▶ 8 Gb switch fabric aggregate bandwidth: 224 Gbps at full duplex.
- ▶ Maximum frame size: 2148 bytes (2112-byte payload).
- ▶ Nonblocking architecture to prevent latency.
- ▶ Support for the 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).
- ▶ Support for optional SFP+ modules.

The QuickTools web interface software feature includes the switch module:

The switch supports the following fabric management (all management connections go through the management module):

- ▶ Web interface through QuickTools
- ▶ Command-line interface through the Telnet program
- ▶ Switch SNMP agent, which enables a network management workstation to receive configuration values, traffic information, and FC failure data through SNMP and the Ethernet interface

For more information, see the *QLogic Virtual Fabric Extension Module Installation Guide*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5083051>

Additional information is available on the QLogic website:

http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/Product_detail_new.aspx?oemid=324&companyid=6

You can find the following publications at that address:

- ▶ *QLogic Virtual Fabric Extension Module QuickTools User Guide*
- ▶ *QLogic Virtual Fabric Extension Module Command Line Interface Guide*
- ▶ *QLogic Virtual Fabric Extension Module Message Guide*
- ▶ *QLogic Virtual Fabric Extension Module CIM Guide*
- ▶ *QLogic Virtual Fabric Extension Module SNMP Guide*

2.10 SAS I/O modules

SAS I/O modules provide affordable storage connectivity for the BladeCenter chassis using SAS technology to create a simple fabric for external shared or non-shared storage attachments. The modules can also perform RAID controller functions inside the BladeCenter S chassis for HDDs installed into DSMs and external EXP3000 and EXP2500 expansions.

Table 2-75 lists the SAS I/O modules available.

Table 2-75 SAS I/O modules

Part number	Feature code	Description	Number of external ports	External port bandwidth	Page
39Y9195	2980	SAS Connectivity Module	4	3 Gbps	169
43W3584	3734	SAS RAID Controller Module ^a	4	3 Gbps	171

a. Supported in the BladeCenter S chassis only

2.10.1 SAS Connectivity Module

The SAS Connectivity Module is used for affordable SAS-based storage attachments in IBM BladeCenter solutions.

At the time of the writing of this book, only the following devices can be attached to this module:

- ▶ IBM System Storage DS3500, as described in 5.4, “SAS storage system interoperability” on page 461
- ▶ 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), as described in 5.10, “IBM BladeCenter S integrated storage” on page 475
- ▶ IBM System Storage EXP3000 and EXP2500, as described in 5.9, “External storage expansion (JBOD) attachments” on page 471
- ▶ Selected tape devices, as described in 5.8, “SAS tape storage for IBM BladeCenter” on page 468

Figure 2-42 shows the SAS Connectivity Module.



Figure 2-42 SAS Connectivity Module for IBM BladeCenter

Table 2-76 shows the SAS Connectivity Module ordering information.

Table 2-76 SAS Connectivity Module part numbers for ordering

Description	Part number	Feature code
SAS Connectivity Module for IBM BladeCenter	39Y9195	2980

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
 - AMM web interface for integrated storage configuration
 - CLI through the 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 2-77 lists the SAS cables that can be used with the SAS Connectivity Module. There are four cables in each order, and the cables have mini-SAS (SFF-8088) male connectors on both sides.

Table 2-77 SAS cables supported by SAS Connectivity Module

Part number	Feature code	Description
39R6531	3707	IBM 3 m SAS Cable
39R6529	3708	IBM 1 m SAS Cable

For more information, see *SAS Connectivity Module Installation and User Guide*:

<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5072374>

2.10.2 SAS RAID Controller Module

The SAS RAID Controller Module for IBM BladeCenter S provides integrated RAID and SAS-based storage area networking capabilities for IBM BladeCenter S-based solutions.

For detailed compatibility information about supported combinations of chassis, blades, storage devices, and operating systems, see 5.10, “IBM BladeCenter S integrated storage” on page 475. Figure 2-43 shows the SAS RAID Controller Module.



Figure 2-43 SAS RAID Controller Module for IBM BladeCenter S

The SAS RAID Controller includes, as standard, a battery backup unit (BBU) (Figure 2-44).



Figure 2-44 Battery Backup Unit for SAS RAID Controller Module

Also, it is possible to order IBM BladeCenter S SAS RAID Controller Module Battery Option, which provides a battery replacement for the IBM BladeCenter RAID Battery Backup Module.

Table 2-78 shows the SAS RAID Controller Module ordering information.

Table 2-78 SAS RAID Controller Module part numbers for ordering

Description	Part number	Feature code
SAS RAID Controller Module for IBM BladeCenter S	43W3584	3734
IBM BladeCenter S SAS RAID Controller Module Battery	00Y3447	5476

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 a 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 a BladeCenter S chassis.
- ▶ 1 Gb Ethernet interface for RAID management.
- ▶ RAID levels supported: 0, 1, 5, and 10.
- ▶ 512 MB of ECC battery-backed cache per controller.

- ▶ A 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, and so on).
- ▶ 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 methods (all management connections go through the management module except RAID management, which goes through the Ethernet switch installed in I/O bay 1):

- ▶ IBM Start Now Advisor.
- ▶ AMM 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 the status of the module components.
- ▶ 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.

IP addresses: Each SAS RAID Controller Module requires two IP addresses: one for RAID management, and one for the management of I/O module itself.

For more information, see the following documentation:

- ▶ *SAS RAID Controller Module Installation and User Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5078040>
- ▶ *SAS RAID Controller Module Interoperability Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5078491>
- ▶ *SAS RAID Controller Module Detailed Host Attachment Guide (Remote Boot included):*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5078491>

2.11 Pass-through and interconnect modules

The following topics are described in this section:

- ▶ Intelligent Copper Pass-thru Module for IBM BladeCenter
- ▶ Multi-Switch Interconnect Module
- ▶ Multi-Switch Interconnect Module for BladeCenter HT
- ▶ 10 Gb Ethernet Pass-thru Module

Pass-through modules are I/O options that can be used for direct connections of blades ports to the external infrastructure devices, such as network or storage switches. The pass-through

module is almost like a traditional network patch-panel, but it routes internal blade ports to the outside.

There are two pass-through modules available:

- ▶ Intelligent Copper Pass-thru Module (ICPM)
- ▶ 10 Gb Ethernet Pass-thru Module

Important: The QLogic 4 Gb and 8 Gb Intelligent Pass-thru Modules are listed in Table 2-79, but they are more like Fibre Channel switch modules, which are described in 2.9, “Fibre Channel switch modules” on page 153.

ICPM is used in standard I/O bays (bays 1 – 4), and it provides Gigabit Ethernet pass-through connectivity with 14 external RJ-45 connectors, one for every blade server. There is no need for specific octopus cables, and standard Ethernet patch cables can be used.

The 10 Gb Ethernet Pass-thru Module is used in high-speed I/O bays (bays 7 - 10). It routes 10 Gb Ethernet signals from 10 Gb Ethernet high-speed expansion cards on blades to the outside with one-to-one mapping, and it occupies a single high-speed bay.

Interconnect modules are represented by the Multi-Switch Interconnect Modules (MSIM and MSIM-HT), which are used with the BladeCenter H (MSIM) and BladeCenter HT (MSIM-HT) chassis. They provide a means for existing standard switch modules to communicate with blade servers using the high-speed fabric. These modules are installed into two adjacent high-speed switch slots, which increases the number of standard ports per blade up to eight without density losses.

Table 2-79 lists the available pass-through and interconnect modules.

Table 2-79 Pass-through and interconnect modules for IBM BladeCenter

Part number	Feature code	Description	Page
44X1907	5482	QLogic 8 Gb Intelligent Pass-thru Module	163
88Y6410	A24D	QLogic 4/8 Gb Intelligent Pass-thru Module	163
44W4483	5452	Intelligent Copper Pass-thru Module	174
39Y9314	1465	Multi-Switch Interconnect Module	176
44R5913	5491	Multi-Switch Interconnect Module for BladeCenter HT	181
46M6181	1641	10 Gb Ethernet Pass-thru Module for IBM BladeCenter	184

2.11.1 Intelligent Copper Pass-thru Module for IBM BladeCenter

The Intelligent Copper Pass-thru Module for IBM BladeCenter (ICPM) 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.

For supported I/O bays and expansion cards, see Table 2-31 on page 80 and Table 2-32 on page 82.

Figure 2-45 shows the pass-through module.



Figure 2-45 Intelligent Copper Pass-thru Module for IBM BladeCenter

Table 2-80 shows the Intelligent Copper Pass-thru Module ordering information.

Table 2-80 Intelligent Copper Pass-thru Module part numbers for ordering

Description	Part number	Feature code
Intelligent Copper Pass-thru Module for IBM BladeCenter	44W4483	5452

The module's features are as follows:

- ▶ Ports:
 - Fourteen external RJ-45 copper ports for making connections to the network infrastructure. The external ports can operate at 100 or 1000 Mbps.
 - Fourteen internal bidirectional 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, the corresponding internal link goes down also. If an internal link goes down, the corresponding external link goes down also.
 - Serial over LAN (SOL) and cKVM technologies are supported by this module.

- ▶ Monitoring:
 - ICPM has informational LEDs on the front panel for indicating the I/O module status. Each Ethernet port on ICPM has two LEDs to indicate the link status and port activity.
 - The module supports self-test diagnostic tests.
- ▶ Management
 - This module is managed through the AMM management interface.
- ▶ Cables:
 - 100BASE-TX:
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568 100-ohm STP (100 m maximum)
 - 1000BASE-T:
 - UTP Category 6
 - UTP Category 5e (100 m maximum)
 - UTP Category 5 (100 m maximum)
 - EIA/TIA-568B 100-ohm STP (100 m maximum)

For more information, see the *Intelligent Copper Pass-thru Module Installation Guide*:

<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5076328>

2.11.2 Multi-Switch Interconnect Module

The Multi-Switch Interconnect Module (MSIM) 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 supported CFFh-type adapter installed. See Table 2-83 on page 178 for the expansion cards currently supported and the I/O modules that are supported with them.

Figure 2-46 shows the MSIM.




Figure 2-46 Multi-Switch Interconnect Module

I/O modules for the MSIM: 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.

Table 2-81 shows the Multi-Switch Interconnect Module ordering information.

Table 2-81 Multi-Switch Interconnect Module part numbers for ordering

Description	Part number	Feature code
Multi-Switch Interconnect Module for IBM BladeCenter	39Y9314	1465

Each MSIM occupies high-speed slots 7 and 8 or 9 and 10, and up to two MSIMs can be installed into the IBM BladeCenter H chassis. To use the MSIM, each blade must have certain CFFh expansion cards installed. For more information, see Table 2-31 on page 80 and Table 2-32 on page 82.

The ports on the CFFh expansion cards in each server are hardwired to specific bays in the switch modules in each MSIM. Table 2-82 lists the mappings of expansion card ports to the I/O bays of MSIMs.

Table 2-82 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 2-83 on page 178 lists the supported I/O module options for the MSIM.

Currently, only two expansion cards support connectivity to the MSIM. The placement of modules in the MSIM varies between these expansion cards, as shown in the rightmost two columns of Table 2-83.

Table 2-83 Supported I/O module options for the MSIM

Part number	Feature code		MSIM I/O bay ^a	
		Supported expansion card → Description	QLogic Ethernet and 8 Gb Fibre Channel Card (CFFh), 00Y3270	2/4 Port Ethernet Expansion Card (CFFh), 44W4479
Ethernet				
69Y1909	7656	Brocade Converged 10GbE Switch Module	Not supported	Not supported
00Y3254 ^b	A3FD	Cisco Catalyst Switch Module 3110G	Left	Both bays
00Y3250 ^c	A3FC	Cisco Catalyst Switch Module 3110X	Left	Both bays
46C9272 ^d	A3FE	Cisco Catalyst Switch Module 3012	Left	Both bays
46C9270 ^e	A3FF	Cisco Nexus 4001I Switch Module	Not supported	Not supported
39Y9324	1484	Server Connectivity Module	Left	Both bays
32R1859	1494	IBM Layer 2-7 Gb Ethernet Switch	Not supported	Not supported
32R1860	1495	IBM Layer 2/3 Copper Gb Ethernet Switch	Left	Both bays
32R1861	1496	IBM Layer 2/3 Fiber Gb Ethernet Switch	Left	Both bays
44W4404	1590	IBM 1/10Gb Uplink ESM	Left	Both bays
46C7191	1639	IBM Virtual Fabric 10Gb Switch Module	Not supported	Not supported
Fibre Channel				
42C1828	5764	Brocade Enterprise 20-port 8Gb SAN SM ^f	Right	Not supported
44X1920	5481	Brocade 20-port 8Gb SAN Switch Module ^f	Right	Not supported
44X1921	5483	Brocade 10-port 8Gb SAN Switch Module ^f	Right	Not supported
44E5696 ^g	A3FH	Cisco Systems 4Gb 20 port Fibre Channel Switch	Right	Not supported
44E5692 ^h	A3FG	Cisco Systems 4Gb 10 port Fibre Channel Switch	Right	Not supported
44X1905	5478	QLogic 20-Port 8Gb SAN Switch Module	Right	Not supported
88Y6406	A24C	QLogic 20-Port 4/8Gb SAN Switch Module	Right	Not supported
44X1907	5482	QLogic 8Gb Intelligent Pass-thru Module	Right	Not supported
88Y6410	A24D	QLogic 4/8Gb Intelligent Pass-thru Module	Right	Not supported
46M6172	4799	QLogic Virtual Fabric Extension Module	Not supported	Not supported

Part number	Feature code		MSIM I/O bay ^a	
		Supported expansion card → Description	QLogic Ethernet and 8 Gb Fibre Channel Card (CFFh), 00Y3270	2/4 Port Ethernet Expansion Card (CFFh), 44W4479
SAS				
39Y9195	2980	SAS Connectivity Module	Not supported	Not supported
43W3584	3734	SAS RAID Controller Module	Not supported	Not supported
InfiniBand				
46M6005	0057	Voltaire 40Gb InfiniBand Switch Module	Not supported	Not supported
Pass-thru Modules				
44W4483	5452	Intelligent Copper Pass-thru Module	Left	Both bays
44X1907	5482	QLogic 8Gb Intelligent Pass-thru Module	Right	Not supported
88Y6410	A24D	QLogic 4/8Gb Intelligent Pass-thru Module	Right	Not supported
46M6181	1641	10Gb Ethernet Pass-thru Module	Not supported	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.

b. Replaces 41Y8523.

c. Replaces 41Y8522.

d. Replaces 43W4395.

e. Replaces 46M6071.

f. Brocade 8 GB SAN Switch modules support only QLogic Ethernet and 8 Gb FC Expansion Card (CFFh), 00Y3270, when used in MSIM.

g. Replaces 39Y9280.

h. Replaces 39Y9284.

Figure 2-47 shows 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 figure shows one sample configuration.

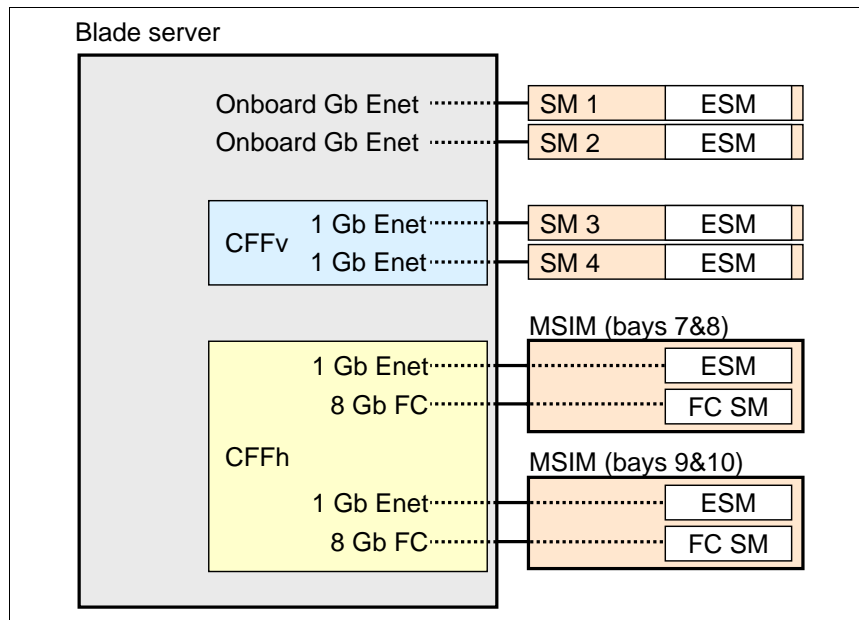


Figure 2-47 Example of eight I/O paths: six 1 Gb Ethernet and two 8 Gb Fibre Channel

The connections are as follows:

- ▶ The blade server has two onboard Gigabit Ethernet adapters that connect to ESMs in bay 1 and bay 2.
- ▶ The blade server has a CFFv (or CIOv) Gigabit Ethernet expansion card installed. The CFFv or CIOv 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 8 Gb Fibre Channel ports. See 3.1.12, "QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)" on page 238 for a description of this card.

The ports on the CFFh card connect to switch modules that are installed in an 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.

Important: A BladeCenter chassis cannot have both an MSIM and high-speed switch module (HSSM) installed at the same time.

For more information, see *Installation and User's Guide Multi-Switch Interconnect Module - IBM BladeCenter*.

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5070476>

2.11.3 Multi-Switch Interconnect Module for BladeCenter HT

The Multi-Switch Interconnect Module for BladeCenter HT (MSIM-HT) 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 2-46 on page 176, and enables a blade server to have a total of eight network connections.

MSIM-HT and I/O modules: 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 2-48 Multi-Switch Interconnect Module for BladeCenter HT (without interposers)

Table 2-84 shows Multi-Switch Interconnect Module for BladeCenter HT ordering information.

Table 2-84 Multi-Switch Interconnect Module-HT part numbers for ordering

Description	Part number	Feature code
Multi-Switch Interconnect Module for BladeCenter HT	44R5913	5491

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. For more information, see Table 2-31 on page 80 and Table 2-32 on page 82.

The ports on the CFFh expansion cards in each server are hardwired to specific bays in the switch modules in each MSIM-HT. Table 2-85 lists mappings of expansion card ports to the I/O bays of MSIMs.

Table 2-85 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)

An 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

installed in the upper bays (7 and 8). The other interposer is used when the MSIM is installed in the lower bays (9 and 10). Figure 2-49 shows the interposer for use with the lower bays.



Figure 2-49 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 2-86 lists the supported I/O module options for the MSIM-HT.

Important: At the time of writing, HS22 (7870) is the only supported blade server for use with the MSIM-HT.

Currently, only two expansion cards support connectivity to the MSIM-HT. As illustrated in Table 2-86, the placement of modules in the MSIM-HT varies between these expansion cards, as shown in the rightmost two columns.

Table 2-86 Supported I/O module options for the MSIM-HT

Part number	Feature code		MSIM-HT I/O bay ^a	
		Supported expansion cards →	QLogic Ethernet and 8 Gb Fibre Channel Card (CFFh), 00Y3270	2/4 Port Ethernet Expansion Card (CFFh), 44W4479
		Description		
Ethernet				
69Y1909	7656	Brocade Converged 10 GbE Switch Module	Not supported	Not supported
00Y3254 ^b	A3FD	Cisco Catalyst Switch Module 3110G	Left	Both bays
00Y3250 ^c	A3FC	Cisco Catalyst Switch Module 3110X	Left	Both bays
46C9272 ^d	A3FE	Cisco Catalyst Switch Module 3012	Left	Both bays
46C9270 ^e	A3FF	Cisco Nexus 4001I Switch Module	Not supported	Not supported
39Y9324	1484	Server Connectivity Module	Not supported	Not supported
32R1859	1494	IBM Layer 2-7 Gb Ethernet Switch	Not supported	Not supported
32R1860	1495	IBM Layer 2/3 Copper Gb Ethernet Switch	Left	Both bays
32R1861	1496	IBM Layer 2/3 Fiber Gb Ethernet Switch	Left	Both bays

Part number	Feature code	Description	MSIM-HT I/O bay ^a	
			QLogic Ethernet and 8 Gb Fibre Channel Card (CFFh), 00Y3270	2/4 Port Ethernet Expansion Card (CFFh), 44W4479
44W4404	1590	IBM 1/10Gb Uplink Ethernet Switch Module	Left	Both bays
46C7191	1639	IBM Virtual Fabric 10Gb Switch Module	Not supported	Not supported
Fibre Channel				
42C1828	5764	Brocade Enterprise 20-port 8Gb SAN SM	Not supported	Not supported
44X1920	5481	Brocade 20-port 8Gb SAN Switch Module	Right	Not supported
44X1921	5483	Brocade 10-port 8Gb SAN Switch Module	Right	Not supported
44E5696 ^f	A3FH	Cisco Systems 4Gb 20 port Fibre Channel Switch	Right	Not supported
44E5692 ^g	A3FG	Cisco Systems 4Gb 10 port Fibre Channel Switch	Right	Not supported
44X1905	5478	QLogic 20-Port 8Gb SAN Switch Module	Right	Not supported
88Y6406	A24C	QLogic 20-Port 4/8Gb SAN Switch Module	Right	Not supported
44X1907	5482	QLogic 8Gb Intelligent Pass-thru Module	Not supported	Not supported
88Y6410	A24D	QLogic 4/8Gb Intelligent Pass-thru Module	Not supported	Not supported
46M6172	4799	QLogic Virtual Fabric Extension 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				
46M6005	0057	Voltaire 40Gb InfiniBand Switch Module	Not supported	Not supported
Pass-thru Modules				
44W4483	5452	Intelligent Copper Pass-thru Module	Not supported	Not supported
44X1907	5482	QLogic 8Gb Intelligent Pass-thru Module	Not supported	Not supported
88Y6410	A24D	QLogic 4/8Gb Intelligent Pass-thru Module	Not supported	Not supported
46M6181	1641	10Gb Ethernet 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.

b. Replaces 41Y8523.

c. Replaces 41Y8522.

d. Replaces 43W4395.

e. Replaces 46M6071.

f. Replaces 39Y9280.

g. Replaces 39Y9284.

For more information, see *Multi-Switch Interconnect Module Installation Instructions - IBM BladeCenter HT*:

<http://ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5073195>

2.11.4 10 Gb Ethernet Pass-thru Module

The 10 Gb Ethernet Pass-thru Module for IBM BladeCenter is ideal for clients who want to enable an end-to-end non-blocking 10 Gb setup within the chassis. This device supports both Ethernet and Converged Enhanced Ethernet (CEE) packets, which allows clients to connect a BladeCenter Chassis to an FCoE-capable top-of-rack switch.

The 14 10-Gb Uplink ports are based on optical SFP+ technology to offer the highest performance while maintaining industry standard connectivity. This offering also works with BladeCenter Open Fabric Manager, providing all the benefits of I/O Virtualization at 10 Gb speeds.

Clients can use this module along with the two-port 10 Gb Converged Network Adapter (CNA) and connect to a converged top-of-rack switch, such as the IBM Converged Switch B32. This setup allows clients to reduce hardware, and power and cooling costs, while boosting performance by operating at 10 Gb bandwidth.

The 10 Gb Ethernet Pass-thru Module is shown on Figure 2-50.



Figure 2-50 10 Gb Ethernet Pass-thru Module

Table 2-87 shows the 10 Gb Ethernet Pass-thru Module ordering information.

Table 2-87 10 Gb Ethernet Pass-thru Module part numbers for ordering

Description	Part number	Feature code
10 Gb Ethernet Pass-thru Module for BladeCenter	46M6181	1641
IBM 10GBase-SR 10 GbE 850 nm Fiber SFP+ Transceiver	44W4408	4942

The module part numbers include the following items:

- ▶ 10 Gb Ethernet Pass-thru Module
- ▶ One HSSM filler module
- ▶ *IBM Important Notices* document
- ▶ Documentation CD-ROM

SFP+ transceivers: Small form-factor pluggable plus (SFP+) transceivers are not included and must be purchased separately.

The 10 Gb Ethernet Pass-thru Module includes the following features and functions:

- ▶ Single-wide high-speed switch module.
- ▶ Fourteen internal 10 Gb ports to the server blades (no auto-negotiation).
- ▶ Up to fourteen 10 Gb SFP+ uplink ports (SFP+ modules are not included).
- ▶ Support for SR, LR, and DAC copper cables.
- ▶ Direct one-to-one mappings of external and internal ports (no configuration required).
- ▶ Simple setup and monitoring through AMM.
- ▶ A management-module inter-integrated Circuit (I2C) interface that provides vital product data (VPD) and register access.
- ▶ Provides self-test and diagnostics capability.
- ▶ No Serial over LAN (SOL) support.

The following network cable is required for the 10 Gb Ethernet Pass-thru Module 10GBASE-SR for 10 Gb ports:

- ▶ 850 nm wavelength, multimode fiber, 50 μ or 62.5 μ (300 m maximum), with LC duplex connector.

For supported I/O bays and expansion cards, see Table 2-31 on page 80 and Table 2-32 on page 82.

For more information, see the *10 Gb Ethernet Pass-thru Module Installation Guide*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5080966>

2.12 Advanced Management Module

An AMM is a hot-swap device that you use to configure and manage all installed BladeCenter components. It provides system management functions and KVM multiplexing for all the blade servers in the BladeCenter chassis. It controls an Ethernet and serial port connections for remote management access.

Every BladeCenter chassis includes, as standard, one or two (model-dependent) AMMs. Only one management module is currently available (Table 2-88):

- ▶ The AMM for BladeCenter S, BladeCenter E, BladeCenter H, and BladeCenter HT

Important: If your existing BladeCenter E has a Standard Management Module, then you might need to replace it with an AMM to support newer blades, I/O modules, and software functions that are used in this chassis.

Table 2-88 Management module supported in each BladeCenter chassis

Management module	Part number	Feature code
Advanced Management Module	2019A1X	1604

To obtain the latest firmware for your management module, see the AMM firmware release matrix:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=SERV-AMM>

Table 2-89 lists all the possible supported combinations of BladeCenter chassis and management modules, and summarizes a list of features that are supported by each configuration.

Table 2-89 Features supported by chassis and management modules installed

	BC S	BC E	BC H	BC HT
Feature	AMM	AMM	AMM	AMM
Blade servers	Yes	Yes	Yes	Yes
Standard switch modules	Yes	Yes	Yes	Yes
High-speed switch modules	No	No	Yes	Yes
Concurrent KVM (cKVM) ^a	Yes	Yes	Yes	Yes
External USB ports	Yes	Yes	Yes	Yes
External PS/2 ports	No	No	No	No
External serial port	Yes	Yes	Yes	Yes
Active Energy Manager	Yes	Yes	Yes	Yes
IBM Fabric Manager	Yes	Yes	Yes	Yes

a. The blade itself must support this feature either by installing a Concurrent KVM feature card, such as on HS21 (8853), or by having such built-in functionality with an Integrated Management Module, such as on HS22 (7870).

The term *standard switch module* refers to the traditional switch module, such as the IBM Layer 2/3 Gigabit Ethernet Switch Module, which has a form-factor compatible with I/O bays 1, 2, 3, and 4 of each BladeCenter chassis. See the drawings of the rear of each chassis in 1.2, “BladeCenter chassis” on page 16.

The term *high-speed switch module* refers to a high-speed switch module, such as IBM Virtual Fabric 10 Gb Switch Module, which has a form-factor compatible with I/O bays 7, 8, 9, and 10 of IBM BladeCenter H and HT chassis. See Figure 2-8 on page 48.

Mixing management modules: 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 AMM has various 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 using IPv4 or IPv6 protocols.

IPv6: The AMM support of IPv6 is from firmware level BPET54D and later.

- The serial port in the front of the AMM 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 through an Ethernet port.

- ▶ The AMM 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 AMM and the Management Module allows remote control of the mouse and keyboard (remote KVM).

The AMM also supports concurrent remote KVM access if each blade server you want to be accessible concurrently has such built-in functionality with IMM or the cKVM feature card installed. It 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, although the maximum number of remote control session per blade can be limited to one if required.

The AMM enables the Chassis Internal Network (CIN). The CIN is an internal communication path between the blade servers and the AMM. When the CIN is enabled and configured properly, the AMM can fully use all the resources, such as LDAP, SMTP, Director, SNMP, DNS, or NTP services, installed on the blade servers. Conversely, the blade servers can directly connect to the AMM through the CIN path. The CIN is internal to the BladeCenter chassis and is on a private user-defined VLAN. For more information about the CIN, see the document *Enabling and Configuring BladeCenter Chassis Internal Network*.

ftp://ftp.software.ibm.com/systems/support/system_x_pdf/cin_whitepaper_7.1.pdf

Additionally, the AMM supports Active Energy Manager and IBM Fabric Manager (IFM). IFM is described in 4.8, “IBM Fabric Manager” on page 433.

2.12.1 Advanced Management Module

The Advanced Management Module (AMM) is a hot-swap module that you use to configure and manage all installed BladeCenter components. The AMM provides system management functions and 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 2-51 shows an Advanced Management Module.




Figure 2-51 Advanced Management Module

All BladeCenter chassis come standard with at least one AMM. Each chassis also supports a second management module for redundancy, excluding BladeCenter S chassis. One of the management modules is active, and the second one, if installed, remains standby until the management functions are manually switched over to it, or if the primary management module fails.

Table 2-90 shows the AMM ordering information.

Table 2-90 Advanced Management Module part numbers for ordering

Description	Part number	Feature code
Advanced Management Module for BladeCenter	2019A1X	1604

The service processor in the management module communicates with the service processor in each blade server to support 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 the 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.

The following tasks can be performed with AMM:

- ▶ Defining the login IDs and passwords
- ▶ Configuring security settings, such as data encryption and user account security
- ▶ Selecting recipients for alert notification of specific events

- ▶ Monitoring the status of the BladeCenter unit, blade servers, and other BladeCenter components:
 - Event log
 - LEDs
 - Hardware and firmware VPD
 - Fan speeds
 - Temperatures
 - Power usage
- ▶ Discovering other BladeCenter units in the network and enabling access to them through their management-module web interfaces.
- ▶ Controlling the BladeCenter unit, blade servers, and other BladeCenter components:
 - Power on/off
 - Firmware update
 - Configuration settings
 - Serial over LAN
- ▶ Configuring power management for BladeCenter unit
- ▶ Accessing the I/O modules to configure them
- ▶ Changing the startup sequence in a blade server
- ▶ Setting the date and time
- ▶ Using a remote console for the blade servers
- ▶ Mounting remote virtual media for the blade servers
- ▶ Changing ownership of the keyboard, video, and mouse
- ▶ Changing ownership of removable media drives and USB ports (the removable media drives in the BladeCenter unit are viewed as USB devices by the blade server operating system)
- ▶ Setting the active color of the critical (CRT) and major (MJR) alarm LEDs (for BladeCenter T units only)
- ▶ Using BladeCenter Open Fabric Manager functions
- ▶ Using Service Advisor functions to autonomously inform IBM Support about any critical events that happen

AMM supports the following management methods:

- ▶ Web-based interface with SSL support
- ▶ CLI through Telnet/SSH
- ▶ SMASH Command Line Protocol
- ▶ SNMP

For more information about AMM, see the following product documentation:

- ▶ *IBM BladeCenter Advanced Management Module Installation Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5073392>
- ▶ *IBM BladeCenter Advanced Management Module User's Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5073887>
- ▶ *Advanced Management Module Command Line Interface Reference Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-54667>

For more information about hardware management features, see *IBM eServer xSeries and BladeCenter Server Management*, SG24-6495.

2.13 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 configuration, a blade server must have the required wiring, and be installed in a chassis that supports this wiring.

Direct serial connections are supported by the following items:

- ▶ BladeCenter H and HT chassis with Serial Port Breakout Cable
- ▶ BladeCenter S with the Serial Module installed

Table 2-91 lists information about the currently available blades that support this wiring.

Table 2-91 Blade servers that support direct serial connections

Blade	Machine type	Direct serial connection wiring
HS12	8028	Supported
HS22	7870	Supported
HS22V	7871	No
HX5	7872	No
HX5	7873	No
PS700/PS701/PS702	8406	No
PS703/PS704	7891	No

2.13.1 Serial Port Breakout Cable

The Serial Port Breakout Cable for BladeCenter H and BladeCenter HT connects directly to a port on the rear of the BladeCenter H (Figure 2-8 on page 48) or BladeCenter HT (Figure 2-13 on page 58). This cable provides 14 serial connections for terminal access, one to each supported blade server. The cable is shown in Figure 2-52.




Figure 2-52 Serial Port Breakout Cable

Not all blade servers support this option. Table 2-91 lists the blades that support this option when the blade is installed into a BladeCenter H or HT chassis.

Table 2-92 shows Serial Port Breakout Cable ordering information.

Table 2-92 Serial Port Breakout Cable for BladeCenter part numbers for ordering

Description	Part number	Feature code
Serial Port Breakout Cable for IBM BladeCenter	40K9605	4811

Each serial connector is an RJ-45. The pin assignments of the RJ-45s are listed in Table 2-93. The maximum baud rate for serial connection is 19.2 Kbps.

Table 2-93 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

Signals: There are four signals plus a ground signal routed to each of the RJ45 connectors, as listed in Table 2-93 on page 191. 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, do 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.

2.13.2 BladeCenter S Serial Pass-thru Module

The Serial Pass-thru Module provides six serial port connectors that can be used to directly attach to each blade server in the BladeCenter S chassis through 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 19.2 K 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 1-8 on page 24). 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.

The servers listed in Table 2-91 on page 190 support the use of the Serial Pass-thru Module.

Figure 2-53 shows the Serial Pass-thru Module.




Figure 2-53 Serial Pass-thru Module

Table 2-94 shows the Serial Pass-thru Module ordering information.

Table 2-94 Serial Pass-thru Module part numbers for ordering

Description	Part number	Feature code
Serial Pass-thru Module for BladeCenter S	43W3583	1584

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 2-54.

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 2-54 RJ-45 pin location reference diagram

2.14 Installation and physical site plans

The IBM BladeCenter E, H, HT, and S chassis and all blade servers are designated as customer setups. Customer setup instructions are included with the systems but can be downloaded from the IBM Support website. See product publications for installation instructions. These product publications can be downloaded from the BladeCenter Support site:

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/brandmain?brandind=5000020>

This section includes the following topics:

- ▶ Chassis dimensions and physical planning
- ▶ Rack installation
- ▶ Power considerations
- ▶ Cooling considerations
- ▶ Power cable considerations
- ▶ Console Managers and SOL console connectivity

2.14.1 Chassis dimensions and physical planning

This section includes physical planning specifications and weight considerations for the IBM BladeCenter models.

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 102 kg (225 lb.)
 - Fully configured without blades: Approximately 38.6 kg (85 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 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, which 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, which 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 - 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, which 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 W)
 - Full configuration: 11942 Btu/hr (3500 W)
- ▶ Acoustical noise emissions for BladeCenter with six blade servers:
 - Operating: 63 - 69 dB blade-dependent
 - Idling: 63 - 69 dB blade-dependent

Weight considerations

The following factors are important when planning for the physical installation of BladeCenter units:

- ▶ Racks have a maximum weight load limit.
- ▶ Floors, especially raised floors, have maximum weight limits.

Specific requirements are available in the following documentation:

- ▶ BladeCenter E type 8677: *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 at the BladeCenter Support website:

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/brandmain?brandind=5000020>

2.14.2 Rack installation

You should 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 IBM SSCT from the following address:

<http://www.ibm.com/systems/x/configtools.html>

When planning for racks, 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 requires the following rack space:
 - BladeCenter E: 7U
 - 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 be installed horizontally only.
- ▶ 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 HT can be installed into a two-post or a four-post rack.
- ▶ The BladeCenter HT unit is designed to be rack-compatible with NEBS-compliant four-post racks that have size 12 - 24 tapped holes.

Specific requirements are available in the following documentation:

- ▶ BladeCenter type 8677: *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 this documentation under Publications in the relevant product section at the following address:

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/brandmain?brandind=5000020>

Other rack kits: IBM tests 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 needs 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 described in the following sections:

- ▶ Power modules for BladeCenter E
- ▶ Power modules for IBM BladeCenter H
- ▶ Power modules for BladeCenter HT
- ▶ Power modules for BladeCenter S

This section provides additional information.

Voltage: The BladeCenter E chassis type 8677, BladeCenter H chassis type 8852, and BladeCenter HT chassis type 8750 (AC) can be powered only with 220 V AC 50/60 Hz worldwide. Your data center must be able to supply this voltage.

The 48 V DC-powered BladeCenter HT chassis type 8740 must be installed by personnel trained in DC wiring installations. Setup instructions are included with the systems.

IBM BladeCenter power management

The IBM BladeCenter hardware architecture provides energy efficiency through its design. In addition, 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 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.

As part of the Active Energy Manager solution, the BladeCenter power management policies are 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. Power management policy is selectable at the power domain level (that is, two separate power domains can have two separate power management policies).

The power management policies are as follows:

- ▶ AC Power Source Redundancy with Blade Throttling Allowed (only for BladeCenter S)
- ▶ AC Power Source Redundancy (only for BladeCenter S)
- ▶ Basic Power Management
- ▶ Power Module Redundancy
- ▶ Power Module Redundancy with Blade Throttling Allowed

These policies control how much electrical power is available to power domain components, and 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 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.

In general, *Basic Power Management* means that maximum power is available for the domain (restricted by the maximum power domain rating), and power supply failure does not necessarily mean that blades are 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.

Power Module Redundancy means N+1 redundancy for the domain. It also means that the power capacity available is equal to the capacity of one or three 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 the maximum power capacity is equal to the capacity of one power supply (essentially, N+1 redundancy becomes N+N because N=1). If a power domain has four power supplies (only applicable to the BladeCenter S), then the maximum capacity is the sum of capacities for three power supplies. No throttling occurs in this scenario. If the blade exceeds the power capacity of the domain, it is not powered up.

Power Module 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+N redundancy (with dual AC power sources) for four power supplies (2+2) in the same manner as described previously. When **AC Power Source Redundancy** is selected, then maximum power capacity is equal to the sum of capacities for two power supplies, and no throttling occurs. *AC Power Source Redundancy with Blade Throttling Allowed* means that the 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, and other characteristics. This tool can be downloaded from the following address:

<http://www.ibm.com/systems/bladecenter/powerconfig/>

Throttling: If throttling is unacceptable, then you must fit the configuration into the available power envelope. You accomplish this task by 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 Power Module Redundancy power management policy.

Additional power considerations for specific chassis are covered later in this book.

For a detailed description of power management in the IBM BladeCenter, see *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 that were shipped earlier than at the time of the writing this book might have different types of power supplies:

- ▶ 1,200 W
- ▶ 1,400 W
- ▶ 1,800 W
- ▶ 2,000 W
- ▶ 2,320 W

Only 2,320 W power supplies are fully supported with all blade servers shipping currently. If you plan to install new blades into all bays of an existing BladeCenter E chassis, you must plan to upgrade your power supplies to 2,320 W 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 2,000 W and 2,320 W power supplies are available now. All others were withdrawn from market. For more information about this topic, see 2.1.5, “Power modules for BladeCenter E” on page 45.

To provide true redundant power, BladeCenter E power modules 1 and 3 must be connected to a different 200 - 240 V 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 6,938 W.

IBM BladeCenter E should be attached to high-voltage PDUs, as described in 2.14.5, “Power cable considerations” on page 204.

For additional information, see 2.1.5, “Power modules for BladeCenter E” on page 45.

IBM BladeCenter H specific considerations

A BladeCenter H chassis includes two 2,900 W or 2,980 W power supplies (model-dependent). Additional 2,900 W or 2,980 W power supplies for power domain B are available as an option, as described in 2.2.4, “Power modules for IBM BladeCenter H” on page 53.

Remember that power cables for BladeCenter H are not included 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 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 9,600 W.

IBM BladeCenter H should be attached to high-voltage PDUs, as described in 2.14.5, “Power cable considerations” on page 204.

For additional information, see 2.2.4, “Power modules for IBM BladeCenter H” on page 53.

IBM BladeCenter HT specific considerations

An IBM BladeCenter HT includes two 3,160 W AC (8750) or 2,535 W DC (8740) power supplies. Additional AC or DC power supplies for power domain B are available as an option, as described in 2.3.5, “Power modules for BladeCenter HT” on page 65.

Power supply modules: Models 8740-2Rx and 8750-2Rx include, as standard, four power supply modules.

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.

In any configuration, the total power consumption for the single IBM BladeCenter HT chassis type 8740 must not exceed 5,776 W, and for chassis type 8750, no more than 6,400 W.

IBM BladeCenter HT 8750 should be attached to high-voltage PDUs, as described in 2.14.5, “Power cable considerations” on page 204.

IBM BladeCenter HT 8740 should be attached to the appropriate DC wiring, as described in 2.14.5, “Power cable considerations” on page 204.

For additional information, see 2.3.5, “Power modules for BladeCenter HT” on page 65.

IBM BladeCenter S specific considerations

Unlike the other BladeCenter chassis, the BladeCenter S can operate from a 110 V AC power source and also a 220 V power source. 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 includes two 950 W (110 V)/1450 W (220 V) AC power supplies. An additional power supplies option is available, as described in 2.4.4, “Power modules for BladeCenter S” on page 73.

In any configuration, the total power consumption for the single IBM BladeCenter S chassis must not exceed 3500 W.

BladeCenter S can be attached to a supported PDU or to office wall jacks, as described in 2.14.5, “Power cable considerations” on page 204.

For more information, see 2.4.4, “Power modules for BladeCenter S” on page 73.

2.14.4 Cooling considerations

When planning a data center, consider the following points regarding cooling.

Air temperature

The permissible temperatures and humidity limits for various BladeCenter chassis are as follows:

- ▶ IBM BladeCenter E:
 - On: 10° - 35°C (50° - 95°F) at 0 - 914 m (0 - 3,000 ft.)
 - On: 10° - 32°C (50° - 90°F) at 914 - 2,133 m (3,000 - 7,000 ft.)
 - Off: -40° - 60°C (-40° - 140°F)
 - Relative humidity:
 - On: 8% - 80%
 - Off: 8% - 80%
 - Maximum altitude: 2,133 m (7,000 ft.)
- ▶ IBM BladeCenter H:
 - On: 10° - 35°C (50° - 95°F) at 0 - 914 m (0 - 3,000 ft.)
 - On: 10° - 32°C (50° - 90°F) at 914 - 2,133 m (3,000 - 7,000 ft.)
 - Off: -40° - 60°C (-40° - 140°F)
 - Relative humidity:
 - On: 8% - 80%
 - Off: 8% - 80%
 - Maximum altitude: 2,133 m (7,000 ft.)
- ▶ IBM BladeCenter HT 8740 and 8750
 - On: 5° - 40°C (41° - 104°F) at altitude: -60 - 1,800 m (-197 ft. - 6,000 ft)
 - On: 5° - 30°C (41° - 86°F) at altitude: 1,800 m - 4,000 m (6,000 ft. - 13,000 ft)
 - On (short term): -5° - 55°C (23° - 131°F) at altitude: -60 - 1,800 m (-197 ft. - 6,000 ft.)
 - On (short term): -5° - 45°C (23° - 113°F) at altitude: 1800 m - 4000 m (6,000 ft. - 13,000 ft.)
 - Off: -40° - 70°C (-40° - 158°F), 30°C/hour maximum rate of change
 - Relative humidity: 5% - 85%
 - Relative humidity (short term): 5% - 90% not to exceed 0.024 water/kg of dry air
 Short term is a period of not more than 96 consecutive hours and a total of not more than 15 days in one year. (A total of 360 hours in any year, but no more than 15 occurrences during that one-year period.)
- ▶ IBM BladeCenter S:
 - Temperature:
 - On: 10° - 35°C (50° - 95°F) at 0 - 914 m (0 - 3,000 ft.)
 - On: 10° - 32°C (50° - 90°F) at 914 - 2,133 m (3,000 - 7,000 ft.)
 - Off: -40° - 60°C (-40° - 140°F)
 - Relative humidity: 8% - 80%
 - Maximum altitude: 2133 m (7,000 ft.)

Air flow

Air flow is critical for ensuring that the operating air temperature stays within permissible limits. Keep in mind the following information when planning your installation:

- ▶ 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 1 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 setup 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 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 maximum amount of heat output of the BladeCenter chassis in BTU per hour is shown in Table 2-95.

Table 2-95 Maximum heat output for IBM BladeCenter chassis configurations

Chassis	Maximum heat output
IBM BladeCenter E	23672 BTU/hour (6,938 W)
IBM BladeCenter H	32757 BTU/hour (9,600 W)
IBM BladeCenter HT 8740	19680 BTU/hour (5,766 W)
IBM BladeCenter HT 8750	21850 BTU/hour (6,400 W)
IBM BladeCenter S	11942 BTU/hour (3,500 W)

Preventing air recirculation

Consider these factors when planning for single or multiple rack installations:

- ▶ When racks are positioned near 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 in mind these points 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 that has a cooling system that does not have sufficient airflow volume and cooling capacity. This situation 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, the Rear Door Heat eXchanger (Figure 2-55) is a more cost-effective solution than adding another air conditioning unit.

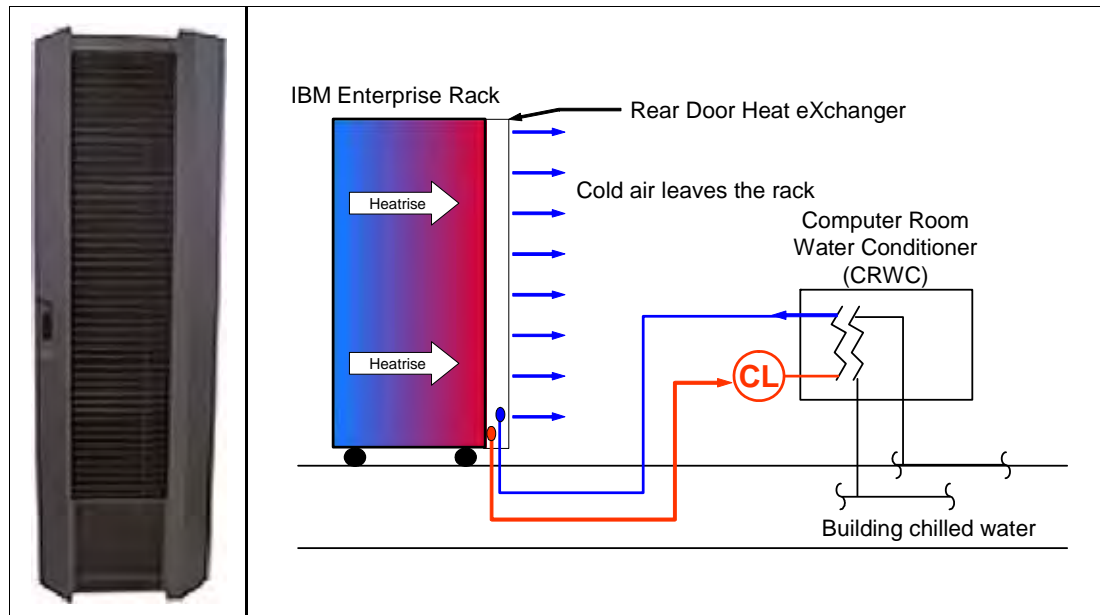


Figure 2-55 Rear Door Heat eXchanger (left) and functional diagram

BladeCenter cooling: 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 because of 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 in a 9308 Enterprise rack.
- It offers a 1-year limited warranty.

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 a convenient way to handle dangerous “hot spots” in your data center.

2.14.5 Power cable considerations

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, see 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. These connections can be either one-phase or three-phase connections with different ampere loads.

Power cords: Power cords are not included for countries in the EMEA (Europe, Middle East, Africa) region.

Some general rules can be followed for choosing the 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 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 enough 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 2-12 on page 55 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 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 or four 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 enough 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-C14 power connectors (model-dependent) when connected to each installed power module.
 - The BladeCenter S unit comes standard with two IEC C13 - C14 power cables, and four country-specific power cords. The additional power supplies option contains two additional 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 the following address:

<http://www.ibm.com/systems/x/configtools.html>

For more information about PDUs and PDU planning and cabling considerations, see:

- ▶ *IBM BladeCenter Power Guide - North America:*
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101706>
- ▶ *IBM System x PDU Guide - North America:*
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101526>
 - *IBM System x PDU Guide - International:*
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101615>
- ▶ *IBM System x Uninterruptable Power Supply (UPS) Guide:*
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101655>

We also suggest that you read the information available in the following guides:

- ▶ *BladeCenter type 8677 - 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, BladeCenter H 8852, or BladeCenter HT 8740 or 8750) at:

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

2.14.6 Console Managers and SOL console connectivity

The IBM BladeCenter management module acts as a KVM switch for blade servers installed into the chassis. The console (keyboard, video, and mouse) can be connected to the management module directly, or through the KVM switch.

The advanced management module for IBM BladeCenter E, BladeCenter H, BladeCenter HT, and BladeCenter S has USB ports for keyboard and mouse.

Ensure that you select the correct 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 2-96.

Table 2-96 Supported console switches and cables

	IBM 1x8 Console Switch (17353LX) IBM 2x16 Console Switch (17354LX) IBM Local 1x8 Console Manager (1754A1X) IBM Local 2x16 Console Manager (1754A2X)	IBM Local 2x8 Console Manager (17351GX) IBM Global 2x16 Console Manager (17352GX) IBM Global 4x16 Console Manager (17354GX)
IBM BladeCenter E with AMM	39M2895 IBM USB Conversion Option	39M2895 IBM USB Conversion Option
IBM BladeCenter H with AMM	39M2895 IBM USB Conversion Option	39M2895 IBM USB Conversion Option
IBM BladeCenter HT with AMM	39M2895 IBM USB Conversion Option	39M2895 IBM USB Conversion Option

For more information, see the following publications:

- *IBM 1754 LCM8 and LCM16 Local Console Managers:*
<http://www.ibm.com/redbooks/abstracts/tips0788.html>
- *IBM 1754 GCM16 and GCM32 Global Console Managers:*
<http://www.redbooks.ibm.com/abstracts/tips0772.html>
- *IBM Rack-Based Local Console Switches:*
<http://www.redbooks.ibm.com/abstracts/tips0730.html>

Providing an SOL console for the blade servers

Most blade servers support Serial over LAN (SOL) technology to get remote access to the text console from the administrator side instead of using a traditional KVM console.

The SOL remote text console function involves several components in the BladeCenter infrastructure, as illustrated in Figure 2-56.

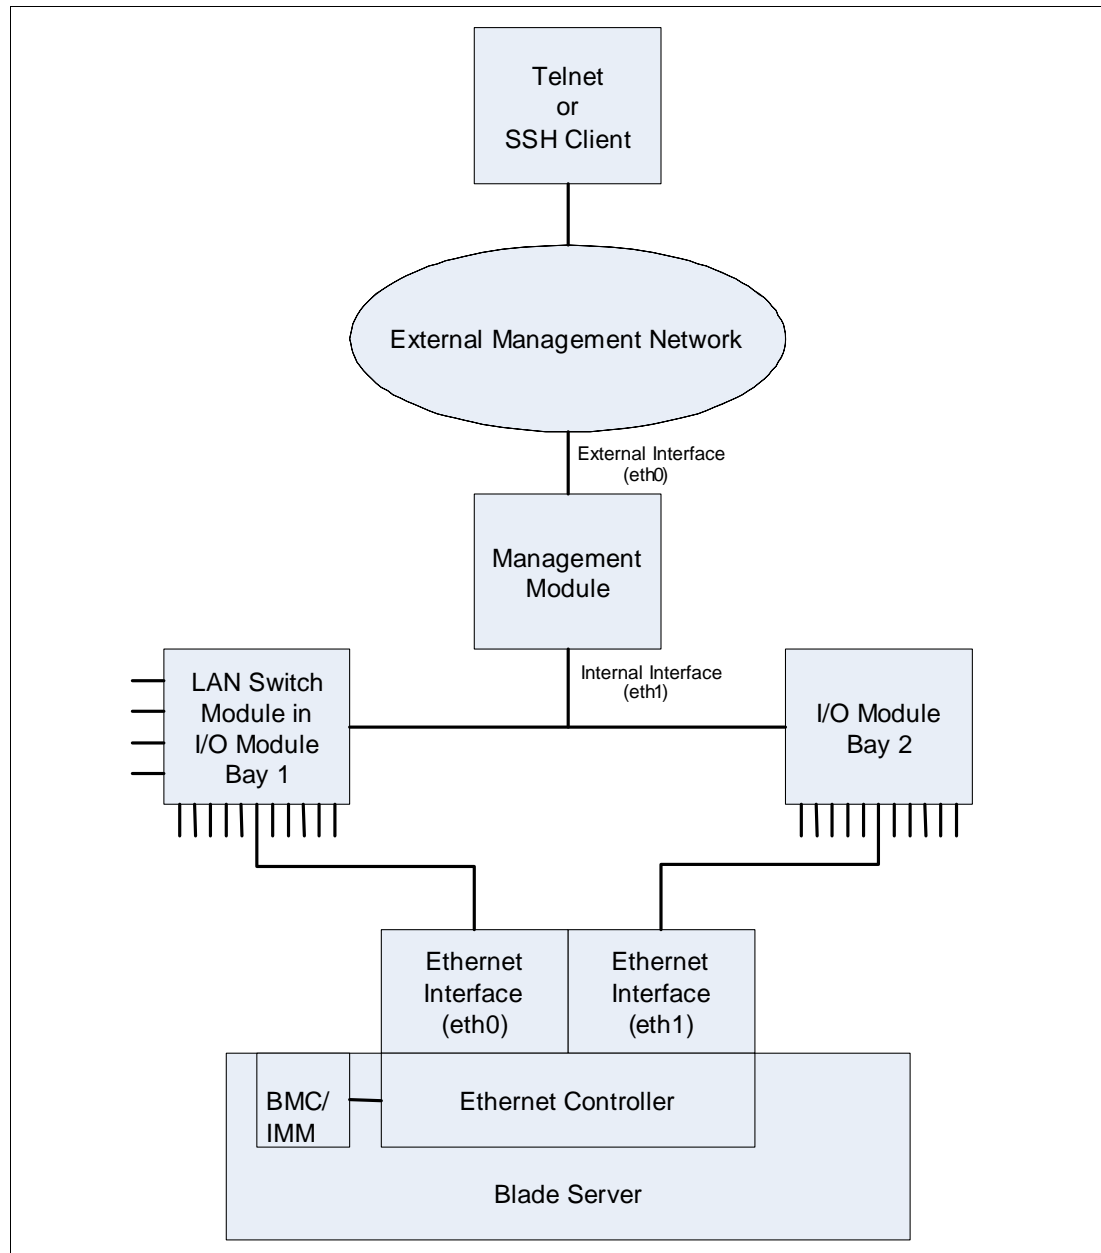


Figure 2-56 Serial over LAN components

In the BladeCenter environment, the integrated service processor (BMC or IMM) 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 configuration 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 CLI session with the Management Module. When this Telnet CLI 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, you can 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 setup allows you to have eight simultaneous SOL sessions active (one for each session, and up to eight blade servers), with 12 additional CLI sessions available for BladeCenter unit management. If security is a concern, you can use Secure Shell (SSH) sessions to establish secure Telnet CLI sessions with the BladeCenter Management Module before starting an SOL console redirect session with a blade server.

To accomplish this task, complete the following steps:

1. Using a Telnet or SSH client, connect to the BladeCenter Management Module CLI. This CLI is accessed through 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 wanted 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.

Consideration: You can have only one active SOL remote console connection to each server.

Consider the following points when you use 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 Optical Pass-thru Module (OPM) or the Copper Pass-thru Module (CPM) 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 system board 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 through the Management Module CLI. The CLI is documented in *IBM BladeCenter Management Module Command-Line Interface Reference Guide*:
<http://www.redbooks.ibm.com/abstracts/tips0756.html>
- ▶ The Management Module can support up to 20 concurrent CLI connections. This capacity 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, see *Serial over LAN (SOL) Setup Guide - IBM BladeCenter T*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-54666>



Blade server hardware configuration

This chapter describes hardware configuration of various blade servers and expansion options.

This chapter includes the following topics:

- ▶ 3.1, “Expansion cards” on page 212
- ▶ 3.2, “IBM BladeCenter HS12” on page 252
- ▶ 3.3, “IBM BladeCenter HS22” on page 260
- ▶ 3.4, “BladeCenter HS23 (E5-2600)” on page 283
- ▶ 3.5, “BladeCenter HS23 (E5-2600 v2)” on page 304
- ▶ 3.6, “BladeCenter HS23E” on page 324
- ▶ 3.7, “IBM BladeCenter HX5” on page 341
- ▶ 3.8, “IBM BladeCenter PS700, PS701, and PS702” on page 377
- ▶ 3.9, “IBM BladeCenter PS703 and PS704” on page 389
- ▶ 3.10, “Expansion blades” on page 400

The following compatibility information is provided in 1.1, “Support matrixes” on page 2:

- ▶ 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 4)
- ▶ The operating systems supported on each server (Table 1-9 on page 11)

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 to connect to two I/O modules.

The routing of the expansion card's two ports is done through the BladeCenter midplane to the I/O modules. One port from the expansion card is wired through 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 first describes the five expansion card form factors and then describes in detail each expansion card that is currently available. See 3.1.1, "Form factors" on page 213.

Table 1-3 on page 4 is a matrix that shows which expansion cards are supported in each blade server.

Ethernet expansion cards are described in the following sections:

- ▶ 3.1.1, "Form factors" on page 213
- ▶ 3.1.2, "Comparing the features of the 10 Gb Ethernet expansion cards" on page 220
- ▶ 3.1.3, "Ethernet Expansion Card (CIOv)" on page 221
- ▶ 3.1.4, "The 2/4 Port Ethernet Expansion Card (CFFh)" on page 222
- ▶ 3.1.5, "Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh)" on page 223
- ▶ 3.1.6, "Broadcom 10 Gb Gen 2 4-port Ethernet Expansion Card (CFFh)" on page 225
- ▶ 3.1.7, "Broadcom 2-port 10 Gb Virtual Fabric Adapter for IBM BladeCenter" on page 226
- ▶ 3.1.8, "Emulex 10 GbE Virtual Fabric Adapter II (CFFh)/Advanced II (CFFh)" on page 229
- ▶ 3.1.9, "Emulex 10 GbE Virtual Fabric Adapter II/Advanced II for HS23" on page 231
- ▶ 3.1.10, "QLogic 10Gb Virtual Fabric Adapter / QLogic 10Gb Virtual Fabric Converged Network Adapter (CFFh)" on page 234
- ▶ 3.1.11, "Intel 10 Gb 2-port Ethernet Expansion Card (CFFh)" on page 237

Combo Ethernet and Fibre Channel expansion cards are described in 3.1.12, "QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)" on page 238.

Fibre Channel expansion cards are described in the following sections:

- ▶ 3.1.13, "QLogic 4 Gb Fibre Channel Expansion Card (CIOv)" on page 239
- ▶ 3.1.14, "QLogic 8 Gb Fibre Channel Expansion Card (CIOv)" on page 240
- ▶ 3.1.15, "Emulex 8 Gb Fibre Channel Expansion Card (CIOv)" on page 242

Converged Network Adapters are described in the following sections:

- ▶ 3.1.16, "QLogic 2-port 10 Gb Converged Network Adapter (CFFh)" on page 243
- ▶ 3.1.17, "Brocade 2-port 10 GbE Converged Network Adapter (CFFh)" on page 244

InfiniBand expansion cards are described in the following sections:

- ▶ 3.1.18, "The 2-port 40 Gb InfiniBand Expansion Card (CFFh)" on page 246

SAS storage expansion cards are described in the following sections:

- ▶ 3.1.19, "SAS Expansion Card (CFFv)" on page 248
- ▶ 3.1.20, "SAS Connectivity Card (CIOv)" on page 249
- ▶ 3.1.21, "ServeRAID H1135 (CIOv) Controller" on page 250

3.1.1 Form factors

Current IBM blade servers either have a PCI-X connector and a PCI Express connector (HS12 blades), or two PCI Express connectors (HS22, HS22V, HX5, and PS70x blades). The expansion card form factors that these connectors support are listed in Table 3-1.

The PCI-X connector supports any of the three PCI-X form factors listed in the table. Depending on the location and shape of the PCI Express connector on the blade server, the connector either supports CIOv cards or it supports CFFh or HSFF cards.

The different form factors for expansion cards are listed here. Most cards that are currently available are the CIOv, CFFv, or CFFh form factor.

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	CFFv (vertical) ^a	PCI-X
CIOv	CIOv (vertical)	PCI Express
CFFh	CFFh (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 listed in Table 3-2.

For specific supported expansion cards and blade servers, see Table 1-3 on page 4 or the sections in this chapter for each blade.

Table 3-2 Blade support matrix for expansion cards

Blade	Type	StFF	SFF	CFFv	CIOv	CFFh	HSFF
HC10	7996	No	No	No	No	No	No
HS12	8014	No	Yes	Yes	No	Yes	Yes
HS12	8028	No	Yes	Yes	No	Yes	Yes
HS20	8678	Yes	Yes	No	No	No	No
HS20	8832	Yes	Yes	No	No	No	No
HS20	8843	Yes	Yes	No	No	No	No
HS20	7981	Yes	Yes	No	No	No	No
HS21	8853	Yes	Yes	Yes	No	Yes	Yes
HS21 XM	7995	Yes	Yes	Yes	No	Yes	Yes
HS22	7870	No	No	No	Yes	Yes	No
HS22V	7871	No	No	No	Yes	Yes	No
HS23	7875	No	No	No	Yes	Yes	No

Blade	Type	StFF	SFF	CFFv	CIOv	CFFh	HSFF
HS23E	8038	No	No	No	Yes	Yes	No
HX5	7872	No	No	No	Yes	Yes	No
HX5	7873	No	No	No	Yes	Yes	No
HS40	8839	Yes	Yes	No	No	No	No
LS20	8850	Yes	Yes	No	No	No	No
LS21	7971	Yes	Yes	Yes	No	Yes	Yes
LS22	7901	Yes	Yes	Yes	No	Yes	Yes
LS41	7972	Yes	Yes	Yes	No	Yes	Yes
LS42	7902	Yes	Yes	Yes	No	Yes	Yes
JS12	7998-60x	Yes	Yes	Yes	No	Yes	Yes
JS20	8842	Yes	No	No	No	No	No
JS21	8844	Yes	Yes	Yes	No	Yes	Yes
JS22	7998-61x	Yes	Yes	Yes	No	Yes	Yes
JS23/43	7778-23x	No	No	No	Yes	Yes	No
PS700/1/2	8406-7xx	No	No	No	Yes	Yes	No
PS703/4	7891-73X 7891-74X	No	No	No	Yes	Yes	No
QS21	0792	No	Yes	Yes	No	Yes	Yes
QS22	0793	No	Yes	Yes	No	Yes	Yes
PN41	3020	No	No	No	No	No	No

This book describes only the form factors of the cards that are currently available for ordering.

CFFv form factor

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 3-28 on page 261). A CFFv adapter requires compatible switch modules to be installed in bay 3 and bay 4 of the BladeCenter chassis.

Figure 3-1 shows the placement of a CFFv expansion card on a blade server.




Figure 3-1 Rear section of a blade server showing CFFv expansion card placement

Table 3-3 lists the available CFFv expansion card.

Table 3-3 Current CFFv expansion cards

Part number	Feature code	Description
44E5688 ^a	A3J9	SAS Expansion Card

a. Replaces 39Y9190.

CIOv form factor

A CIOv expansion card uses the PCI Express 2.0 x8 connector that is present on the HS22, HS22V, HX5, and JS70x blade servers. A CIOv adapter requires compatible switch modules to be installed in bay 3 and bay 4 of the BladeCenter chassis. The CIOv card can be used in conjunction with a CFFh card in supported servers.

Figure 3-2 shows the placement of a CIOv expansion card on an HS22 blade server.




Figure 3-2 HS22 blade server showing CIOv expansion card placement

Table 3-4 lists the available CIOv expansion cards.

Table 3-4 Current CIOv expansion cards

Part number	Feature code	Description
44W4475	1039	Ethernet Expansion Card (CIOv)
46M6065	3594	QLogic 4 Gb Fibre Channel Expansion Card (CIOv)
43W4068	1041	SAS Connectivity Card (CIOv)
90Y4570	A1XJ	ServeRAID-H1135 Controller for Flex System and BladeCenter
46M6140	3598	Emulex 8 Gb Fibre Channel Expansion Card (CIOv)
44X1945	1462	QLogic 8 Gb Fibre Channel Expansion Card (CIOv)

CFFh form factor

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 be used only 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 or CIOv adapter (see Figure 3-4 on page 219). Only one CFFh card can be installed in a single-wide blade server.

Depending on an adapter type, 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 is installed in one or several high-speed bays (bays 7 - 10).
- ▶ In the BladeCenter S, a compatible switch module is installed in bay 2.

The MSIM must contain compatible switch modules. For more information about the MSIM, see 2.11.2, "Multi-Switch Interconnect Module" on page 176.

Figure 3-3 shows the placement of a CFFh expansion card on a blade server.




Figure 3-3 Rear section of a blade server showing CFFh expansion card placement

Table 3-5 lists the expansion cards that are available in the CFFh form factor.

Table 3-5 Current CFFh expansion cards

Part number	Feature code	Description
44W4479	5476	2/4 Port Ethernet Expansion Card (CFFh)
46M6001	0056	2-port 40 Gb InfiniBand Expansion Card (CFFh)
46M6164	0098	Broadcom 4-port 10 Gb Gen 2 Ethernet Expansion Card (CFFh)
81Y3133	A1QR	Broadcom 2-port 10 Gb Virtual Fabric Adapter (CFFh)
81Y1650	3593	Brocade 2-port 10 Gb Converged Network Expansion Card (CFFh)
00Y3266	A3NV	Emulex Virtual Fabric Adapter II (CFFh)
00Y3264	A3NW	Emulex Virtual Fabric Adapter Advanced II (CFFh)
81Y3120	A287	Emulex 10 GbE Virtual Fabric Adapter II (CFFh) for HS23
90Y9332	A2ZN	Emulex 10 GbE Virtual Fabric Adapter Advanced II (CFFh) for HS23
42C1810	3593	Intel 10 Gb 2-port Ethernet Expansion Card (CFFh)
90Y3570	A1NW	Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh)
00Y3280 ^a	A3JB	QLogic 2-port 10 Gb Converged Network Adapter (CFFh)
00Y3332	A4AC	Qlogic 10Gb Virtual Fabric Adapter (CFFh)
00Y5618	A4AD	QLogic 10Gb Virtual Fabric CNA (CFFh)
00Y3270 ^b	A3JC	QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)

a. Replaces 42C1830.

b. Replaces 44X1940.

Figure 3-4 shows the CFFh expansion card installed together with a CFFv expansion card. A CIOv expansion card can also be installed together with a CFFh card at the same time. These combinations are the only combinations that can be installed at the same time.




Figure 3-4 Rear section of a blade server showing CFFh and CFFv expansion cards installed together

3.1.2 Comparing the features of the 10 Gb Ethernet expansion cards

Table 3-6 summarizes the features of the 10 Gb Ethernet expansion cards that are currently available.

Table 3-6 Comparing features of 10 Gb Ethernet expansion cards

Expansion Card →													
Feature	Integrated 10 Gb Ethernet	Broadcom 10Gb 4-pt. Eth., 46M6164	Broadcom 2-port 10 Gb VFA, 81Y3133	Brocade 2-Port 10 Gb CNA, 81Y1650	Emulex VFA II, 00Y3266	Emulex VFA Advanced II, 00Y3264	Emulex VFA II for HS23, 81Y3120	Emulex VFA Adv. II for HS23, 90Y9332	Intel 10 Gb 2-port Ethernet, 42C1810	Mellanox 2-port 10 Gb Eth., 90Y3570	QLogic 2-port 10 Gb CNA, 00Y3280	QLogic 10 Gb Virtual Fabric, 00Y3332	QLogic 10Gb VF CNA, 00Y5618
Number of physical ports	2	4	2	2	2	2	2	2	2	2	2	2	2
1 Gbps port speed	N	N	N	N	Y	Y	Y	Y	N	N	N	N	N
10 Gbps port speed	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
pNIC mode	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
vNIC Virtual Fabric mode ^a	Y	N	N	N	Y	Y	Y	Y	N	N	N	N	N
vNIC Switch Independent mode	Y	N	Y	N	Y	Y	Y	Y	N	N	N	Y	Y
TCP offload engine	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Wake on LAN	Y	Y	Y	N	Y	Y	Y	Y	Y	Y	Y	Y	Y
Serial over LAN/cKVM	Y	Y	Y	N	Y	Y	Y	Y	Y	Y	Y	Y	Y
Converged Enhanced Ethernet	Y	N	N	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
Fibre Channel over Ethernet (FCoE)	Y ^b	N	N	Y	Y ^c	Y	Y ^d	Y	N	N	Y	Y ^e	Y
iSCSI HW initiator/offload	Y ^b	N	N	N	Y ^c	Y	Y ^d	Y	N	N	N	Y ^e	Y
PXE boot	Y	Y	Y	N	Y	Y	Y	Y	Y	Y	Y	Y	Y
iSCSI BladeBoot	N	Y	N	N	N	N	N	N	Y	N	Y	N	N
FCoE boot (Boot from SAN)	Y ^b	N	N	Y	Y ^c	Y	Y ^d	Y	N	N	Y	Y ^e	Y
VLAN tagging	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Jumbo frames	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Failover	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
IBM Fabric Manager	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

a. Requires IBM Virtual Fabric 10 Gb switch module, 46C7191.

b. Requires the optional Advanced Upgrade feature, 90Y9310.

c. Requires the optional Advanced Upgrade feature, 49Y4265.

d. Requires the optional Advanced Upgrade feature, 90Y9350.

e. Requires the optional Advanced Upgrade feature, 00Y5622.

For end-to-end FCoE and iSCSI compatibility, and specific requirements and limitations for SAN Boot support, see IBM System Storage Interoperation Center (IBM SSIC) at:

<http://www.ibm.com/systems/support/storage/ssic>

3.1.3 Ethernet Expansion Card (CIOv)

The Ethernet Expansion Card (CIOv) is an Ethernet expansion card with two 1 Gb Ethernet ports designed for BladeCenter servers with CIOv expansion slots (Figure 3-5). It is based on proven Broadcom 5709S ASIC technology and offers value added features like TCP offload engine (TOE) and software-based iSCSI support for high performance. Table 3-7 lists the ordering information.

Table 3-7 Ordering part number and feature code

Description	Part number	Feature code
Ethernet Expansion Card (CIOv)	44W4475	1039

Figure 3-5 shows the Ethernet Expansion Card (CIOv).



Figure 3-5 The Ethernet Expansion Card (CIOv) for IBM BladeCenter

The expansion card has the following features:

- ▶ CIOv form factor
- ▶ PCI Express host interface
- ▶ Broadcom BCM5709S communication module
- ▶ BladeCenter Open Fabric Manager (BOFM) support
- ▶ Connection to 1000BASE-X environments using BladeCenter Ethernet switches
- ▶ Full-duplex (FDX) capability, enabling simultaneous transmission and reception of data on the Ethernet local area network (LAN)
- ▶ Failover support
- ▶ Preboot Execution Environment (PXE) support
- ▶ Support for direct memory access (DMA)
- ▶ Wake on LAN support for both ports

The Ethernet Expansion Card (CIOv) requires a supported I/O module installed in bay 3 and bay 4 of the chassis in which the cards and servers are installed.

For more information, see the following publications:

- ▶ *Ethernet Expansion Card (CIOv) for IBM BladeCenter*.
<http://www.redbooks.ibm.com/abstracts/tips0697.html>
- ▶ *Ethernet Expansion Card (CIOv) Installation and User's Guide*.
<http://ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-080268>

3.1.4 The 2/4 Port Ethernet Expansion Card (CFFh)

The 2/4 Port Ethernet Expansion Card (CFFh) is based on the Broadcom 5709 Gigabit Ethernet chipset. It enables two ports in the BladeCenter S Chassis. It enables four ports in the BladeCenter H and HT chassis.

Table 3-8 lists the ordering information.

Table 3-8 Ordering part number and feature code

Description	Part number	Feature code
2/4 Port Ethernet Expansion Card (CFFh)	44W4479	5476

The 2/4 Port Ethernet Expansion Card (CFFh), shown in Figure 3-6 on page 223, features:

- ▶ CFFh form factor
- ▶ 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 (CFFh) is used in conjunction with the Multi-Switch Interconnect Module (MSIM). This configuration requires that Gigabit Ethernet Switch Modules be installed in the MSIM. 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 needs to be installed in bay 2 of the BladeCenter S chassis.

Figure 3-6 shows the 2/4 Port Ethernet Expansion Card (CFFh).



Figure 3-6 2/4 Port Ethernet Expansion Card (CFFh)

For more information, see the following documents:

- *Installation Guide for 2/4 Port Ethernet Expansion Card (CFFh) for IBM BladeCenter.*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5082171>
- *2/4 Port Ethernet Expansion Card (CFFh) for IBM BladeCenter.*
<http://www.redbooks.ibm.com/abstracts/tips0698.html>

3.1.5 Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh)

The Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh) for IBM BladeCenter delivers the industry's lowest latency performance of 10 Gigabit Ethernet connectivity with stateless offloads for converged fabrics in enterprise data centers, high-performance computing, and embedded environments.

The adapter is based on Mellanox ConnectX-2 EN technology, which improves network performance by increasing available bandwidth to the processor, especially in virtualized server environments.

Table 3-9 lists the ordering information.

Table 3-9 Ordering part number and feature code

Description	Part number	Feature code
Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh)	90Y3570	A1NW

The Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh) has the following specifications:

- ▶ CFFh form-factor
- ▶ Two full-duplex 10 Gb Ethernet ports
- ▶ PCI Express x8 Gen 2 host interface
- ▶ Mellanox ConnectX-2 EN ASIC
- ▶ Connectivity to high-speed I/O module bays in BladeCenter H and HT chassis
- ▶ Power consumption: 6.4 W (typical)
- ▶ BladeCenter Open Fabric Manager support
- ▶ PXE v2.1 boot support
- ▶ Wake on LAN support (first port only)
- ▶ Serial over LAN and cKVM support

The Ethernet support and standards are as follows:

- ▶ IEEE Std 802.3ae 10 Gigabit Ethernet
- ▶ IEEE Std 802.3ad Link Aggregation and Failover
- ▶ IEEE Std 802.3x Pause
- ▶ IEEE 802.1Q VLAN tags
- ▶ IEEE 802.1p QoS/CoS
- ▶ Multicast
- ▶ Jumbo frame support (10 KB)
- ▶ 128 MAC/VLAN addresses per port

Data Center Bridging (DCB) support is as follows:

- ▶ Enhanced Transmission Selection (ETS) (IEEE 802.1Qaz)
- ▶ Quantized Congestion Notification (QCN) (IEEE 802.1Qau)
- ▶ Priority-based Flow Control (PFC) (IEEE 802.1Qbb)
- ▶ Data Center Bridging Capability eXchange (DCBX) protocol

TCP/UDP/IP stateless offload is as follows:

- ▶ TCP/UDP/IP checksum offload
- ▶ TCP Large Send Offload (< 64 KB) or Giant Send Offload (64 KB - 16 MB) for segmentation
- ▶ Receive Side Scaling (RSS) up to 32 queues
- ▶ Line rate packet filtering

Additional processor offloads are as follows:

- ▶ RDMA over Converged Ethernet support
- ▶ FC checksum offload
- ▶ VMDirect Path support
- ▶ Traffic steering across multiple cores
- ▶ Intelligent interrupt coalescence
- ▶ Compliant with Microsoft RSS and NetDMA

Hardware-based I/O virtualization is as follows:

- ▶ Address translation and protection
- ▶ Dedicated adapter resources and guaranteed isolation
- ▶ Multiple queues per virtual machine
- ▶ Hardware switching between guest OSes
- ▶ Enhanced QoS
- ▶ VMware NetQueue support

Storage support: Compliant with T11.3 FC-BB-5 FCoE frame format

The expansion card is shown in Figure 3-7.




Figure 3-7 Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh)

For more information, see *Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh)*:

<http://www.redbooks.ibm.com/abstracts/tips0820.html>

3.1.6 Broadcom 10 Gb Gen 2 4-port Ethernet Expansion Card (CFFh)

The Broadcom 10 Gb Gen 2 4-port Ethernet Expansion Card is based on the Broadcom 57711 module, which has support for a PCI Express 2.0 (Gen 2) x8 host interface. It has four 10 Gigabit Ethernet interfaces. It is a high-speed expansion card especially designed to interface with a 10 Gb Ethernet Switch Module.

Table 3-10 lists the ordering information.

Table 3-10 Ordering part number and feature code

Description	Part number	Feature code
Broadcom 10 Gb Gen 2 4-port Ethernet Expansion Card (CFFh)	46M6164	0098

The expansion card has the following features and benefits:

- ▶ Offers an end-to-end solution up to 10 Gb
- ▶ Enables four 10 Gb ports from the blade server to the external network
- ▶ Is based on the Broadcom 57711 module
- ▶ Has two PCI Express 2.0 x8 host interfaces for high-speed connection
- ▶ Has connectivity to high-speed I/O module bays in the BladeCenter H and BladeCenter HT chassis

- ▶ Supports failover
- ▶ Supports IBM BladeCenter Open Fabric Manager
- ▶ Supports iSCSI BladeBoot
- ▶ Supports Wake-on-LAN
- ▶ Has support for SOL and cKVM over the high-speed network

Performance features are as follows:

- ▶ TCP offload engine (TOE)
- ▶ Full fast-path TCP offload
- ▶ TCP/IP checksum offload
- ▶ TCP/IP segmentation offload
- ▶ Improved performance over the Broadcom BCM57710 module

This expansion card can be combined with a CFFv or CIOv expansion card on the same blade server. It connects to two 10 Gb Ethernet Switch Modules in high speed switch bays 7, 8, 9, and 10.

The expansion card is shown in Figure 3-8.



Figure 3-8 Broadcom 10 Gb Gen 2 4-port Ethernet Expansion Card (CFFh)

For more information, see *Broadcom 10 Gb Gen 2 2-Port and 4-Port Ethernet Expansion Cards*:

<http://www.redbooks.ibm.com/abstracts/tips0728.html>

3.1.7 Broadcom 2-port 10 Gb Virtual Fabric Adapter for IBM BladeCenter

The Broadcom 2-port 10 Gb Virtual Fabric Adapter for IBM BladeCenter is a part of the IBM comprehensive networking portfolio of 1 GbE and 10 GbE adapters. It delivers industry-leading performance and scalability per watt, reducing requirements for power and cooling.

The Broadcom 2-port 10 Gb Virtual Fabric Adapter can be configured to operate in virtual NIC mode, where up to eight virtual NIC devices can be enabled. Each physical 10 Gbps port can be divided into four virtual ports with bandwidth allocation in 100 Mbps increments to the maximum 10 Gbps per physical port.

Table 3-11 lists the ordering information.

Table 3-11 Ordering part number and feature code

Description	Part number	Feature code
Broadcom 2-port 10 Gb Virtual Fabric Adapter or IBM BladeCenter	81Y3133	A1QR

The expansion card has the following features:

- ▶ Offers an end-to-end solution up to 10 Gb
- ▶ One PCI Express 2.0 x8 host interface for high-speed connection
- ▶ Connectivity to high-speed I/O module bays in the BladeCenter H and BladeCenter HT chassis
- ▶ Enables two 10 Gb ports from the blade server to the external network
- ▶ Based on the Broadcom 10 Gb 57712 controller
- ▶ Hardware parity, CRC, ECC, and other advanced error checking
- ▶ VLAN insertion and extraction
- ▶ Support for jumbo frames up to 9000 bytes
- ▶ UEFI preboot eXecutive Environment (PXE) 2.0 boot support
- ▶ Supports load balancing and failover
- ▶ Supports IBM BladeCenter Open Fabric Manager (vNIC mode and pNIC mode)
- ▶ Supports Wake-on-LAN
- ▶ Supports SOL and cKVM

Virtual NIC (vNIC) support is as follows:

- ▶ Supports IBM Virtual Fabric Switch in Switch Independent Mode with eight virtual ports (four per physical port)
- ▶ uEFI based vNIC configuration
- ▶ vNIC port-specific quality of service configuration
- ▶ Supports vNIC capability, all included in base function; no firmware key is required

The performance features are as follows:

- ▶ TCP stateful offload (Microsoft Chimney / TOE)
- ▶ TCP and UDP stateless offloads over IPv4 and IPv6
- ▶ Interrupt coalescing
- ▶ Message Signal Interrupts (MSI, MSI-X)
- ▶ RX/TX multiqueue
- ▶ VMware NetQueue and Microsoft VMQ

The expansion card has two modes of operation:

► Switch Independent vNIC mode

In vNIC mode, each physical port appears to the blade server as four virtual NICs. The default bandwidth for each vNIC is 2.5 Gbps. Bandwidth for each vNIC can be configured from 100 Mbps to 10 Gbps, up to a maximum of 10 Gb per virtual port.

The total bandwidth of the four vNICs per port cannot exceed 10 Gbps. However, the Broadcom adapter supports oversubscription, where the aggregate bandwidth can be as high as 40 Gbps (4x10 Gbps). This support allows any of the vNICs to consume as much as 10 Gbps where there is no contention or less when there is contention (based on the weight assigned to that vNIC). vNICs can also be configured to have zero bandwidth if you must allocate the available bandwidth to fewer than eight vNICs.

The adapter works with the following switches:

- IBM Virtual Fabric 10 Gb Switch module
- Cisco Nexus 4001I Switch Module

These switches are installed in I/O module bays 7 and 9 in the chassis. Switch Independent Mode extends the existing customer VLANs to the virtual NIC interfaces. The IEEE 802.1Q VLAN tag is essential to the separation of the vNIC groups by the NIC adapter or driver and the switch. The VLAN tags are added to the packet by the applications or drivers at each endstation rather than by the switch.

Important: The Broadcom 2-port 10 Gb Virtual Fabric Adapter does not support IBM Virtual Fabric Mode.

► Physical NIC (pNIC) mode

In pNIC mode, the expansion card can operate as a standard 10 Gbps or 1 Gbps 2-port Ethernet expansion card. When in pNIC mode, the expansion card functions with any supported 10 Gb switch or 10 Gb Pass-thru Module installed in I/O module bays 7 and 9.

The expansion card is shown in Figure 3-9.



Figure 3-9 Broadcom 2-port 10 Gb Virtual Fabric Adapter for IBM BladeCenter

For more information, see *Broadcom 2-port 10 Gb Virtual Fabric Adapter*:

<http://www.redbooks.ibm.com/abstracts/tips0827.html>

3.1.8 Emulex 10 GbE Virtual Fabric Adapter II (CFFh)/Advanced II (CFFh)

The Emulex 10 GbE Virtual Fabric Adapter II and Emulex 10 GbE Virtual Fabric Adapter Advanced II are enhancements to the existing IBM BladeCenter Virtual Fabric portfolio. These adapters allow IBM clients to simplify their I/O infrastructure by reducing the number of switches needed inside the chassis. The Emulex 10 GbE Virtual Fabric Adapter II is a dual-port 10 Gb Ethernet card that supports 1 Gbps or 10 Gbps traffic, or up to eight virtual NIC (vNIC) devices.

The virtual NICs are configured to meet the client's mix of network connectivity and throughput demands for today's complex server application environments. Each physical 10 Gbps port can be divided into four virtual ports with bandwidth allocation in 100 Mbps increments to the maximum 10 Gbps per physical port. The Emulex 10 GbE Virtual Fabric Adapter II Advanced adds FCoE and iSCSI hardware initiator functions.

Table 3-12 lists the Emulex 10 GbE Virtual Fabric Adapter II (CFFh) and Emulex 10 GbE Virtual Fabric Adapter Advanced II (CFFh) ordering part numbers and feature codes.

Table 3-12 Adapter ordering part numbers and feature codes

Description	Part number	Feature code
Emulex 10 GbE Virtual Fabric Adapter II (CFFh)	00Y3266	A3NV
Emulex 10 GbE Virtual Fabric Adapter Advanced II (CFFh)	00Y3264	A3NW
Emulex 10 GbE Virtual Fabric Advanced Upgrade	49Y4265	2436

Emulex Virtual Fabric Adapter II and Emulex Virtual Fabric Adapter Advanced II have these features:

- ▶ Connection to either 1 Gb or 10 Gb data center infrastructure (1 Gb and 10 Gb auto-negotiation).
- ▶ PCI Express 2.0 x8 host interface.
- ▶ IBM BladeCenter CFFh form factor. Can be combined with a CIOv expansion card on the same blade server.
- ▶ It operates either as an 8-port virtual NIC or as a 2-port 1/10 Gb Ethernet adapter.
- ▶ Virtual port bandwidth allocation in 100 Mbps increments.
- ▶ Supports up to eight virtual ports.
- ▶ Wake On LAN support.
- ▶ Full-duplex (FDX) capability.
- ▶ Bus-mastering support.
- ▶ Direct memory access (DMA) support.
- ▶ Preboot Execution Environment (PXE) support.
- ▶ IPv4/IPv6 TCP, UDP checksum offload.
- ▶ Large send offload (LSO).
- ▶ Large receive offload.
- ▶ Receive side scaling (RSS).
- ▶ IPv4 TCP Chimney Offload.
- ▶ TCP Segmentation Offload.
- ▶ VLAN insertion and extraction.

- ▶ Jumbo frames up to 9000 bytes.
- ▶ Load balancing and failover support, including adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), teaming support, and IEEE 802.3ad.
- ▶ Enhanced Ethernet (draft):
 - Enhanced Transmission Selection (ETS) (P802.1Qaz)
 - Priority-based Flow Control (PFC) (P802.1Qbb)
 - Data Center Bridging Capabilities eXchange Protocol, CIN-DCBX and CEE-DCBX (P802.1Qaz)
- ▶ The Emulex Virtual Fabric Adapter Advanced II provides FCoE and iSCSI hardware initiator support.
- ▶ Supports Serial over LAN (SoL) and concurrent KVM (cKVM).
- ▶ Two vNIC modes: IBM Virtual Fabric Mode (vNIC1) and Switch Independent Mode (vNIC2).

The Emulex 10 GbE Virtual Fabric Adapter II now offers two vNIC modes of operation in addition to pNIC mode:

- ▶ IBM Virtual Fabric Mode (vNIC1) works only in conjunction with the IBM Virtual Fabric 10 Gb Switch Module installed in I/O module bays 7 and 9 in the chassis. In this mode, the Emulex adapter communicates with the switch module to obtain vNIC parameters (using DCBX). Also, a special tag within each data packet is added and later removed by the NIC and switch for each vNIC group to maintain separation of the virtual channels.

In vNIC mode, each physical port is divided into four virtual ports, providing a total of eight virtual NICs per adapter. The default bandwidth for each vNIC is 2.5 Gbps. Bandwidth for each vNIC can be configured at the IBM Virtual Fabric 10 Gb Switch Module, from 100 Mbps to 10 Gbps, up to a total of 10 Gb per physical port. The vNICs can also be configured to have 0 bandwidth if you must allocate the available bandwidth to fewer than eight vNICs. In IBM Virtual Fabric Mode, you can change the bandwidth allocations through the switch user interfaces without requiring a reboot of the server.

When storage protocols are enabled on the Emulex 10 GbE Virtual Fabric Adapter Advanced II, six ports are Ethernet, and two ports are either iSCSI or FCoE

- ▶ Switch Independent Mode (vNIC2), where the adapter works with the following switches:
 - IBM Virtual Fabric 10 Gb Switch Module
 - Cisco Nexus 4001I Switch Module
 - 10 Gb Ethernet Pass-thru Module and a top-of-rack (TOR) switch

These switches must be installed in I/O module bays 7 and 9 in the chassis. Beyond the way that the adapter communicates with the switch, Switch Independent Mode offers the same capabilities as IBM Virtual Fabric Mode in terms of the number of vNICs and the bandwidth each vNIC can be configured to have. Switch Independent Mode extends the existing customer VLANs to the virtual NIC interfaces. The IEEE 802.1Q VLAN tag is essential to the separation of the vNIC groups by the NIC adapter or driver and the switch. The VLAN tags are added to the packet by the applications or drivers at each endstation rather than by the switch.

- ▶ In pNIC mode, the expansion card can operate as a standard 10 Gbps or 1 Gbps 2-port Ethernet expansion card. When in pNIC mode, the expansion card functions with any available 10 Gb switch or 10 Gb pass-through module installed in I/O module bays 7 and 9.

In pNIC mode, the Advanced adapter or the Advanced upgrade operates in traditional Converged Network Adapter (CNA) mode with two ports of Ethernet and two ports of storage (iSCSI or FCoE) available to the operating system.

Figure 3-10 shows the Emulex 10 GbE Virtual Fabric Adapter II.



Figure 3-10 Emulex 10 GbE Virtual Fabric Adapter II (CFFh)

For more information, see *Emulex 10 GbE Virtual Fabric Adapter II (CFFh)*:

<http://www.redbooks.ibm.com/abstracts/tips0828.html>

3.1.9 Emulex 10 GbE Virtual Fabric Adapter II/Advanced II for HS23

The Emulex 10 GbE Virtual Fabric Adapter II (VFA II) and Emulex 10 GbE VFA Advanced II for IBM BladeCenter HS23 are the new options available to the existing IBM BladeCenter Virtual Fabric portfolio. These adapters are supported on the HS23 blade to enable up to four ports for increased I/O bandwidth and maximum performance. The combination of HS23 and Emulex options enables clients to simplify their I/O infrastructure by reducing the number of switches needed inside the chassis while supporting Ethernet and virtual NICs (vNICs) using the same hardware components.

The virtual NICs are configured to meet the client's mix of network connectivity and throughput demands for today's complex server application environments. Both physical 10 Gbps ports on the adapter can be divided into four virtual ports with bandwidth allocation in 100 Mbps increments up to the maximum of 10 Gbps per physical port. The Emulex 10 GbE VFA Advanced II adds FCoE and iSCSI hardware initiator functionality to enable data center infrastructure convergence.

The Emulex 10 GbE VFA II routes onboard 10 Gb Ethernet ports on HS23 to I/O bays 7 and 9 of BladeCenter H or HT chassis, and the adapter's 10 GbE ports are routed to the I/O bays 8 and 10. With four 10 Gb Ethernet ports, the solution offers the flexibility to operate as a Virtual Fabric with up to 14 vNICs or as traditional 10 Gb Ethernet NICs. The adapter enables you to configure up to 18 ports on a single blade server: eight virtual NICs with the Emulex 10 GbE VFA II, plus six vNICs with the onboard Emulex 10 GbE NICs, plus the two onboard Gigabit Ethernet NICs, plus two GbE NICs with a CIOv expansion card.

Table 3-13 lists the Emulex 10 GbE Virtual Fabric Adapter II (CFFh) and Emulex 10 GbE Virtual Fabric Adapter Advanced II (CFFh) for HS23 ordering information.

Table 3-13 Adapter ordering part numbers and feature codes

Description	Part number	Feature code
Emulex 10 GbE VFA II for IBM BladeCenter HS23	81Y3120	A287
Emulex 10 GbE VFA Advanced II for IBM BladeCenter HS23	90Y9332	A2ZN
Virtual Fabric Advanced FOD Upgrade for IBM BladeCenter	90Y9350	A2ZP

Emulex Virtual Fabric Adapter II for HS23 and Emulex Virtual Fabric Adapter Advanced II for HS23 have the following features:

- ▶ Two 10 Gb Ethernet ports (1 Gb and 10 Gb auto-negotiation).
- ▶ Based on Emulex BE3 chip - 81Y3120 and 90Y9332.
- ▶ Based on Emulex BE3R chip - 00Y3298 and 00Y3302.
- ▶ PCIe x8 Gen 2 host interface.
- ▶ CFFh form factor can be combined with a CIOv expansion card on the same blade server.
- ▶ BladeCenter HS23 onboard dual-port 10 GbE NIC pass-through capability
- ▶ Full-duplex (FDX) capability (40 Gbps throughput).
- ▶ 10 Gb ports operating either in a virtual NIC (vNIC) or physical NIC (pNIC) mode:
 - vNIC mode: Up to eight vNICs (up to four vNICs per one 10 Gb port):
 - Virtual Fabric or Switch Independent mode of operation.
 - Virtual port bandwidth allocation in 100 Mbps increments.
 - Up to two vNICs can be configured as an iSCSI or FCoE vNICs (one per port) with optional Advanced Upgrade (90Y9350) or with Emulex 10 GbE VFA Advanced II (90Y9332 or 00Y3302).
 - pNIC mode: Dual-port 1/10 Gb Ethernet adapter

With optional Advanced Upgrade (90Y9350) or with Emulex 10 GbE VFA Advanced II (90Y9332 or 00Y3302), the pNIC mode enables two LAN (10 GbE) ports and two storage ports (iSCSI or FCoE).
- ▶ IPv4/IPv6 offload:
 - TCP, UDP checksum offload
 - Large send offload (LSO)
 - Large receive offload (LRO)
 - Receive side scaling (RSS)
- ▶ IPv4 TCP Chimney Offload.
- ▶ IEEE 802.1Q VLAN tagging.
- ▶ VLAN insertion and extraction.
- ▶ Jumbo frames up to 9000 bytes.
- ▶ Load balancing and failover teaming support, including adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), and IEEE 802.3ad.
- ▶ Enhanced Ethernet (draft) support:
 - Enhanced Transmission Selection (ETS) (P802.1Qaz)
 - Priority-based Flow Control (PFC) (P802.1Qbb)
 - Data Center Bridging eXchange Protocol (DCBX) (P802.1Qaz)

- ▶ Supports Serial over LAN (SoL) and concurrent KVM (cKVM).
- ▶ Supports Preboot Execution Environment (PXE).
- ▶ Supports Wake On LAN.
- ▶ Supports Message Signal Interrupt (MSI-X).

The Emulex 10 GbE Virtual Fabric Adapter II for HS23 offers two vNIC modes of operation in addition to pNIC mode:

- ▶ IBM Virtual Fabric Mode works only in conjunction with IBM Virtual Fabric 10 Gb Switch Module installed in I/O module bays 8 and 10 in the chassis. In this mode, the Emulex adapter communicated with the switch module to obtain vNIC parameters (using DCBX). Also, a special tag within each data packet is added and later removed by the NIC and switch for each vNIC group to maintain separation of the virtual channels.

In vNIC mode, each physical port is divided into four virtual ports, providing a total of eight virtual NICs per adapter. The default bandwidth for each vNIC is 2.5 Gbps. Bandwidth for each vNIC can be configured at the IBM Virtual Fabric 10 Gb Switch Module from 100 Mbps to 10 Gbps, up to a total of 10 Gb per physical port. The vNICs can also be configured to have 0 bandwidth if you must allocate the available bandwidth to fewer than eight vNICs. In IBM Virtual Fabric Mode, you can change the bandwidth allocations through the switch user interfaces without requiring a reboot of the server.

When storage protocols are enabled on the Emulex 10 GbE Virtual Fabric Adapter Advanced II for HS23, six ports are Ethernet, and two ports are either iSCSI or FCoE

- ▶ Switch Independent Mode, where the adapter works with the following I/O modules:
 - IBM Virtual Fabric 10 Gb Switch Module
 - Cisco Nexus 4001I Switch Module
 - 10 Gb Ethernet Pass-thru Module and a top-of-rack (TOR) switch

These modules must be installed in I/O module bays 8 and 10 in the chassis. Beyond the way that the adapter communicates with the switch, Switch Independent Mode offers the same capabilities as IBM Virtual Fabric Mode in terms of the number of vNICs and the bandwidth each vNIC can be configured to have. Switch Independent Mode extends the existing customer VLANs to the virtual NIC interfaces. The IEEE 802.1Q VLAN tag is essential to the separation of the vNIC groups by the NIC adapter or driver and the switch. The VLAN tags are added to the packet by the applications or drivers at each endstation rather than by the switch.

- ▶ In pNIC mode, the expansion card can operate as a standard 10 Gbps or 1 Gbps 2-port Ethernet expansion card. When in pNIC mode, the expansion card functions with the following I/O modules:
 - IBM Virtual Fabric 10 Gb Switch Module
 - Cisco Nexus 4001I Switch Module
 - 10 Gb Ethernet Pass-thru Module and a top-of-rack (TOR) switch

These modules must be installed in I/O module bays 8 and 10.

In pNIC mode, the Advanced adapter or the Advanced upgrade operates in traditional Converged Network Adapter (CNA) mode with two ports of Ethernet and two ports of storage (iSCSI or FCoE) available to the operating system.

Figure 3-10 on page 231 shows the Emulex 10 GbE Virtual Fabric Adapter II for HS23.



Figure 3-11 Emulex 10 GbE VFA II for IBM BladeCenter HS23 (CFFh)

For more information, see *Emulex 10 GbE VFA II for IBM BladeCenter HS23*:

<http://www.redbooks.ibm.com/abstracts/tips0875.html>

3.1.10 QLogic 10Gb Virtual Fabric Adapter / QLogic 10Gb Virtual Fabric Converged Network Adapter (CFFh)

QLogic Virtual Fabric Adapter and Qlogic Virtual Fabric CNA are based on the QLogic 8200 VFA family of Converged Network Controllers, supporting simultaneous LAN (TCP/IP) and SAN (FCoE and iSCSI) traffic at line-rate, 10 Gbps Ethernet speed. This convergence of networking traffic lowers data center costs by eliminating the need for separate LAN and SAN infrastructure.

Table 3-14 lists the ordering information.

Table 3-14 Ordering part number and feature code

Description	Part number	Feature code
QLogic 10Gb Virtual Fabric Adapter (CFFh)	00Y3332	A4AC
QLogic 10 Gb Virtual Fabric CNA (CFFh)	00Y5618	A4AD
Qlogic Virtual Fabric Advanced FoD Upgrade	00Y5622	A4AE

The QLogic Virtual Fabric Adapter (CFFh) can be upgraded to the same features as the QLogic Virtual Fabric CNA (CFFh) with the addition of the Qlogic Virtual Fabric Advanced FoD Upgrade option. Table 3-14 lists the upgrade part number.

Figure 3-12 shows the QLogic Virtual Fabric Adapter (CFFh)/QLogic Virtual Fabric CNA (CFFh).



Figure 3-12 The QLogic Virtual Fabric Adapter / Virtual Fabric CNA (CFFh) for IBM BladeCenter

The expansion cards have the following features:

- ▶ 10 Gbps full-duplex line rate per-port throughput
- ▶ Combo Form Factor (CFFh) PCI Express 2.0 x8 adapter
- ▶ Communication module: QLogic 8200 ASIC
- ▶ 10 Gb ports operating either in a virtual NIC (vNIC) or physical NIC (pNIC) mode:
 - vNIC mode: Up to eight vNICs (up to four vNICs per one 10 Gb port):
 - Switch Independent mode of operation.
 - Virtual port bandwidth allocation in 100 Mbps increments.
 - Up to four vNICs can be configured as a NIC, iSCSI, or FCoE vNICs (up to four NIC vNICs, one FCoE vNIC, and one iSCSI vNIC per port) with optional Advanced Upgrade (00Y5622) or with QLogic 10Gb Virtual Fabric CNA (00Y5618).
 - pNIC mode: dual-port 10 Gb Ethernet adapter
 - With optional Advanced Upgrade (00Y5622) or with QLogic 10Gb Virtual Fabric CNA (00Y5618), the pNIC mode enables two LAN (10 GbE) ports and four storage ports (2 iSCSI and 2 FCoE).
- ▶ Serial over LAN (SoL) and concurrent KVM (cKVM) support
- ▶ IBM Fabric Manager support
- ▶ Preboot Execution Environment (PXE) support

Switch Independent vNIC mode is also known as NIC Partitioning or NPAR, where each physical port appears to the server as four virtual NICs. Bandwidth for each vNIC can be configured from 100 Mbps to 10 Gbps. This vNIC mode extends the existing VLANs to the virtual NIC interfaces. It uses IEEE 802.1Q VLAN tags to separate data traffic; the VLAN tags are added to the packet by the applications or drivers rather than by the switch.

Note: The adapter supports only fixed 10 GbE connectivity (no auto-negotiation); 1 Gb Ethernet connectivity is not supported.

Stateless offload features include the following items:

- ▶ IP, TCP, and UDP checksum offloads
- ▶ Large and giant send offload (LSO and GSO)
- ▶ Large receive offload (LRO)
- ▶ Receive Side Scaling (RSS)
- ▶ Header-data split
- ▶ Interrupt coalescing
- ▶ VMware NetQueue and Microsoft VMQ

Enhanced Ethernet compliance is as follows:

- ▶ 802.1Qbb (Priority-based flow control)
- ▶ 802.1Qaz (Enhanced transmission selection)
- ▶ 802.1Qaz (DCBX protocol)

Protocols are as follows:

- ▶ 802.3ae (10 Gb Ethernet)
- ▶ 802.1q (VLAN)
- ▶ 802.3ad (Link Aggregation)
- ▶ 802.1p (Priority Encoding)
- ▶ 802.3x (Flow Control)
- ▶ IPv4 Specification (RFC 791)
- ▶ IPv6 Specification (RFC 2460)

FCoE compliance is as follows:

- ▶ Logins: Support for 2,048 concurrent logins and 2,048 active exchanges
- ▶ N_Port ID virtualization (NPIV)
- ▶ SCSI-3 Fibre Channel Protocol (SCSI-FCP)
- ▶ Fibre Channel Tape (FC-TAPE) Profile
- ▶ SCSI Fibre Channel Protocol-2 (FCP-2)
- ▶ Second Generation FC Generic Services (FC-GS-2)
- ▶ Third Generation FC Generic Services (FC-GS-3)
- ▶ FCoE and FIP (FC-BB-5)

iSCSI compliance is as follows:

- ▶ RFC 3347 (iSCSI Requirements and Design Considerations)
- ▶ CHAP
- ▶ iSNS
- ▶ SLP

Tools and utilities are as follows:

- ▶ Management tools and device utilities:
 - QConvergeConsole: A unified management tool (GUI and CLI) for Fibre Channel/FCoE, iSCSI, and networking
 - Native OS management tools for networking
- ▶ Boot support: pre-execution environment (PXE), FCoE, and iSCSI boot
- ▶ APIs: SNIA HBA API V2, SMI-S

For more information, see the *QLogic 10Gb Virtual Fabric Adapter and CNA Product Guide*:

<http://www.redbooks.ibm.com/abstracts/tips1025.html>

3.1.11 Intel 10 Gb 2-port Ethernet Expansion Card (CFFh)

The Intel 10 Gb 2-port Ethernet Expansion Card for IBM BladeCenter is a dual-port 10 Gb Ethernet card that is powered by an Intel Ethernet 82599 dual-port 10 Gigabit Ethernet Controller.

Table 3-15 lists the Intel 10 Gb 2-port Ethernet Expansion Card (CFFh) ordering information.

Table 3-15 Ordering part number and feature code

Description	Part number	Feature code
Intel 10 Gb 2-port Ethernet Expansion Card (CFFh)	42C1810	3593

The expansion card has the following features:

- ▶ CFFh form-factor
- ▶ PCI Express x8 Gen 2 host interface
- ▶ Intel 82599 ASIC
- ▶ Two full-duplex 10 Gb Ethernet ports that operate in standard physical NIC (pNIC) mode
- ▶ Connectivity to high-speed I/O module bays in BladeCenter H and BladeCenter HT chassis
- ▶ Wake on LAN support
- ▶ Serial over LAN/cKVM support
- ▶ PXE boot support
- ▶ iSCSI BladeBoot (built-in software iSCSI initiator) support
- ▶ VLAN tagging support
- ▶ Jumbo frames support
- ▶ Failover support
- ▶ IBM Fabric Manager support

Performance features are as follows:

- ▶ I/O advanced technology (I/OAT) acceleration
- ▶ Virtual Machine Device Queues (VMDq)
- ▶ TCP/IP checksum offload
- ▶ TCP/IP segmentation offload

Intel I/O Acceleration Technology (Intel I/OAT) is a suite of features that improves data acceleration across the platform, from networking devices to the chipset and processors, which help to improve system performance and application response times:

- ▶ Intel QuickData
- ▶ Direct Cache Access (DCA)
- ▶ MSI-X
- ▶ Low-Latency Interrupts
- ▶ Receive Side Scaling (RSS)

Virtual Machine Device Queues (VMDq) reduces I/O impact on the hypervisor in a virtualized server by performing data sorting and coalescing in the network silicon. VMDq technology uses multiple queues in the network controller. With VMDq, as data packets enter the network adapter, they are sorted, and packets going to the same destination (or virtual machine) are grouped together in a single queue. The packets are then sent to the hypervisor, which directs them to their respective virtual machines. Relieving the hypervisor of packet filtering and sorting improves overall processor usage and throughput levels.

Figure 3-13 shows the Intel 10 Gb 2-port Ethernet Expansion Card.



Figure 3-13 Intel 10 Gb 2-port Ethernet Expansion Card (CFFh)

For more information, see *Intel 10 Gb 2-port Ethernet Expansion Card (CFFh)*:

<http://www.redbooks.ibm.com/abstracts/tips0764.html>

3.1.12 QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)

This card 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.

Table 3-16 lists the ordering information.

Table 3-16 Ordering part number and feature code

Description	Part number	Feature code
QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)	00Y3270 ^a	A3JC

a. Replaces 44X1940.

The QLogic Ethernet and 8 Gb FC Expansion Card, shown in Figure 3-14, is used in conjunction with the Multi-Switch Interconnect Module (MSIM), and requires, in the MSIM, either or both of the following modules:

- ▶ 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)



Figure 3-14 QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)

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 want to enable the Fibre Channel connections of the expansion card, you do not need to install an FC switch module in the MSIM.

For more information, see the following publications:

- ▶ *QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh) for IBM BladeCenter Installation and User's Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5080514>
- ▶ *QLogic Ethernet and 8 Gb Fibre Channel Expansion Card:*
<http://www.redbooks.ibm.com/abstracts/tips0690.html>

3.1.13 QLogic 4 Gb Fibre Channel Expansion Card (CIOv)

The QLogic 4 Gb Fibre Channel Expansion Card (CIOv) for IBM BladeCenter enables you to quickly connect the BladeCenter servers with CIOv expansion slots to a Fibre Channel SAN.

Table 3-17 lists the ordering information.

Table 3-17 Ordering part number and feature code

Description	Part number	Feature code
QLogic 4 Gb Fibre Channel Expansion Card (CIOv)	46M6065	3594

Figure 3-15 shows the QLogic 4 Gb Fibre Channel Expansion Card (CIOv) for IBM BladeCenter.



Figure 3-15 The QLogic 4 Gb Fibre Channel Expansion Card (CIOv) for IBM BladeCenter

The QLogic 4 Gb Fibre Channel Expansion Card (CIOv) has the following features:

- ▶ CIOv form factor
- ▶ QLogic ISP2532 ASIC
- ▶ PCI Express 2.0 host interface
- ▶ Support for two full-duplex Fibre Channel ports at 4 Gbps maximum per channel
- ▶ Support for Fibre Channel Protocol SCSI (FCP-SCSI) and Fibre Channel Internet Protocol (FC-IP)
- ▶ Support for Fibre Channel service (class 3)
- ▶ Support for switched fabric, point-to-point, and Fibre Channel Arbitrated Loop (FC-AL) connections
- ▶ Configuration and boot support in BIOS and UEFI

This card is installed into the PCI Express CIOv slot of a supported blade server. It provides connections to Fibre Channel-compatible modules in bays 3 and 4 of a supported BladeCenter chassis.

For more information, see the following publications:

- ▶ *QLogic 4 Gb Fibre Channel Expansion Card (CIOv):*
<http://www.redbooks.ibm.com/abstracts/tips0695.html>
- ▶ *QLogic 4 Gb Fibre Channel Expansion Card (CIOv) Installation and User Guide:*
<http://ibm.com/support/entry/portal/docdisplay?lnidocid=MIGR-5079686>

3.1.14 QLogic 8 Gb Fibre Channel Expansion Card (CIOv)

The QLogic 8 Gb Fibre Channel Expansion Card (CIOv) for IBM BladeCenter enables high-speed access for IBM blade servers to connect to a Fibre Channel storage area network (SAN). When compared to the previous generation of 4 Gb adapters, the new adapter doubles the throughput speeds for Fibre Channel traffic.

Table 3-18 lists the ordering information.

Table 3-18 Ordering part number and feature code

Description	Part number	Feature code
QLogic 8 Gb Fibre Channel Expansion Card (CIOv)	44X1945	1462

Figure 3-16 shows the QLogic 8 Gb Fibre Channel Expansion Card (CIOv) for IBM BladeCenter.



Figure 3-16 The QLogic 8 Gb Fibre Channel Expansion Card (CIOv) for IBM BladeCenter

The QLogic 8 Gb Fibre Channel Expansion Card (CIOv) has the following features:

- ▶ CIOv form factor
- ▶ QLogic 2532 8 Gb ASIC
- ▶ PCI Express 2.0 host interface
- ▶ Support for two full-duplex Fibre Channel ports at 8 Gbps maximum per channel
- ▶ Support for Fibre Channel Protocol Small Computer System Interface (FCP-SCSI) and Fibre Channel Internet Protocol (FC-IP)
- ▶ Support for Fibre Channel service (class 3)
- ▶ Support for switched fabric, point-to-point, and Fibre Channel Arbitrated Loop (FC-AL) connections
- ▶ Configuration and boot support in BIOS and UEFI

This card is installed in the PCI Express CIOv slot of a supported blade server. It provides connections to Fibre Channel-compatible modules in bays 3 and 4 of the supported BladeCenter chassis.

For more information, see the *QLogic 8 Gb Fibre Channel Expansion Card (CIOv)*:

<http://www.redbooks.ibm.com/abstracts/tips0692.html>

3.1.15 Emulex 8 Gb Fibre Channel Expansion Card (CIOv)

The Emulex 8 Gb Fibre Channel Expansion Card (CIOv) for IBM BladeCenter, shown in Figure 3-17, enables high-performance connection to a storage area network (SAN). Because of the innovative design of the IBM BladeCenter midplane, this Fibre Channel adapter operates without the need for an optical transceiver module, saving you significant hardware cost. Each adapter provides dual paths to the SAN switches to ensure full redundancy. In addition, the exclusive firmware-based architecture allows firmware and features to be upgraded without taking the server offline or rebooting, and without the need to upgrade the driver.



Figure 3-17 The Emulex 8 Gb Fibre Channel Expansion Card (CIOv) for IBM BladeCenter

Table 3-19 lists the ordering information.

Table 3-19 Ordering part number and feature code

Description	Part number	Feature code
Emulex 8 Gb Fibre Channel Expansion Card (CIOv)	46M6140	3598

The expansion card has the following features and benefits:

- ▶ Support of the 8 Gbps Fibre Channel standard
- ▶ Usage of the Emulex “Saturn” 8 Gb Fibre Channel I/O Controller (IOC) chip
- ▶ Enablement of high-speed and dual-port connection to a Fibre Channel SAN
- ▶ Can be combined with a CFFh card on the same blade server
- ▶ Comprehensive virtualization capabilities with support for N_Port ID Virtualization (NPIV) and Virtual Fabric
- ▶ Simplified installation and configuration using common HBA drivers
- ▶ Efficient administration by using HBAnyware for HBAs anywhere in the SAN
- ▶ Common driver model that eases management and enables upgrades independent of HBA firmware
- ▶ Support of BladeCenter Open Fabric Manager

This card is installed in the PCI Express CIOv slot of a supported blade server. It provides 8 Gbps connections to Fibre Channel-compatible modules in bays 3 and 4 of a supported BladeCenter chassis.

For more information, see the following publications:

- ▶ *Emulex 8 Gb Fibre Channel Expansion Card (CIOv)*:
<http://www.redbooks.ibm.com/abstracts/tips0703.html>
- ▶ *Emulex 8 Gb Fibre Channel Expansion Card (CIOv) Installation and User's Guide*:
<http://ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5080195>

3.1.16 QLogic 2-port 10 Gb Converged Network Adapter (CFFh)

The QLogic 2-port 10 Gb Converged Network Adapter (CFFh) for IBM BladeCenter offers robust 8 Gb Fibre Channel storage connectivity and 10 Gb networking over a single Converged Enhanced Ethernet (CEE) link. Because this adapter combines the functions of a Network Interface Card and a Host Bus Adapter on a single converged adapter, clients can realize potential benefits in cost, power and cooling, and data center footprint by deploying less hardware.

Table 3-20 lists the ordering information.

Table 3-20 Ordering part number and feature code

Description	Part number	Feature code
QLogic 2-port 10 Gb Converged Network Adapter (CFFh)	00Y3280 ^a	A3JB

a. Replaces 42C1830.

Figure 3-18 shows the QLogic 2-port 10 Gb Converged Network Adapter (CFFh) for IBM BladeCenter.



Figure 3-18 The QLogic 2-port 10 Gb Converged Network Adapter (CFFh) for IBM BladeCenter

The expansion card has the following features:

- ▶ Combo Form Factor (CFFh) PCI Express 2.0 x8 adapter
- ▶ Communication module: QLogic ISP8112
- ▶ Support for up to two CEE HSSMs in a BladeCenter H or HT chassis
- ▶ Support for 10 Gb Converged Enhanced Ethernet (CEE)
- ▶ Support for Fibre Channel over Converged Enhanced Ethernet (FCoCEE)
- ▶ Full hardware offload for FCoCEE protocol processing
- ▶ Support for IPv4 and IPv6
- ▶ Support for SAN boot over FCoCEE and iSCSI boot
- ▶ Support for Wake on LAN
- ▶ Support for BladeCenter Open Fabric Manager for BIOS, UEFI, and FCode

Stateless offload features include:

- ▶ IP, TCP, and UDP checksum offloads
- ▶ Large and Giant Send Offload (LSO and GSO)
- ▶ Receive Side Scaling (RSS)
- ▶ Header-data split
- ▶ Interrupt coalescing
- ▶ NetQueue

IEEE standards compliance:

- ▶ 802.1Qbb rev. 0 (Priority-based flow control)
- ▶ 802.1Qaz rev. 0 (Enhanced transmission selection)
- ▶ 802.1Qaz rev. 0 (DCBX protocol)
- ▶ 802.3ae (10 Gb Ethernet)
- ▶ 802.1q (VLAN)
- ▶ 802.3ad (Link Aggregation)
- ▶ 802.1p (Priority Encoding)
- ▶ 802.3x (Flow Control)
- ▶ 802.3ap (KX/KX4)
- ▶ IEEE 1149.1 (JTAG)
- ▶ IPv4 Specification (RFC 791)
- ▶ IPv6 Specification (RFC 2460)
- ▶ TCP/UDP Specification (RFC 793/768)
- ▶ ARP Specification (RFC 826)

The QLogic 2-port 10 Gb Converged Network Adapter requires at least one 10 Gb Ethernet Pass-thru Module (P/N 46M6181) or IBM Virtual Fabric 10Gb Switch Module (P/N 46C7191) to be installed in a BladeCenter H or BladeCenter HT chassis.

For more information, see the following publications:

- ▶ *QLogic 2-port 10 Gb Converged Network Adapter (CFFh):*
<http://www.redbooks.ibm.com/abstracts/tips0716.html>
- ▶ *QLogic 2-port 10 Gb Converged Network Adapter (CFFh) Installation and User's Guide;*
 you must download the "publication release" EXE file and run it to extract the PDF:
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5081072>

3.1.17 Brocade 2-port 10 GbE Converged Network Adapter (CFFh)

The Brocade 2-port 10 GbE Converged Network Adapter, together with the Brocade Converged 10 GbE Switch Module, is a part of a leading Converged Ethernet solution for IBM BladeCenter. This solution offers Fibre Channel migration and investment protection, maximum bandwidth and performance, and simplicity in a converged environment.

This card combines the functions of a Host Bus Adapter (HBA) and Network Interface Card (NIC) on one PCIe 2.0 x8 card. It supports full Fibre Channel over Ethernet (FCoE) protocol offload and allows Fibre Channel traffic to converge onto 10 Gbps Converged Enhanced Ethernet (CEE) networks. FCoE and 10 Gbps CEE operations run simultaneously.

Table 3-21 lists the Brocade 2-port 10 Gb Converged Network Adapter (CFFh) ordering information.

Table 3-21 Ordering part number and feature code

Description	Part number	Feature code
Brocade 2-port 10 Gb Converged Network Adapter (CFFh)	81Y1650	5437

Figure 3-19 shows the Brocade 2-port 10 Gb Converged Network Expansion Card.



Figure 3-19 Brocade 2-port 10 GbE Converged Network Adapter (CFFh)

The expansion card has the following features:

- ▶ CFFh form factor
- ▶ PCI Express x8 Gen 2 host interface
- ▶ Communication module: Two 400 MHz processors ASIC
- ▶ Two ports with XAUI interfaces for PCIe x8 to the high-speed switch module
- ▶ Support for up to two Brocade Converged 10 GbE Switch Modules in a BC-H or BC-HT chassis
- ▶ Support for IPv4 and IPv6
- ▶ Support for BladeCenter Open Fabric Manager for BIOS and UEFI
- ▶ Brocade Host Connectivity Manager (HCM) device management and Brocade Command Line Utility
- ▶ CEE support
 - 10 Gbps throughput per port full duplex
 - 1.8 million packets per second per port in each direction (700-byte packets, latency < 2 μ s)

- Checksum/CRC offloads for FCoE packets, IPv4/IPv6 TCP and UDP packets, and IPv4 header
- VLAN support and tagging
- Jumbo frame support (up to 9600 bytes)
- Header data split (HDS) feature for advanced link layer
- Receive side scaling (RSS) feature for advanced link layer
- TCP segmentation offload (TSO) and large send offload (LSO)
- Link aggregation (NIC teaming)
- Priority-based Flow Control (802.1Qbb)
- Enhanced Transmission Selection (802.1Qaz)
- Data Center Bridging eXchange Protocol (DCBX)
- FCoE support
 - 10 Gbps throughput per port full duplex
 - Full hardware-based FCoE protocol offload
 - 500,000 IOPS per port for maximum IO transfer rates
 - Target rate limiting
 - N_Port ID Virtualization (NPIV)
 - SAN Boot
 - Fabric-based boot LUN discovery
 - Persistent binding
 - Fibre Channel-Security Protocol (FC-SP)
 - FCoE Initialization Protocol (FIP)

The Brocade 2-port 10 Gb Converged Network Adapter is designed to be teamed with the Brocade FCoE Switch Module, which is installed in the high-speed switch bays in the BladeCenter H or HT chassis.

For more information, see the *Brocade 2-port 10 Gb Converged Network Adapter for IBM BladeCenter*.

<http://www.redbooks.ibm.com/abstracts/tips0790.html>

3.1.18 The 2-port 40 Gb InfiniBand Expansion Card (CFFh)

The 2-port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter is a dual port InfiniBand Host Channel Adapter (HCA) based on proven Mellanox ConnectX IB technology. This HCA, when combined with the QDR switch, delivers end-to-end 40 Gb bandwidth per port. This solution is ideal for low latency, high bandwidth, and performance-driven server and storage clustering applications in a High Performance Compute environment. The adapter uses the CFFh form factor and can be combined with a CIOv or CFFv adapter to get additional SAS, Fibre Channel, or Ethernet ports.

Table 3-22 lists the ordering information.

Table 3-22 Ordering part number and feature code

Description	Part number	Feature code
2-port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter	46M6001	0056

Figure 3-20 shows the 2-port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter.



Figure 3-20 The 2-port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter

The 2-port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter includes the following features and specifications.

- ▶ The 2-port 40 Gb InfiniBand Expansion Card features:
 - Form-factor: CFFh
 - Host interface: PCI-E x8 Gen 2 (5.0 GTps): 40+ 40 Gbps bidirectional bandwidth
 - Dual 4X InfiniBand ports at speeds 10 Gbps, 20 Gbps, or 40 Gbps per port
 - 6.5 GBps bidirectional performance
 - RDMA and Send/Receive semantics
 - Hardware-based congestion control
 - Atomic operations
 - Sixteen million I/O channels
 - 4 - 256 KB MTU
 - 1 GB messages
 - Nine virtual lanes: Eight data + one management
 - 1 μ s MPI ping latency
 - Processor offload of transport operations
 - End-to-end QoS and congestion control
 - TCP/UDP/IP stateless offload
- ▶ Enhanced InfiniBand features:
 - Hardware-based reliable transport
 - Hardware-based reliable multicast
 - Extended Reliable Connected transport
 - Enhanced Atomic operations
 - Fine grained end-to-end QoS

- ▶ Hardware-based I/O virtualization features:
 - Single Root IOV
 - Address translation and protection
 - Multiple queues per virtual machine
 - VMware NetQueue support
- ▶ Protocol support:
 - Open MPI, OSU MVAPICH, HP MPI, Intel MPI, MS MPI, and Scali MPI
 - IPoIB, SDP, and RDS
 - SRP, iSER, FCoIB, and NFS RDMA

The 2-port 40 Gb InfiniBand Expansion Card requires at least one Voltaire 40 Gb InfiniBand Switch Module (part number 46M6005) to be installed in a BladeCenter H chassis.

For more information, see *2-port 40 Gb InfiniBand Expansion Card (CFFh)*:

<http://www.redbooks.ibm.com/abstracts/tips0700.html>

3.1.19 SAS Expansion Card (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. After you assign disks to a blade using the SAS Connectivity Modules predefined or user-defined configurations, you can create RAID arrays across those disks using the SAS expansion cards LSI configuration utility.

Table 3-23 lists the ordering information.

Table 3-23 Ordering part number and feature code

Description	Part number	Feature code
SAS Expansion Card (CFFv) for IBM BladeCenter	44E5688 ^a	A3J9

a. Replaces 39Y9190.

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.

The SAS Expansion Card (CFFv) key features are as follows:

- ▶ A PCI-X bus adapter.
- ▶ Uses an LSI 1064 SAS Controller ASIC.
- ▶ Provides a blade with two 1X SAS ports for connecting to the SAS Connectivity Module (see 2.10.1, “SAS Connectivity Module” on page 169).
- ▶ 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 3-21 shows the SAS Expansion Card.



Figure 3-21 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 10 HDDs, plus two optional hot spares. An IS volume supports up to 10 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 are as follows:

- ▶ 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, and intermixing SAS/SATA drives is also supported. However, each volume must have hard disks of the same type, SAS or SATA.

For more information, see the *SAS Expansion Card (CFFv) Installation and User's Guide*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5072377>

3.1.20 SAS Connectivity Card (CIOv)

The SAS Connectivity Card (CIOv) for IBM BladeCenter is an expansion card that offers the ideal way to connect the supported BladeCenter servers to a wide variety of SAS storage devices. Built on proven 3 Gbps, full-duplex, SAS technology, the SAS Connectivity Card can connect to the IBM System Storage DS3200 from the BladeCenter E or H chassis, and to multiple Disk Storage Modules in the BladeCenter S. The card routes the pair of SAS channels from the blade's onboard SAS controller to SAS switches installed in the BladeCenter chassis.

Table 3-24 lists the ordering information.

Table 3-24 Ordering part number and feature code

Description	Part number	Feature code
SAS Connectivity Card (CIOv) for IBM BladeCenter	43W4068	1041

The SAS Connectivity Card (CIOv), shown in Figure 3-22, is a SAS bridge card that is used in conjunction with blade servers possessing a CIOv interface. The SAS Connectivity Card (CIOv) connects unused SAS ports of the onboard SAS controller on the blade server to the blade's midplane connector. The onboard SAS controller is then able to connect to SAS storage through SAS modules in switch bays 3 and 4 in the BladeCenter chassis.



Figure 3-22 The SAS Connectivity Card (CIOv) for IBM BladeCenter

The SAS Connectivity Card (CIOv) has the following features:

- ▶ CIOv form factor
- ▶ Provides external connections for the two SAS ports of blade server's onboard LSI 1064e disk storage controller
- ▶ Support for two full-duplex SAS ports at 3 Gbps maximum per channel
- ▶ Support for SAS, SSP, and SMP protocols

For more information, see the following publications:

- ▶ *SAS Connectivity Card (CIOv) for IBM BladeCenter:*
<http://www.redbooks.ibm.com/abstracts/tips0701.html>
- ▶ *SAS Connectivity Card (CIOv) Installation and User Guide:*
<http://ibm.com/support/entry/portal/docdisplay?lnidocid=MIGR-5079707>

3.1.21 ServeRAID H1135 (CIOv) Controller

The IBM ServeRAID H1135 Controller enables you to use RAID technology to manage hard disk drive configurations and integrate serial-attached SCSI (SAS) technology into your storage infrastructure. You can install the optional IBM ServeRAID H1135 Controller in an IBM BladeCenter HS23E blade server or an IBM Flex System™ x220 compute node.

Table 3-25 lists the ordering information.

Table 3-25 Ordering information for the ServeRAID H1135

Description	Part number	Feature code
ServeRAID H1135 (CIOv) Controller	90Y4570	A1XJ

The ServeRAID H1135 Controller has the following specifications:

- ▶ Based on the LSI SAS2004 6 Gbps SAS 4-port controller
- ▶ Up to 6 Gbps throughput per port
- ▶ CIOv form factor
- ▶ PCIe 2.0 x4 host interface
- ▶ Two SAS ports routed internally to the two hot-swap drive bays
- ▶ Two SAS ports routed externally to the chassis I/O bays 3 and 4
- ▶ Support for SAS/SATA HDD and SSD drives
- ▶ Support for RAID 0, 1, 1E, and 10, and non-RAID
- ▶ Support for up to two RAID volumes
- ▶ Support for up to 10 drives in one RAID volume
- ▶ Support up to 14 volume drives, including up to two hot-spare drives
- ▶ Support for virtual drive sizes greater than 2 TB
- ▶ Fixed stripe size of 64 KB
- ▶ S.M.A.R.T. support
- ▶ Support for MegaRAID Storage Manager management software
- ▶ Support for connectivity to the EXP2512 and EXP2524 storage expansion enclosures
- ▶ Support for connectivity to the BladeCenter S disk storage modules (through SAS Connectivity Modules or SAS RAID Controller Modules), tape drives, and external storage systems
- ▶ Support for operations as a RAID controller for the internal drives and as an HBA for the external storage at the same time

Figure 3-23 shows the ServeRAID H1135 (CIOv) Controller.




Figure 3-23 The ServeRAID H1135 (CIOv) Controller

For more information, see the following publications:

- ▶ *ServeRAID H1135 SAS/SATA Controller:*
<http://www.redbooks.ibm.com/abstracts/tips0054.html?Open#H1135>
- ▶ *ServeRAID H1135 SAS/SATA Controller Installation and User Guide:*
<http://www-947.ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5088601>

3.2 IBM BladeCenter HS12

The BladeCenter HS12 server (machine types 8028) is the first single processor socket blade server offered by IBM. With the increase in performance of dual and quad core processors, single socket servers today have enough processing power to run most single application server workloads.

The following topics are covered in this section:

- ▶ Processor
- ▶ Memory
- ▶ Onboard network adapters
- ▶ Integrated systems management processor
- ▶ Local storage options
- ▶ I/O expansion options

The HS12 is a single-socket Intel processor-based blade with the RAS features that clients expect from IBM BladeCenter, at a lower price point than a two-socket blade server. These features make 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.

Figure 3-24 shows the internal layout of the HS12 type 8028 (with hot-swap drives).




Figure 3-24 The HS12 blade server (machine type 8028)

Features of the HS12 include the following items:

- ▶ A 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 six DIMM sockets.
- ▶ Dual Gigabit Ethernet connections.
- ▶ Support for PCI-X or PCI-Express I/O expansion cards.
- ▶ Two 2.5-inch small form factor (SFF) 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. For more information, see 3.10, "Expansion blades" on page 400.
- ▶ Integrated systems management processor.
- ▶ Supported in all IBM BladeCenter chassis.
- ▶ Three-year, customer-replaceable unit (CRU) and onsite limited warranty on the type 8028 models.

Table 3-26 lists the features of the HS12 blade server.

Table 3-26 Features of the HS12

Feature	Description
Processor options	Intel Xeon 3300 series quad-core
Number of processors (std/max ^a)	1/1
Front-side bus	1066 MHz or 1333 MHz
Cache	6 MB or 12 MB
Memory	Six DIMM slots / 24 GB maximum
Memory Type	PC2-5300 VLP 667 MHz double data rate (DDR2) ECC memory
Internal hard disk drives (std/max)	0/2
Internal storage type	<ul style="list-style-type: none"> ▶ Up to two 2.5-inch hot swap SAS SFF drives ▶ Up to two 2.5-inch hot swap SATA SFF drives ▶ Up to two 2.5-inch hot-swap SFF solid-state drives (SSD)
Network	Two ports, Integrated Dual Gigabit Ethernet (Broadcom 5714)
I/O upgrade	<ul style="list-style-type: none"> ▶ One PCI-X expansion connector ▶ One PCI-Express expansion connector
Warranty	Three-year customer-replaceable unit (CRU) and onsite limited warranty

a. Standard/maximum (std/max)

Table 3-27 lists the HS12 models that are currently available. Models are available with single core, dual core, or quad core processors, and with various clock speeds.

Table 3-27 Models of the HS12

Model	Intel processor	Cores per socket	Processor / FSB speed (GHz/MHz)	L2 cache	Processor (standard/maximum)	Memory (standard/maximum)	Processor power	SAS or SATA card
Machine type 8028 (3-year warranty)								
8028-54Y	Xeon X3323	Quad	2.5/1333	6 MB	1/1	2 GB/24 GB	80 W	SAS
8028-56Y	Xeon X3363	Quad	2.83/1333	12 MB	1/1	2 GB/24 GB	80 W	SAS

Intel Xeon processors: For further information regarding Intel Xeon processor features, go to the following address:

<http://ark.intel.com/products/family/28330/Intel-Xeon-Processor-3000-Sequence?ui=BIG&wapkw=xeon+3000>

3.2.1 Processor

The HS12 blade server supports one Intel Xeon 3300 series quad-core processor, depending on the model.

See Table 3-27 for specifics. The HS12 has only one processor socket.

3.2.2 Memory

There are a total of six DIMM slots in the HS12. The memory used in the HS12 blade server has the following features:

- ▶ PC2-5300 667 MHz high-performance double data rate (DDR2)
- ▶ Very low profile (VLP) form factor
- ▶ ECC
- ▶ Chipkill (on DIMMs 1 GB or larger)

Figure 3-25 shows the physical layout of the memory DIMMs.




Figure 3-25 DIMM layout of the HS12

The 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.

The HS12 supports two mutually exclusive memory configurations: normal and sparing. The configuration is enabled in the 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 is used for sparing.) In an HS12 with 24 GB installed, up to 20 GB (using 6 x dual rank 4 GB 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.

3.2.3 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.

Ethernet ports: 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 software includes a command-line interface (CLI) for scripting purposes.

3.2.4 Integrated systems management processor

The blade server HS12 has an integrated baseboard management controller (BMC) that provides the following benefits:

- ▶ 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, and heartbeat)
- ▶ IPMI over LAN
- ▶ LAN messaging and alerting

BMC access: 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.

3.2.5 Local storage options

HS12 8028 has a SAS interface card installed on a dedicated connector in the server. The interface card comes fitted from the factory and cannot be ordered separately.

Figure 3-26 shows the SAS interface card installed on the HS12 type 8028.




Figure 3-26 SAS Interface Card on an HS12 type 8028

Local drives are 2.5-inch small form factor (SFF) drives. Drive support is indicated in Table 3-28.

Table 3-28 Hard disk drive support on the HS12

Machine type	SAS	Solid state	SATA
8028	Supported (hot-swap)	Supported (hot-swap)	Supported (hot-swap)

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 setup provides a 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.

Local disk drives for HS12 type 8028

Table 3-29 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.

SAS hard disk drives: A SAS hard disk drive cannot be part of a RAID array with a SATA or solid-state drive.

Table 3-29 Small form factor hard disk drive options for HS12 type 8028

Part number	Feature code	Description	Compatible with
81Y9650	A282	IBM 900 GB 10 K 6 Gbps SAS 2.5-inch SFF HS HDD	8028
81Y9670	A283	IBM 300 GB 15 K 6 Gbps SAS 2.5-inch SFF HS HDD	8028
42D0637	5599	IBM 300 GB 10 K 6 Gbps SAS 2.5-inch SFF Slim-HS HDD	8028
42D0677	5536	IBM 146 GB 15 K 6 Gbps SAS 2.5-inch SFF Slim-HS HDD	8028
42D0707	5409	IBM 500 GB 7200 K 6 Gbps NL SAS 2.5-inch SFF Slim-HS HDD	8028
49Y2003	5433	IBM 600 GB 10 K 6 Gbps SAS 2.5-inch SFF Slim-HS HDD	8028
81Y9690	A1P3	IBM 1 TB 7.2 K 6 Gbps NL SAS 2.5-inch SFF HS HDD	8028
81Y9722	A1NX	IBM 250 GB 7.2 K 6 Gbps NL SATA 2.5-inch SFF HS HDD	8028
81Y9726	A1NZ	IBM 500 GB 7.2 K 6 Gbps NL SATA 2.5-inch SFF HS HDD	8028
81Y9730	A1AV	IBM 1 TB 7.2 K 6 Gbps NL SATA 2.5-inch SFF HS HDD	8028

Figure 3-27 shows the hot-swap disk drive bays on the front of the HS12 type 8028.



Figure 3-27 Showing the hot swap disk drive bays on the front of the HS12 type 8028

3.2.6 I/O expansion options

The blade server HS12 has two onboard connectors for various I/O expansion card options:

- ▶ A PCI-X connector
- ▶ 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 through 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
- ▶ An HSFF expansion card

For more information about this topic, see 3.1, “Expansion cards” on page 212.

Table 3-30 lists the available I/O expansion card options for the HS12.

Table 3-30 Blade server HS12 I/O expansion card options

Part number	Feature code	Description
44W4479	5476	2/4 Port Ethernet Expansion Card (CFFh)
46M6164	0098	Broadcom 10 Gb Gen2 4-port Ethernet Exp Cd (CFFh)
00Y3270 ^a	A3JC	QLogic Eth and 8 Gb Fibre Channel Exp Card (CFFh)
00Y3280 ^b	A3JB	QLogic 2-port 10 Gb Converged Network Adapter (CFFh)
46M6001	0056	2-port 40 Gb InfiniBand Expansion Card (CFFh)
44E5688 ^c	A3J9	SAS Expansion Card (CFFv)

a. Replaces 44X1940.

b. Replaces 42C1830.

c. Replaces 39Y9190.

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.

For more information about the PEU3e, see 3.10.1, “PCI Express I/O Expansion Unit” on page 401.

3.3 IBM BladeCenter HS22

The IBM BladeCenter HS22 is a two-socket blade server running the Intel Xeon processors. It is ideal for infrastructure, virtualization, and enterprise business applications and is compatible with the IBM BladeCenter H, E, S, and HT chassis.

The following topics are covered in this section:

- ▶ Model information
- ▶ Chassis support
- ▶ Processor
- ▶ Memory
- ▶ Standard onboard features
- ▶ Local storage
- ▶ Integrated virtualization
- ▶ I/O expansion

The HS22 supports up to two Intel Xeon 5600 series multi-core processors, 12 DIMMs modules, two hot-swap drives, two PCI Express connectors, and one internal USB connector.

Figure 3-28 shows the HS22.



Figure 3-28 The HS22 type 7870

Table 3-31 lists the features of the HS22.

Table 3-31 Features of the HS22 type 7870

Feature	Description
Processor options	Intel Xeon 5600 Processor Series hex-core processors Intel Xeon 5600 Processor Series quad-core processors
Number of processors (std/max)	1/2
Memory Speed	800 MHz, 1033 MHz, or 1333 MHz (dependent on processor model)
Cache	4 MB, 8 MB, or 12 MB (shared between cores)
Memory slots	12 DIMM slots
Memory Maximums	Models with Intel Xeon 5600 series processors: 192 GB using 16 GB DIMMs Models with Intel Xeon 5500 series processors: 96 GB using 8 GB DIMMs
Memory Types	DDR 3 ECC Very Low Profile (VLP) Registered DIMM DDR 3 ECC Very Low Profile (VLP) Low Power Registered DIMM
Hot Swap Disk Drives (std/max)	0/2 (SAS, SATA, or SSD)
Maximum internal storage	Optional: Up to two 2.5-inch hot-swap SAS or SATA HDDs Optional: Up to two 2.5-inch hot-swap solid-state drives
Network	Two ports, one Broadcom BCM5709S dual-port Gigabit Ethernet controller, TOE
I/O upgrade	Two PCI Express connectors (CFFh and CIOv form factors)

Figure 3-29 shows the components on the system board of the HS22.




Figure 3-29 Layout of the HS22 system board

3.3.1 Model information

The current model is listed in Table 3-32.

Table 3-32 Models of the HS22 type 7870

Server models ^a	Processor model	Clock speed	Number of processors (std/max)	Number of cores	Shared cache size	Maximum memory speed	Memory (std/max ^b)
Intel Xeon 5600 series 4-core processors							
7870-G2x	Intel Xeon E5620	2.40 GHz	1/2	4	12 MB	1066 MHz	3x 2 GB / 192 GB

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. The US GAV model numbers use the following nomenclature: xxU. For example, the US orderable part number for 7870-A2x is 7870-A2U. See the product-specific official IBM announcement letter for other country-specific GAV model numbers.

b. Maximum system memory capacity of 192 GB is based on using 12 x 16 GB memory DIMMs.

3.3.2 Chassis support

The HS22 type 7870 is supported in the BladeCenter chassis, as listed in Table 3-33. See the following subsections for details about the support.

Table 3-33 HS22 chassis support

Blade	TDP ^a	BC S (8886)	BC E (8677)	BC T	BC H (8852)	BC HT DC (8740)	BC HT AC (8750)
HS22	130 W	Some limits	No	No	Some limits	No	Some limits
HS22	Up to 95 W	Full	Some limits	No	Full	Full	Full

a. Thermal Design Power

The HS22 type 7870 chassis support depends on the type and model of the chassis and the power supply and cooling module (blower) in the chassis, as shown in Table 3-34 on page 264.

- ▶ A green cell (■) means that the chassis can be filled with HS22 blade servers up to the maximum number of blade bays in the chassis (14 blades in the BladeCenter H).
- ▶ A yellow cell (■) means that the maximum number of HS22 blades that the chassis can hold is fewer than the total available blade bays. *All other bays must remain empty.* Empty bays must be distributed evenly between the two power domains of the chassis (bays 1 - 6 and bays 7 - 14).

BladeCenter E considerations:

- ▶ The HS22 is not supported in BladeCenter E chassis with power supplies that have a capacity less than 2000 W.
- ▶ The HS22 requires an Advanced Management Module to be installed in the BladeCenter E.

Table 3-34 HS22 chassis compatibility

CPU TDP ^a	Maximum number of HS22 supported in each chassis								
	BC-E with AMM (8677) (14 bays)		BC-S (8886) (6 bays)	BC-H (14 bays)				BC-HT DC (8740) (12 bays) ^b	BC-HT AC (8750) (12 bays) ^b
	2000 W power supplies	2320 W power supplies		2900 W power supplies		2980 W power supplies ^c			
				Std. blower	Enh. blower ^d	Std. blower	Enh. blower ^d		
Intel Xeon processors									
130 W	None ^e	None ^e	5	None ^e	14	None ^e	14	4+4	5+5
95 W	5+6	14	6	6+6	14	14	14	12	12
80 W	5+7	14	6	14	14	14	14	12	12
60 W	6+7	14	6	14	14	14	14	12	12
40 W	14	14	6	14	14	14	14	12	12

a. Thermal design power (TDP)

b. For non-NEBS (Enterprise) environments, HS22 models with up to 130 W CPUs are supported. For NEBS environments, HS22 models with up to 80 W CPUs are supported (only specific CPUs are supported: E5620, E5645, L5638, and L5618).

c. IBM BladeCenter H 2980 W AC Power Modules, 68Y6601 (standard in 4Tx, 5Tx, and 9xx BC-H chassis models; optional with all other BC-H chassis models).

d. IBM BladeCenter H Enhanced Cooling Modules, 68Y6650 (standard in 4Tx, 5Tx, and 9xx BC-H chassis models; optional with all other BC-H chassis models).

e. Not supported.

3.3.3 Processor

The HS22 type 7870 features the Intel Xeon 5600 series processors. The Xeon 5600 series processor has models with either four or six cores per processor. With the Xeon 5600 series processors, the memory controller is built into the processor. The memory is directly connected to the processor through three memory controller buses.

The two Xeon processors in the HS22 are connected through a serial coherency link called QuickPath Interconnect (QPI), as shown in Figure 3-30 on page 265. The QPI link is capable of up to 6.4 giga-transfers per second (GTps), depending on the processor model. The I/O hub chip is also connected to each processor through another QPI link.

Figure 3-30 shows the major components of the HS22 system.

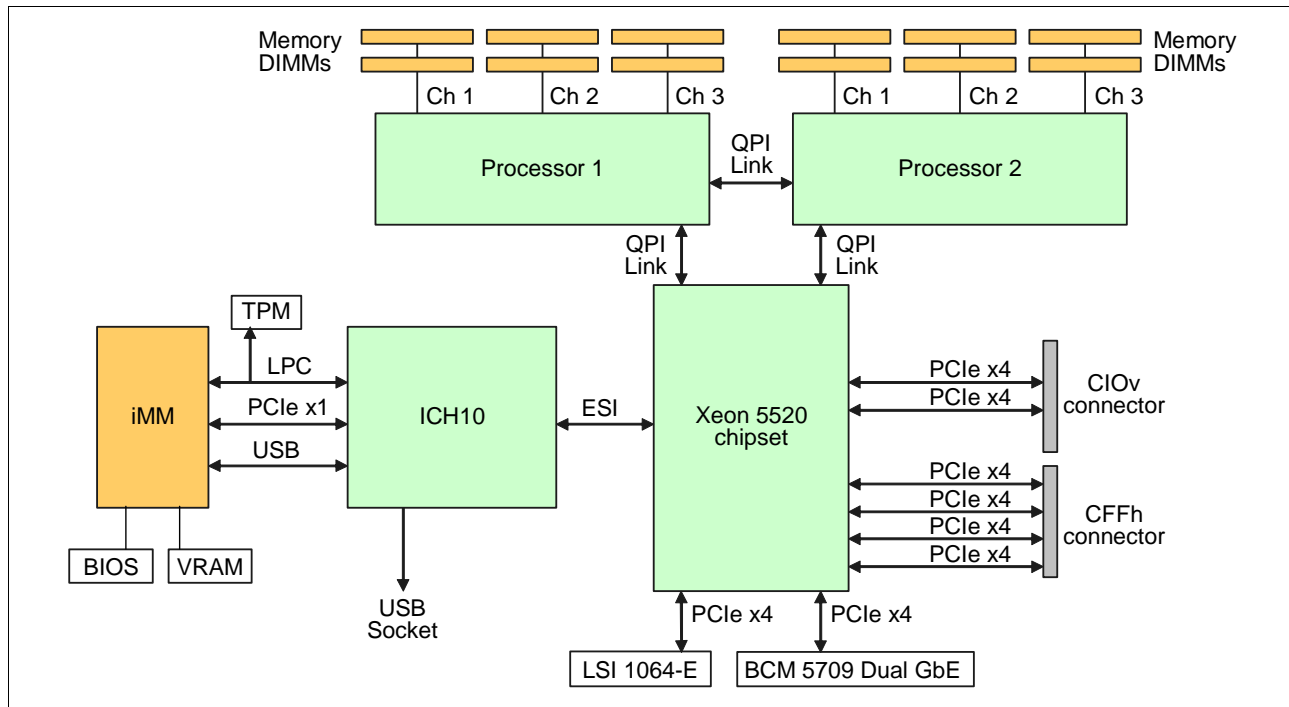


Figure 3-30 HS22 processor board block diagram

The blade server HS22 type 7870 supports either configuration:

- ▶ Up to two Intel Xeon Processor 5600 series hex-core processors
- ▶ Up to two Intel Xeon Processor 5600 series quad-core processors

In a two-socket configuration, both processors must be identical.

The Intel Xeon 5600 series processors are available in various clock speeds and have standard and low-power offerings. The processors support the following Intel technologies:

- ▶ QuickPath Technology
- ▶ Intel Turbo Boost Technology
- ▶ Intel Hyper-Threading Technology
- ▶ Intel Virtualization Technology
- ▶ Intelligent Power Capability
- ▶ Intel 64 Technology

Table 3-35 and Table 3-36 on page 268 list the processor models available for the HS22.

Table 3-35 Intel Xeon 5600 Series processor features

Processor number	Processor frequency	Turbo	HT	L3 cache	Cores	Power	QPI Link speed ^a	Max DDR3 speed	Supports 2DPC/1333 ^b
Advanced									
Xeon X5690	3.46 GHz	Yes	Yes	12 MB	6	130 W	6.4 GTps	1333 MHz	Yes
Xeon X5687	3.60 GHz	Yes	Yes	12 MB	4	130 W	6.4 GTps	1333 MHz	Yes
Xeon X5680	3.33 GHz	Yes	Yes	12 MB	6	130 W	6.4 GTps	1333 MHz	Yes
Xeon X5677	3.46 GHz	Yes	Yes	12 MB	4	130 W	6.4 GTps	1333 MHz	Yes
Xeon X5675	3.06 GHz	Yes	Yes	12 MB	6	95 W	6.4 GTps	1333 MHz	Yes
Xeon X5672	3.20 GHz	Yes	Yes	12 MB	4	95 W	6.4 GTps	1333 MHz	Yes
Xeon X5670	2.93 GHz	Yes	Yes	12 MB	6	95 W	6.4 GTps	1333 MHz	Yes
Xeon X5667	3.06 GHz	Yes	Yes	12 MB	4	95 W	6.4 GTps	1333 MHz	Yes
Xeon X5660	2.80 GHz	Yes	Yes	12 MB	6	95 W	6.4 GTps	1333 MHz	Yes
Xeon X5650	2.66 GHz	Yes	Yes	12 MB	6	95 W	6.4 GTps	1333 MHz	Yes
Xeon X5647	2.93 GHz	Yes	Yes	12 MB	4	130 W	5.86 GTps	1066 MHz	No
Standard									
Xeon E5649	2.53 GHz	Yes	Yes	12 MB	6	80 W	5.86 GTps	1333 MHz	No
Xeon E5640	2.66 GHz	Yes	Yes	12 MB	4	80 W	5.86 GTps	1066 MHz	No
Xeon E5630	2.53 GHz	Yes	Yes	12 MB	4	80 W	5.86 GTps	1066 MHz	No
Xeon E5620	2.40 GHz	Yes	Yes	12 MB	4	80 W	5.86 GTps	1066 MHz	No
Xeon E5607	2.26 GHz	No	No	8 MB	4	80 W	4.8 GTps	1066 MHz	No
Xeon E5603	1.60 GHz	No	No	4 MB	4	80 W	4.8 GTps	1066 MHz	No
Low power									
Xeon L5640	2.26 GHz	Yes	Yes	12 MB	6	60 W	5.86 GTps	1066 MHz	No

a. GTps: giga-transfers per second

b. Whether the processor supports the ability to run two DIMMs per channel at 1333 MHz. For details, see 3.3.4, "Memory" on page 269.

The Xeon 5600 series of processors have several features:

- Nehalem architecture

The latest-generation Intel microarchitecture represents the next step in faster, multi-core technology that intelligently maximizes performance to match your workload. It is designed to take advantage of Intel 45 nm (for the 5500 series) and 32 nm (for the 5600 series) high-k metal gate silicon technology. The architecture uses parallel processing enabled by an integrated memory controller and Intel QuickPath Technology, providing high-speed interconnects between independent processing cores.

- QuickPath Technology

Intel QuickPath Technology is a platform architecture that provides high-speed (up to 25.6 GBps), point-to-point connections between processors, and between processors and the I/O hub. Each processor has its own dedicated memory that it accesses directly through an Integrated Memory Controller. In cases where a processor needs to access the dedicated memory of another processor, it can do so through a high-speed QuickPath Interconnect (QPI) that links all the processors.

- Intel Turbo Boost Technology

Intel Turbo Boost Technology dynamically turns off unused processor cores and increases the clock speed of the cores in use. It increases the frequency in steps of 133 MHz (to a maximum of three steps or 400 MHz) if the processor's predetermined thermal and electrical requirements are still met.

For example, with three cores active, a 2.26 GHz processor can run the cores at 2.4 GHz. With only one or two cores active, the same processor can run those cores at 2.53 GHz. Similarly, a 2.93 GHz processor can run at 3.06 GHz or even 3.33 GHz. When the cores are needed again, they are dynamically turned back on and the processor frequency is adjusted so. This feature can be enabled or disabled in the UEFI BIOS.

For more information about Intel Turbo Boost Technology, go to the following address:

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

- Intel Hyper-Threading Technology

Intel Hyper-Threading Technology boosts performance for parallel, multi-threaded applications. It enables simultaneous multi-threading within each processor core, up to two threads per core or eight threads per quad-core processor. Hyperthreading reduces computational latency, making optimal use of every clock cycle. This feature can be enabled or disabled in the UEFI BIOS.

- Intelligent Power Capability

Intelligent Power Capability powers individual processor elements on and off as needed, to reduce power draw.

- Intel Virtualization Technology

Improvements to Intel Virtualization Technology (Intel VT-x) provide hardware-assisted page-table management, allowing the guest operating system more direct access to the hardware and reducing compute-intensive software translation from the VMM. Intel VT-x also includes Intel VT FlexMigration and Intel VT Flex-Priority, which are capabilities for flexible workload migration and performance optimization across the full range of 32-bit and 64-bit operating environments.

For more information about Intel Virtualization Technology, go to the following address:

<http://www.intel.com/technology/virtualization/technology.htm>

The Xeon 5600 series of processors have an additional feature: Advanced Encryption Standard New Instructions (AES-NI). These new instructions allow the processor to perform

hardware-accelerated encryption for the AES algorithm. By significantly accelerating encryption, AES- NI enables more secure systems. Increased speed means that longer encryption keys can be used, strengthening the protection of data. More network traffic can be encrypted with less impact on server performance. Database and storage encryption can be more widely and be cost-effectively deployed.

Table 3-36 lists the processor options for the HS22.

Table 3-36 Processors for the HS22 type 7870

Description	Part number	Feature code	Supported model
Intel Xeon 5600 series processors			
Intel Xeon X5690 6C 3.46 GHz 12 MB 1333 MHz 130 W	81Y9331	A16C	C8x
Intel Xeon X5687 4C 3.60 GHz 12 MB 1333 MHz 130 W	81Y9330	A16D	C7x
Intel Xeon X5680 HC 3.33 GHz 12 MB Cache 1333 MHz	59Y5713	7751	F2x
Intel Xeon X5677 QC 3.46 GHz 12 MB Cache 1333 MHz	59Y5714	7750	F3x
Intel Xeon X5675 6C 3.06 GHz 12 MB 1333 MHz 95 W	81Y9329	A16E	C6x
Intel Xeon X5672 4C 3.20 GHz 12 MB 1333 MHz 95 W	81Y9328	A16F	C5x
Intel Xeon X5670 HC 2.93 GHz 12 MB Cache 1333 MHz	59Y5711	7753	H4x
Intel Xeon X5667 QC 3.06 GHz 12 MB Cache 1333 MHz	59Y5712	7752	H5x
Intel Xeon X5660 HC 2.80 GHz 12 MB Cache 1333 MHz	59Y5710	7754	CTO
Intel Xeon X5650 HC 2.66 GHz 12 MB Cache 1333 MHz	59Y5709	7755	H2x, HAx
Intel Xeon E5649 6C 2.53 GHz 12 MB 1333 MHz 80 W	81Y9327	A16G	B6x, D4x
Intel Xeon X5647 4C 2.93 GHz 12 MB 1066 MHz 130 W	81Y9326	A16H	B5x
Intel Xeon E5645 HC 2.40 GHz 12 MB Cache 1333 MHz	68Y8125	7668	CTO, 91x, 92x
Intel Xeon E5640 QC 2.53 GHz 12 MB Cache 1066 MHz	59Y5708	7756	G4x, GCx
Intel Xeon L5638 HC 2.0 GHz 12 MB Cache 1333 MHz	68Y8124	7669	CTO
Intel Xeon E5630 QC 2.40 GHz 12 MB Cache 1066 MHz	59Y5707	7757	CTO
Intel Xeon E5620 QC 2.26 GHz 12 MB Cache 1066 MHz	59Y5705	7758	G2x
Intel Xeon L5640 HC 2.26 GHz 12 MB Cache 1333 MHz	59Y5706	7759	N2x
Intel Xeon L5630 QC 2.13 GHz 4 MB Cache 800 MHz	59Y5704	7760	CTO
Intel Xeon L5609 QC 1.60 GHz 12 MB Cache 1066 MHz	49Y5184	7761	CTO
Intel Xeon E5607 4C 2.26 GHz 8 MB 1066 MHz 80 W	81Y9325	A16J	A7x
Intel Xeon E5603 4C 1.60 GHz 4 MB 1066 MHz 80 W	81Y9323	A16L	A5x

For more information about the Intel Xeon 5600 series processors, go to the following address:

http://www.intel.com/p/en_US/products/server/processor/xeon5000

3.3.4 Memory

The HS22 has six DIMM sockets per processor (12 DIMMs in total). The HS22 with the Xeon 5600 series processor can support up to 192 GB of memory in total when using 16 GB DIMMs and with both processors installed. The HS22 uses Double Data Rate-3 (DDR-3) very-low-profile (VLP) registered DIMMs.

Figure 3-31 shows the location of the 12 memory DIMM sockets.




Figure 3-31 DIMM layout on the HS22 system board

The Xeon 5600 series processors have an integrated memory controller. The memory DIMMs connect directly to one of the processors. With one processor installed, only six DIMMs can be accessed. To access memory in the other six DIMM connectors, the second processor must be installed.

Table 3-37 shows which memory DIMM connectors are connected to each processor.

Table 3-37 Memory channel configuration

Processor	Memory channel	DIMM connector
1	Channel 0	DIMM connector 1 and 2
	Channel 1	DIMM connector 5 and 6
	Channel 2	DIMM connector 3 and 4
2	Channel 3	DIMM connector 7 and 8
	Channel 4	DIMM connector 11 and 12
	Channel 5	DIMM connector 9 and 10

Memory DIMM installation recommendations

The HS22 boots with one memory DIMM installed; however, the suggested memory configuration is to balance the memory across all the memory channels on each processor to use the available memory bandwidth. Therefore, the suggested memory configuration is either:

- ▶ Three or six memory DIMMs in a single processor HS22 blade server
- ▶ Six or 12 memory DIMMs in a dual processor HS22 blade server

This sequence spreads the DIMMs across as many memory channels as possible. For best performance, install the DIMMs in the sockets as shown in the following tables:

- Table 3-38, if you have one processor installed

Table 3-38 Suggested DIMM installation for the HS22 with one processor installed

Number of processors	Number of DIMMs	Processor 1						Processor 2					
		Channel		Channel		Channel		Channel		Channel		Channel	
		DIMM 1	DIMM 2	DIMM 3	DIMM 4	DIMM 5	DIMM 6	DIMM 7	DIMM 8	DIMM 9	DIMM 10	DIMM 11	DIMM 12
1	1		x										
1	2		x		x								
1	3		x		x		x						
1	4	x	x		x		x						
1	5	x	x	x	x		x						
1	6	x	x	x	x	x	x						

- Table 3-39, if you have two processors installed

Table 3-39 Suggested DIMM installation for the HS22 with two processors installed

Number of processors	Number of DIMMs	Processor 1						Processor 2					
		Channel		Channel		Channel		Channel		Channel		Channel	
		DIMM 1	DIMM 2	DIMM 3	DIMM 4	DIMM 5	DIMM 6	DIMM 7	DIMM 8	DIMM 9	DIMM 10	DIMM 11	DIMM 12
2	2		x						x				
2	3		x		x				x				
2	4		x		x				x		x		
2	5		x		x		x		x		x		
2	6		x		x		x		x		x		x
2	7	x	x		x		x		x		x		x
2	8	x	x		x		x	x	x		x		x
2	9	x	x	x	x		x	x	x		x		x
2	10	x	x	x	x		x	x	x	x	x		x
2	11	x	x	x	x	x	x	x	x	x	x		x
2	12	x	x	x	x	x	x	x	x	x	x	x	x

Terminology: The term *DIMMs per channel* is commonly referred to as the abbreviation DPC. One DIMM per channel can be referred to as 1DPC and, likewise, two DIMMs per channel can be referred to as 2DPC. The ability to run two DIMMs per channel at 1333 MHz is often referred to as the abbreviation 2DPC@1333MHz.

Memory modes

The HS22 type 7870 supports two memory modes.

These modes can be selected in the UEFI BIOS setup:

- ▶ Independent channel mode

Independent channel mode gives a maximum of 96 GB of usable memory with one processor installed, and 192 GB of usable memory with two processors installed (using 16 GB DIMMs). The DIMMs can be installed without matching sizes. Independent channel mode is the most appropriate memory mode for most HS22 configurations.

- ▶ Memory mirroring mode

Memory mirroring works much like disk mirroring and provides high server availability in the event of a memory DIMM failure. The total memory is divided into two channels. The third channel is unused. Data is written concurrently to both channels. With mirroring enabled, only one-third of total memory is available for use, because two DIMM slots are the primary channel, two DIMM slots are the backup channel, and two DIMM slots are not used. Because of the double writes to memory, performance is affected. Figure 3-32 illustrates the HS22 memory mirroring layout.

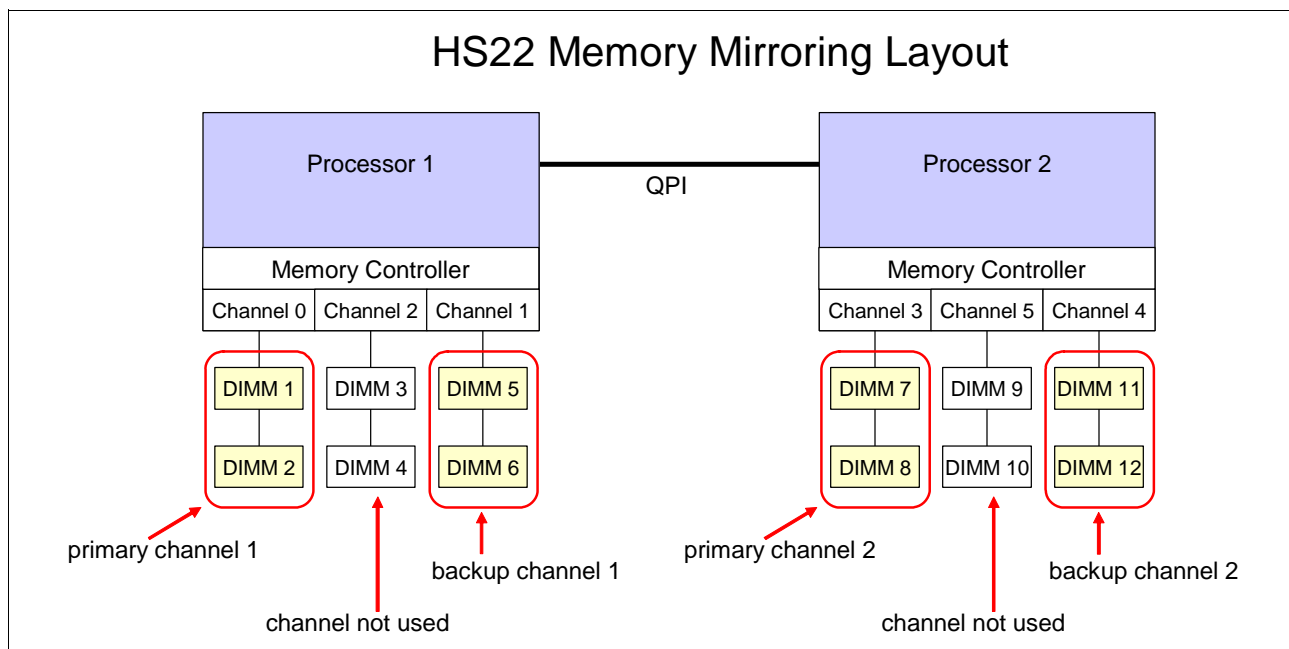


Figure 3-32 The HS22 type 7870 memory mirroring layout

If a failed DIMM is one of the DIMMs in the primary channel, it is instantly disabled and the mirrored (backup) memory in the other channel becomes active (primary) until the failing DIMM is replaced. Mirroring is handled at the hardware level, and no operating system support is required.

Memory for HS22 servers with Xeon 5600 series processors

Standard DIMMs supported in the HS22 operate at 1.5 V. The HS22 server with the Xeon 5600 series processors installed also supports low-power DIMMs that can operate at 1.35 V. The 1.35 V low-power memory DIMMs can be up to 10% more energy efficient than existing 1.5 V memory DIMMs.

The HS22 models with one of the Advanced (X) Xeon 5600 series processors installed can run specific DIMMs at 1333 MHz even when two DIMMs per channel are populated. Table 3-35 on page 266 shows which processors support two DIMMs per channel. Table 3-41 on page 274 shows which DIMMs support two DIMMs per channel at 1333 MHz.

The maximum speed at which the memory channel runs on the HS22 is a combination of the following factors:

- ▶ The processor:
 - The X (Advanced) range of Xeon 5600 series processors has a maximum memory channel speed of 1333 MHz. This range is the only range of Xeon processors capable of running memory DIMMs that are rated at 1333 MHz that are actually at 1333 MHz.
 - The E and L range of Xeon 5600 series processor models has a maximum memory speed of 1066 MHz or 800 MHz and cannot run the memory any faster.
- ▶ The memory DIMM module:
 - Not all DIMMs support running at 1333 MHz.
 - Certain memory DIMMs operate only up to 1066 MHz. Even if these 1066 MHz DIMMs are used with a processor that supports 1333 MHz, the memory channel speed still does not exceed 1066 MHz.
 - Dual-rank DIMMs offer better interleaving and thus better performance than single-rank DIMMs. See Table 3-41 on page 274 to determine the maximum memory channel speed that each DIMM supports and the number of ranks (note the Rank column in those tables).
- ▶ The number of DIMM modules installed:
 - The X (Advanced) range of Xeon 5600 series processors can run two DIMMs per memory channel and still operate the memory bus speed at 1333 MHz.
 - With the E and L range of the Xeon 5600 series range of processors, the more memory DIMMs per channel, the lower the speed at which the memory channels operate.
 - DIMM modules should be installed in lots of three, which for the HS22 is either three or six DIMM modules per processor. This configuration makes the best use of the bandwidth available to the processor's three memory channels.

See Table 3-35 on page 266 to determine the maximum memory channel speed for Xeon 5600 series processor.

Key information:

- ▶ The HS22 with Intel Xeon 5600 series processors supports 1.35 V or 1.5 V DIMMs.
- ▶ If the low-power 1.35 V DIMMs are installed with standard 1.5 V DIMMs, then these low-power DIMMs also operate at 1.5 V and thus lose their energy-efficient capability.
- ▶ For models with the Xeon 5600 processor, 1.35 V DIMMs operate at 1.5 V if you enable 2DPC at 1333 MHz in UEFI.
- ▶ Low-power 1.35 V DIMMs support 1DPC@1333MHz or 2DPC@1066MHz, and 1.5-volt DIMMs support 2DPC@1333MHz (if the processor also supports these memory speeds).

Figure 3-33 is a graphical representation of one and two DIMM modules per processor memory channel.




Figure 3-33 Intel Xeon 5600 series processor with one and two DIMMs per channel

Table 3-40 lists the supported memory speed for a Xeon 5600 series processor type with one or two DIMM modules per channel installed.

Table 3-40 Xeon 5600 series processor memory speeds and DIMMs per channel

Processor type ^a	QPI link speed ^b	Max DDR3 speed	Memory voltage	1DPC ^c	2DPC ^c
Advanced	6.4 GTps	1333 MHz	1.5 V	1333 MHz	1333 MHz
			1.35 V	1333 MHz	1066 MHz
Standard	5.86 GTps	1066 MHz	1.5 V	1066 MHz	1066 MHz
			1.35 V	1066 MHz	1066 MHz
Basic	4.8 GTps	800 MHz	1.5 V	800 MHz	800 MHz
			1.35 V	800 MHz	800 MHz

a. For more information about processor type, see Table 3-35 on page 266.

b. GTps: Giga-transfers per second.

c. Maximum memory speed possible. Actual memory speed also depends on other factors.

Table 3-41 lists the memory options for the HS22 using the Xeon 5600 series processors (ordered by voltage and then DIMM capacity) available for order at time of writing.

Table 3-41 List of memory modules for the HS22 type 7870 using the Xeon 5600 series processors

Part number	Feature code	Description	2DPC @1333 ^a	Rank
Low Power DIMMs (1.35 V)^b				
46C0560	A0WX	2 GB (1x 2 GB, 1Rx8, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	Yes ^c	Single
46C0563	A0WY	4 GB (1x 4 GB, 1Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	Yes ^c	Single
46C0564	A0WZ	4 GB (1x 4 GB, 2Rx8, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	Yes ^c	Dual
46C0568	8644	8 GB (1x 8 GB, 2Rx4, 1.35 V) PC3-10600 CL9 ECC DDR3 1333 MHz VLP	Yes ^c	Dual
00D4981	A3BT	8GB (1x8GB, 1Rx4, 1.35V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	Yes ^c	Single
00D4985	A3BU	8GB (1x8GB, 2Rx8, 1.35V) PC3L-10600 CL9 ECC DDR3 1333MHz VLP	Yes ^c	Dual
46C0599	2422	16 GB (1x 16 GB, 2Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	Yes ^c	Dual

- a. These DIMMs are supported to run two DIMMs per channel at a memory bus speed of 1333 MHz ("2DPC@1333MHz") with the X (Advanced) range of Intel Xeon 5600 series processors. See Table 3-35 to determine the processors in the X range.
- b. Low-power 1.35 V DIMMs operate at 1.5 V if they are mixed with 1.5 V DIMMs.
- c. Supports 2DPC@1333 MHz only when operating at 1.5 V.

Memory performance

For best memory performance consider the following information:

- ▶ Choose a Xeon processor model that can support the highest required maximum memory speed. Some Xeon processor models run memory only at 800 MHz, as shown in the rightmost column in Table 3-35 on page 266 or Table 3-36 on page 268.
- ▶ Choose a Xeon 5600 processor model and memory DIMM combination that supports running 2DPC@1333MHz.
- ▶ Always populate both processors with equal amounts of memory to ensure a balanced system.
- ▶ Always populate all three memory channels on each processor with equal memory capacity.
- ▶ Ensure that an even number of ranks are populated per channel.
- ▶ Use dual-rank DIMMs when possible.
- ▶ Disable memory mirroring mode.

Preferred practice: Memory in Intel servers has historically been installed in pairs (because of interleaving requirements). It was common to see servers configured with 2 GB, 4 GB, 8 GB, 16 GB, or 32 GB of memory installed.

With the Xeon 5600 series processors, for optimal performance, install memory in triplets. Therefore, as an example, configure the HS22 with 6 GB, 12 GB, 24 GB, or 48 GB of memory.

For more information about optimizing memory performance, see *Understanding Intel Xeon 5600 Series Memory Performance and Optimization in IBM System x and BladeCenter Platforms*:

<ftp://public.dhe.ibm.com/common/ssi/ecm/en/xsw03075usen/XSW03075USEN.PDF>

3.3.5 Standard onboard features

The HS22 blade server has the following standard onboard features.

UEFI

The HS22 uses an integrated Unified Extensible Firmware Interface (UEFI) next-generation BIOS. Its new capabilities include:

- ▶ Human readable event logs. No more beep codes!
- ▶ Complete setup solution by allowing the adapter configuration function to be moved into the UEFI.
- ▶ Complete out-of-band coverage by the Advance Settings Utility to simplify remote setup.

Using all of the features that UEFI offers requires a UEFI-aware operating system and adapters. UEFI is fully compatible with earlier versions of the BIOS.

For more information about UEFI, see *Introducing uEFI-Compliant Firmware on IBM System x and BladeCenter Servers*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5083207>

Onboard network adapter

The HS22 includes a dual-port Gigabit Ethernet controller with the following specifications:

- ▶ Broadcom BCM5709S Gigabit Ethernet controller
- ▶ Supports TCP offload engine (TOE)
- ▶ Supports failover and load balancing for better throughput and system availability
- ▶ Supports highly secure remote power management using IPMI 2.0
- ▶ Supports Wake on LAN and Preboot Execution Environment (PXE)
- ▶ Support for IPv4 and IPv6

Integrated systems management processor

The HS22 type 7870 blade server includes an Integrated Management Module (IMM) to monitor server availability, perform Predictive Failure Analysis, and trigger IBM Systems Director alerts. The IMM performs the functions of the Baseboard Management Controller (BMC) of earlier blade servers, and adds the features of the Remote Supervisor Adapter found in System x servers, and remote control and remote media. A separate card, like the cKVM card, is no longer required for these functions (the cKVM card is not supported on the HS22).

For more information about the IMM, see *Transitioning to UEFI and IMM*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5079769>

Video controller

The Video subsystem in the HS22 is a component of the Integrated Management Module (IMM) and is based on the Matrox G200eV video core with 16 MB of video memory. Table 3-42 lists the supported video resolutions.

Table 3-42 Supported video resolutions

Resolution	Maximum refresh
640 x 480	85 Hz
800 x 600	85 Hz
1024 x 768	75 Hz

Trusted Platform Module

Trusted computing is an industry initiative that provides a combination of secure software and secure hardware to create a trusted platform. It is a specification that increases network security by building unique hardware IDs into computing devices.

The HS22 implements Trusted Platform Module (TPM) Version 1.2 support.

3.3.6 Local storage

The HS22 blade server features an onboard LSI Logic 53C1064E SAS controller with two hot-swap drives bays that are accessible from the front of the blade server (Figure 3-34). The LSI Logic 53C1064E SAS controller provides RAID 0 or RAID 1 capability and supports up to two internal hot-swap SAS or SATA HDDs or two internal hot-swap solid-state drives. The LSI 53C1064E SAS controller ports operate at a burst rate of 1.5 or 3.0 Gbps.



Figure 3-34 The HS22 showing the front hot-swap disk drive bays

Local SAS or SATA disk drives

The HS22 type 7870 has support for up to two hot-swap SFF SAS or SATA hard disk drives. These two hot-swap hard disk drives are accessible from the front of the blade server without removing the blade server from the BladeCenter chassis. See Table 3-43 for a list of supported SAS and SATA HDDs.

Table 3-43 Supported SAS and SATA HDDs available for order at the time of writing

Part number	Feature code	Description
SAS hard disk drives		
42D0677 ^a	5536	IBM 146 GB 15K 6 Gbps SAS 2.5-inch SFF Slim-HS HDD
42D0637 ^a	5599	IBM 300GB 10K 6 Gbps SAS 2.5-inch SFF Slim-HS HDD
81Y9670 ^a	A283	IBM 300 GB 15K 6 Gbps SAS 2.5-inch SFF HS HDD
49Y2003 ^a	5433	IBM 600GB 10 K 6 Gbps SAS 2.5-inch SFF Slim-HS HDD
81Y9650 ^a	A282	BM 900GB 10K 6 Gbps SAS 2.5-inch SFF HS HDD
00AD075	A48S	IBM 1.2 TB 10K 6Gbps SAS 2.5in G2HS HDD
42D0707 ^a	5409	IBM 500GB 7200K 6 Gbps NL SAS 2.5-inch SFF Slim-HS HDD
81Y9690 ^a	A1P3	IBM 1 TB 7.2K 6 Gbps NL SAS 2.5-inch SFF HS HDD
SATA hard disk drives		
81Y9722 ^a	A1NX	IBM 250GB 7.2K 6 Gbps NL SATA 2.5-inch SFF HS HDD
81Y9726 ^a	A1NZ	IBM 500GB 7.2K 6 Gbps NL SATA 2.5-inch SFF HS HDD
81Y9730 ^a	A1AV	IBM 1TB 7.2 K6 Gbps NL SATA 2.5-inch SFF HS HDD

a. Operates only up to 3 Gbps when installed in the HS22.

Solid-state drives

Solid-state drives (SSD) are more reliable than spinning hard disk drives and they consume much less power than a standard SAS drive. Target Applications for SSDs include video surveillance, transaction-based databases, and other applications that have high performance but moderate space requirements.

Table 3-44 lists the supported SSDs.

Table 3-44 Supported solid-state drives (SSDs) available for order at the time of writing

Part number	Feature code	Description
49Y5839		IBM 64GB SATA 2.5in MLC HS Enterprise Value SSD
90Y8648		IBM 128GB SATA 2.5in MLC HS Enterprise Value SSD
90Y8643		IBM 256GB SATA 2.5in MLC HS Enterprise Value SSD
49Y5844		IBM 512GB SATA 2.5in MLC HS Enterprise Value SSD
00W1125		IBM 100GB 2.5in SATA MLC HS Enterprise SSD
49Y6129		IBM 200GB 2.5in SAS MLC HS Enterprise SSD
49Y6134		IBM 400GB 2.5in SAS MLC HS Enterprise SSD
49Y6139		IBM 800GB 2.5in SAS MLC HS Enterprise SSD

IBM High IOPS SSD adapters

The HS22 type 7870 can support the installation of one or two IBM High IOPS SSD adapters in a PCI Express expansion unit attached to the blade server. The combination of the HS22 and the expansion unit takes two blade slots in a BladeCenter chassis. This configuration allows self-contained high-performance servers and computing appliances in a BladeCenter environment. These adapters deliver throughput of up to 100,000 I/O operations per second (IOPS), while providing the added benefits of lower power, cooling, and management impact, and a smaller storage footprint.

One or two IBM High IOPS SSD adapters can be installed in the PCI Express I/O expansion unit. Table 3-45 lists the IBM High IOPS SSD adapters and the supported PCI Express expansion units.

Table 3-45 IBM High IOPS SSD PCIe adapters and the supported PCI Express expansion units

Part number	Feature code	Description	PCI Express I/O Expansion Unit 43W4391	PCI Express Gen 2 Expansion Blade 46M6730
46M0877	0096	IBM 160 GB High IOPS SS Class SSD PCIe Adapter	Y	N
46M0898	1649	IBM 320 GB High IOPS MS Class SSD PCIe Adapter	Y	N
81Y4519	5985	640 GB High IOPS MLC Duo Adapter	Y	N

The IBM High IOPS SSD PCIe adapters have the following specifications:

- ▶ 160 GB, 320 GB, or 640 GB storage capacity
- ▶ Functions as a PCIe storage and controller device
- ▶ PCI Express low-profile MD2 form factor
- ▶ PCI Express 1.1 x4 host interface
- ▶ Power dissipation: 6 W (nominal), 9 W (maximum)

For more information about the expansion units, see:

- ▶ 3.10.1, "PCI Express I/O Expansion Unit" on page 401
- ▶ 3.10.2, "PCI Express Gen 2 Expansion Blade/Expansion Blade II" on page 403

For more information about the IBM High IOPS SSD PCIe Adapters, see *IBM High IOPS SSD PCIe Adapters*:

<http://www.redbooks.ibm.com/abstracts/tips0729.html>

3.3.7 Integrated virtualization

The HS22 offers an IBM standard USB flash drive option preinstalled with VMware ESXi. This is an embedded version of VMware ESXi. It is fully contained on the flash drive, without requiring any disk space. The IBM USB Memory Key for VMware Hypervisor plugs into the USB port at the rear of the HS22 system board, as shown in Figure 3-35. The CIOv expansion card is installed over the top of the USB memory key.




Figure 3-35 Rear view of HS22 system board showing position of USB port with embedded hypervisor memory key installed

Table 3-46 lists the part number, feature code, and description of the available IBM USB Memory Keys for VMware Hypervisor.

Table 3-46 IBM USB Memory Key for VMware Hypervisor

Part number	Feature code	Description
41Y8287	3033	IBM USB Memory Key for VMware ESXi 4.1
41Y8296	A1NP	IBM USB Memory Key for VMware ESXi 4.1 Update 1
41Y8300	A2VC	IBM USB Memory Key for VMware ESXi 5.0
41Y8307	A383	IBM USB Memory Key for VMware ESXi 5.0 Update 1
41Y8311	A2R3	IBM USB Memory Key for VMware ESXi 5.1
41Y8298	A2G0	IBM Blank USB Memory Key for VMware ESXi Downloads

For a complete description of the features and capabilities of VMware ESX Server, go to the following address:

<http://www.vmware.com/products/vi/esx/>

3.3.8 I/O expansion

The HS22 type 7870 is able to connect to a wide variety of networks and fabrics by installing the appropriate I/O expansion card. Supported networks and fabrics include 1 Gb and 10 Gb Ethernet, 4 Gb and 8 Gb Fibre Channel, Serial Attached SCSI (SAS), and InfiniBand. An HS22 blade server with an I/O expansion card installed must be installed in a BladeCenter chassis with switch modules (or pass-through) that are compatible with the I/O expansion card.

The HS22 type 7870 blade server supports two formats of I/O expansion cards: the new CIOv format and the CFFh format, as described in the following sections.

CIOv

The CIOv I/O expansion connector provides I/O connections to modules in bays 3 and 4 of a supported BladeCenter chassis. The CIOv slot is a second-generation PCI Express 2.0 x8 slot. A maximum of one CIOv I/O expansion card is supported per single-wide blade server. A CIOv I/O expansion card can be installed on a blade server at the same time as a CFFh I/O expansion card is installed. Table 3-47 lists the CIOv expansion cards that are supported in the HS22 type 7870.

Table 3-47 Supported CIOv expansion cards for the HS22 server

Part number	Feature code	Description
44W4475	1039	Ethernet Expansion Card (CIOv)
46M6065	3594	QLogic 4 Gb Fibre Channel Expansion Card (CIOv)
43W4068	1041	SAS Connectivity Card (CIOv)
46M6140	3598	Emulex 8 Gb Fibre Channel Expansion Card (CIOv)
44X1945	1462	QLogic 8 Gb Fibre Channel Expansion Card (CIOv)

For the latest information about the expansion cards supported by the HS22, see the IBM ServerProven website:

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

CIOv expansion cards are installed in the CIOv slot in the HS22 (Figure 3-36).




Figure 3-36 The HS22 type 7870 showing the CIOv I/O Expansion card position

The CIOv Expansion Cards require a supported I/O module that is installed in bay 3 and bay 4 of the chassis where the cards and servers are installed.

CFFh

The CFFh I/O expansion connector provides I/O connections to high speed switch modules in bays 7, 8, 9, and 10 of a BladeCenter H or BladeCenter HT chassis or to switch bay 2 in a BladeCenter S chassis. The CFFh slot is a second-generation PCI Express x16 (PCIe 2.0 x16) slot. A maximum of one CFFh I/O expansion card is supported per blade server. A CFFh I/O expansion card can be installed on a blade server at the same time as a CIOv I/O expansion card is installed. Table 3-48 lists the available CFFh expansion cards that are supported in the HS22 type 7870.

Table 3-48 Supported CFFh expansion cards for the HS22 type 7870 server

Part number	Feature code	Description
44W4479	5476	2/4 Port Ethernet Expansion Card (CFFh)
44W4479	5476	2/4 Port Ethernet Expansion Card (CFFh)
46M6001	0056	2-port 40 Gb InfiniBand Expansion Card (CFFh)
46M6164	0098	Broadcom 10 Gb Gen2 4-port Ethernet Exp Cd (CFFh)
81Y1650	5437	Brocade 2 port 10GbE Converged Network Adapter
00Y3264	A3NW	Emulex 10 GbE Virtual Fabric Adapter Advanced II
00Y3266	A3NV	Emulex 10 GbE Virtual Fabric Adapter II
42C1810	3593	Intel 10 Gb 2-port Ethernet Expansion Card (CFFh)
90Y3570	A1NW	Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh)

Part number	Feature code	Description
00Y3280 ^a	A3JB	QLogic 2-port 10 Gb Converged Network Adapter (CFFh)
00Y3270 ^b	A3JC	QLogic Eth and 8 Gb Fibre Channel Exp Card (CFFh)

a. Replaces 42C1830.

b. Replaces 44X1940.

For the latest information about the expansion cards supported by the HS22, see the IBM ServerProven website:

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

CFFh expansion cards are installed in the CFFh slot in supported servers, such as the HS22 (Figure 3-37).




Figure 3-37 The HS22 type 7870 showing the CFFh I/O Expansion card position

The CFFh I/O Expansion Card requires a supported high speed I/O module that is installed in bay 7, bay 8, bay 9, or bay 10 of the BladeCenter H or BladeCenter HT chassis. Examples of the types of modules supported in the high speed I/O module bays are the InfiniBand Pass-thru Module and the IBM Virtual Fabric 10Gb Switch Module.

A Multi-switch Interconnect Module can be used in the high speed switch module bays of the BladeCenter H (MSIM) or BladeCenter HT (MSIM-HT) to allow connectivity to standard switch modules.

In a BladeCenter S chassis, the CFFh I/O Expansion Card requires a supported switch module in bay 2. When used in a BladeCenter S chassis, a maximum of two ports are routed from the CFFh I/O Expansion Card to the switch module in bay 2.

For the latest information about the switch modules that are supported with each CFFh I/O expansion card, see *IBM BladeCenter Interoperability Guide*, REDP-BCIG:

<http://www.redbooks.ibm.com/big>

PCI Express I/O Expansion Unit (PEU3e)

The HS22 supports the attachment of one PCI Express I/O Expansion Unit. The PEU3e allows standard PCI-Express type adapters to be attached to the HS22.

For more information about the PEU3e, see 3.10.1, “PCI Express I/O Expansion Unit” on page 401.

PCI Express Gen 2 Expansion Blade (BPE4)

The HS22 supports the attachment of up to four PCI Express Gen 2 Expansion Blades. The BPE4 allows certain standard PCI-Express adapters to be attached to the HS22. If there is a CFFh I/O Expansion Card already installed on the HS22, it must be removed to allow for the attachment of the PCI Express Gen 2 Expansion Blade. The CFFh I/O Expansion Card can then be installed in the topmost PCI Express Gen 2 Expansion Blade. No other type of expansion blade can be mixed with the PCI Express Gen 2 Expansion Blade.

For more information about BPE4, see 3.10.2, “PCI Express Gen 2 Expansion Blade/Expansion Blade II” on page 403.

IBM BladeCenter GPU Expansion Blade (BGE)

The HS22 supports the attachment of up to three IBM BladeCenter GPU Expansion Blades. The BGE is integrated with one NVIDIA Tesla M2070 or Tesla M2070Q GPU adapter that can be attached to the HS22. If there is a CFFh I/O Expansion Card already installed on the HS22, it must be removed to allow for the attachment of the IBM BladeCenter GPU Expansion Blade. The CFFh I/O Expansion Card can then be installed in the topmost IBM BladeCenter GPU Expansion Blade. No other type of expansion blade can be mixed with the IBM BladeCenter GPU Expansion Blade.

For more information about the BGE, see 3.10.3, “IBM BladeCenter GPU Expansion Blade” on page 408.

3.4 BladeCenter HS23 (E5-2600)

The IBM BladeCenter HS23 (E5-2600) is a two-socket blade server running the Intel Xeon processor E5-2600 v2 product family processors. It is ideal for virtualization, hosted client, SAP, and enterprise applications. This blade server is compatible with the IBM BladeCenter E, H, HT, and S.

The following topics are covered in this section:

- ▶ 3.4.1, “Overview” on page 284
- ▶ 3.4.2, “Model information” on page 286
- ▶ 3.4.3, “Chassis support” on page 288
- ▶ 3.4.4, “Processor” on page 289
- ▶ 3.4.5, “Memory” on page 291
- ▶ 3.4.6, “Standard onboard features” on page 295
- ▶ 3.4.7, “Local storage” on page 296
- ▶ 3.4.8, “Integrated virtualization” on page 299
- ▶ 3.4.9, “I/O expansion” on page 299

3.4.1 Overview

The HS23 delivers a robust set of features valued by clients who are focused on performance:

- ▶ Two Intel Xeon processor E5-2600 product family processors
- ▶ Sixteen memory DIMMs
- ▶ Up to 1600 MHz memory speeds
- ▶ Two hot-swap drive bays
- ▶ Internal USB port for embedded hypervisor support
- ▶ Two ports of integrated 1 GbE ports and two ports of integrated 10 GbE with upgrade path to 4x 10 GbE ports
- ▶ Two PCI slots: CIOv and CFFh
- ▶ Up to 18 total I/O ports per blade server (two integrated GbE ports, up to 14 virtual ports and two physical ports with CIOv adapter)
- ▶ Run multiple protocols (FCoE/iSCSI)

Figure 3-38 shows the HS23.



Figure 3-38 The HS23 type 7875

Table 3-49 lists the features of the HS23 (E5-2600).

Table 3-49 Features of the HS23 (7875, E5-2600)

Feature	Description
Processor options	Intel Xeon processor E5-2600 product family (2-, 4-, 6-, or 8-core)
Number of processors (std/max)	1/2
Memory speed	Up to 1600 MHz (dependent on processor model)
Cache	5 MB, 10 MB, 15 MB, or 20 MB (shared between cores)
Memory slots	16 DIMM slots
Memory Maximums	512 GB using 32 GB DIMMs
Memory types	DDR 3 ECC Very Low Profile (VLP) Registered DIMM
Hot swap disk drives (std/max)	0/2 (SAS, SATA, or SSD)
Maximum internal storage	Optional: Up to 2.4 TB with two 2.5-inch SAS hot swap drives Optional: Up to 2 TB with two 2.5-inch NL SAS or NL SATA hot swap drives Optional: Up to 3.2 TB with two 2.5-inch SAS SSDs
Network	Emulex ServerEngines BE3 Dual 10 Gb and Dual 1 Gb
I/O upgrade	Two PCI Express connectors (CFFh and CIOv form factors)

Figure 3-39 shows the components on the system board of the HS23.

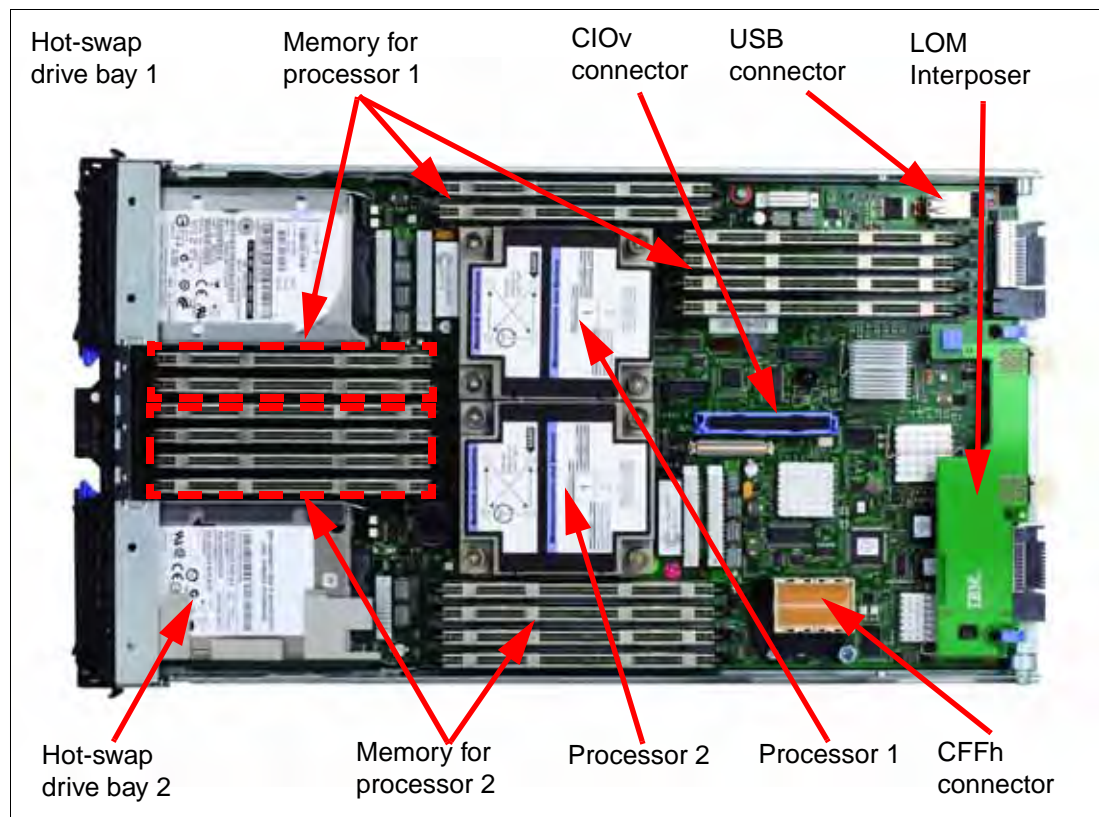


Figure 3-39 Layout of the HS23 system board

3.4.2 Model information

Table 3-50 lists the available models of the HS23.

Table 3-50 Models of the HS23 type 7875

Server models ^a	Intel Xeon processor (two maximum)	Memory (std/ max ^b)	Onboard Disk controller	Drive bays (used /free)	Network	I/O slots (used /max)
7875-A1x	1x E5-2603 4C 1.8 GHz 10 MB 1066 MHz 80 W	1x 4 GB / 512 GB	LSI SAS2004	0/2	2x GbE	0/2
7875-A2x	1x E5-2609 4C 2.4 GHz 10 MB 1066 MHz 80 W	4x 4 GB / 512 GB	LSI SAS2004	0/2	2x GbE	0/2
7875-B1x	1x E5-2620 6C 2.0 GHz 15 MB 1333 MHz 95 W	4x 4 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 2x 10 GbE	1/2 ^c
7875-B2x	1x E5-2640 6C 2.5 GHz 15 MB 1333 MHz 95 W	4x 4 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 2x 10 GbE	1/2 ^c
7875-B3x	1x E5-2630 6C 2.3 GHz 15 MB 1333 MHz 95 W	4x 4 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 2x 10 GbE	1/2 ^c
7875-C1x	1x E5-2650 8C 2.0 GHz 20 MB 1600 MHz 95 W	4x 4 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 2x 10 GbE	1/2 ^c
7875-C6x	1x E5-2650 8C 2.0 GHz 20 MB 1600 MHz 95 W	4x 8 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 2x 10 GbE	1/2 ^c
7875-C7x	1x E5-2660 8C 2.2 GHz 20 MB 1600 MHz 95 W	4x 8 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 2x 10 GbE	1/2 ^c
7875-C3x	1x E5-2665 8C 2.4 GHz 20 MB 1600 MHz 115 W	4x 4 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 2x 10 GbE	1/2 ^c
7875-C4x	1x E5-2670 8C 2.6 GHz 20 MB 1600 MHz 115 W	4x 4 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 2x 10 GbE	1/2 ^c
7875-C8x	1x E5-2680 8C 2.7 GHz 20 MB 1600 MHz 130 W	4x 8 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 2x 10 GbE	1/2 ^c
7875-D1x	1x E5-2650L 8C 1.8 GHz 20 MB 1600 MHz 70 W	4x 4 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 2x 10 GbE	1/2 ^c
7875-F1x	1x E5-2648L 8C 1.8 GHz 20 MB 1600 MHz 70 W	4x 4 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 2x 10 GbE	1/2 ^c
7875-G1x ^d	1x E5-2630 6C 2.3 GHz 15 MB 1333 MHz 95 W	4x 4 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 4x 10 GbE ^d	1/2
7875-G2x ^d	1x E5-2670 8C 2.6 GHz 20 MB 1600 MHz 115 W	4x 4 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 4x 10 GbE ^d	1/2
7875-93x ^e	2x E5-2650 8C 2.0 GHz 20 MB 1600 MHz 95 W	16x 8 GB / 512 GB	LSI SAS2004	0/2	2x GbE + 2x 10 GbE ^e	1/2

a. Model numbers provided are worldwide generally available variant (GAV) model numbers that are not orderable as listed and must be modified by country. The US GAV model numbers use the following nomenclature: xxU. For example, the US orderable part number for 7875-A1x is 7875-A1U. See the product-specific Official IBM announcement letter for other country-specific GAV model numbers.

b. Maximum system memory capacity of 512 GB is based on using 32 GB memory DIMMs.

- c. The 10 Gb Interposer Card in these models provides the same midplane connectors as the adapter installed in the CFFh slot (if any). To use the CFFh slot, you must first remove the Interposer Card.
- d. Models G1x and G2x are optimized for IBM Virtual Fabric and come standard with the Emulex 10GbE Virtual Fabric Adapter Advanced II (CFFh) for HS23 and IBM Virtual Fabric Advanced Software Upgrade (LOM). These adapters enable IBM Virtual Fabric with support for FCoE and iSCSI across four 10 Gb Ethernet ports.
- e. Model 93x is optimized for IBM BladeCenter Foundations for Cloud workloads and includes IBM Virtual Fabric Advanced Software Upgrade (LOM), 90Y9310, which enables FCoE and iSCSI on the onboard two-port 10 Gb controller. They also include an IBM USB Memory Key for VMware ESXi V5.0 (41Y8300) and one year of IBM Systems Director service and support.

3.4.3 Chassis support

The HS23 type 7875 is supported in the BladeCenter chassis, as listed in Table 3-51. Details about this support are in the next subsections.

Table 3-51 HS23 chassis support

Blade	TDP ^a	BC E 8677	BC T 8720/8730	BC S 8886	BC H 8852	BC HT DC 8740	BC HT AC 8750
HS23	115 W or 130 W	No	No	Full	Some limits ^b	Some limits ^b	Some limits ^b
HS23	Up to 95 W	Some limits ^b	No	Full	Some limits ^b	Full	Full

a. Thermal Design Power of the processor in the server. See Table 3-50 to determine to which HS23 server models TDP applies.

b. See Table 3-52 for details.

The specifics of the supported chassis depend on the models of the chassis and the power supplies and cooling modules (blowers) in the chassis, as shown in Table 3-52 on page 289:

- A green cell (■) means that the chassis can be filled with HS22 blade servers up to the maximum number of blade bays in the chassis (14 blades in the BladeCenter H).
- A yellow cell (■) means that the maximum number of HS22 blades that the chassis can hold is fewer than the total available blade bays. *All other bays must remain empty.* Empty bays must be distributed evenly between the two power domains of the chassis (bays 1 - 6 and bays 7 - 14).

Table 3-52 HS23 chassis compatibility

CPU TDP ^a	Maximum number of HS23 supported in each chassis									
	BC-E with AMM (8677) (14 bays)		BC-S (8886) (6 bays)	BC-H (models other than 4Tx)				BC-H (-4Tx)	BC-HT DC (8740) (12 bays) ^b	BC-HT AC (8750) (12 bays) ^b
	2000 W power supplies	2320 W power supplies		2900 W supplies		2980 W supplies ^c		2980 W ^c		
				Std. blower	Enh. blower ^d	Std. blower	Enh. blower ^d	Enh. blower ^d		
Intel Xeon processors										
130 W	None ^e	None ^e	6	None ^e	14	None ^e	14	14	5+5	5+5
115 W	None ^e	None ^e	6	None ^e	14	None ^e	14	14	5+5	5+5
95 W	None ^e	None ^e	6	None ^e	14	None ^e	14	14	12	12
80 W	6+7	14	6	14	14	14	14	14	12	12
70 W	None ^e	None ^e	6	None ^e	14	None ^e	14	14	12	12
60 W	None ^e	None ^e	6	None ^e	14	None ^e	14	14	12	12
Intel Xeon robust thermal profile processors ^f										
95 W	5+7	14	6	14	14	14	14	14	12	12
70 W	14	14	6	14	14	14	14	14	12	12

a. Thermal design power (TDP).

b. Support shown is for non-NEBS environments.

c. IBM BladeCenter H 2980 W AC Power Modules, 68Y6601 (standard in 4Tx, optional with all other BC-H chassis models).

d. IBM BladeCenter H Enhanced Cooling Modules, 68Y6650 (standard in 4Tx, optional with all other BC-H chassis models).

e. Not supported.

f. Intel Xeon E5-2648L (70 W) and E5-2658 (95 W) are robust thermal profile processors used in HS23.

3.4.4 Processor

The HS23 type 7875 features the Intel Xeon processor E5-2600 product family. The new Intel Xeon processor E5-2600 product family has models with four, six, or eight cores per processor, and the memory controller is built into the processor. The memory is directly connected to the processor through four memory controller buses.

The two Intel Xeon processors in the HS23 are connected through a serial coherency link called QuickPath Interconnect (QPI), as shown in Figure 3-40 on page 290. The QPI link is capable of up to 8.0 GTps, depending on the processor model.

Figure 3-40 shows the major components of the HS23 system.

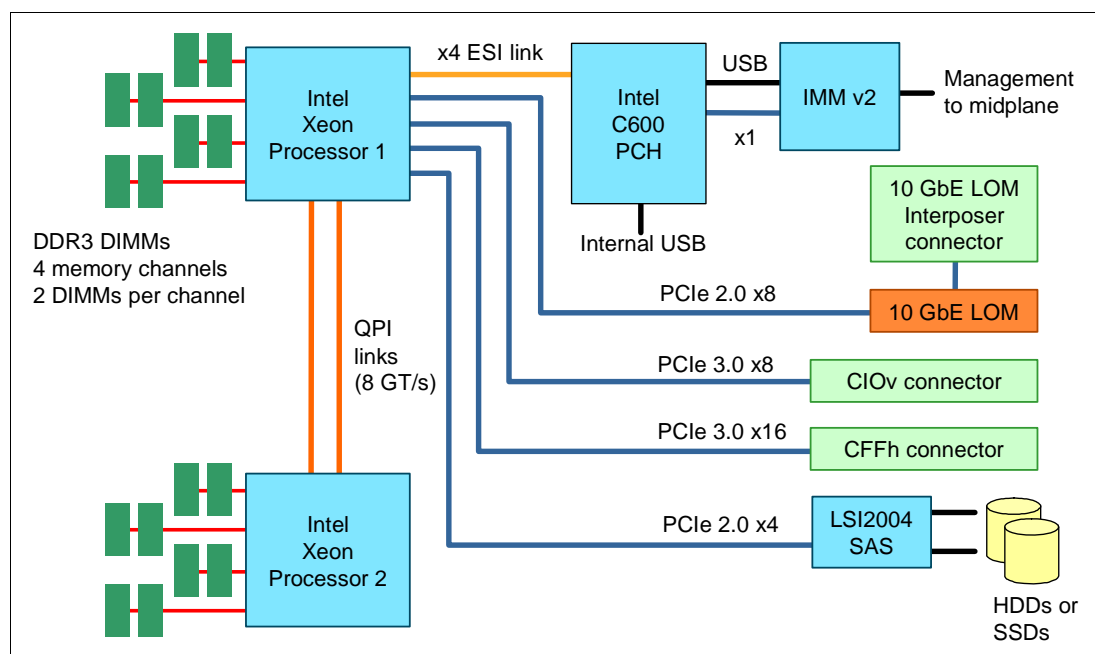


Figure 3-40 HS23 processor block diagram

The processors must be identical in a two-socket configuration. The Intel Xeon processor E5-2600 product family are available in various clock speeds and have standard and lower power offerings.

Table 3-53 lists the processor options for the HS23 and the models that include them (if one exists).

Table 3-53 Available processor options

Part number	Feature code ^a	Description	Supported model
81Y9292	A1S3 / A1SH	Intel Xeon Processor E5-2603 4C 1.8 GHz 10 MB 1066 MHz 80 W	A1x
81Y9294	A1S5 / A1SK	Intel Xeon Processor E5-2609 4C 2.4 GHz 10 MB 1066 MHz 80 W	A2x
81Y9295	A1S6 / A1SL	Intel Xeon Processor E5-2620 6C 2.0 GHz 15 MB 1333 MHz 95 W	B1x, 91x
94Y8572	A23V / A23Y	Intel Xeon Processor E5-2630 6C 2.3 GHz 15 MB 1333 MHz 95 W	B3x, G1x
81Y9304	A1SF / A1SV	Intel Xeon Processor E5-2630L 6C 2.0 GHz 15 MB 1333 MHz 60 W	CTO ^b
94Y8571	A23U/ A23X	Intel Xeon Processor E5-2640 6C 2.5 GHz 15 MB 1333 MHz 95 W	B2x
81Y9298	A1S9 / A1SP	Intel Xeon Processor E5-2650 8C 2.0 GHz 20 MB 1600 MHz 95 W	C1x, 92x
81Y9299	A1SA / A1SQ	Intel Xeon Processor E5-2660 8C 2.2 GHz 20 MB 1600 MHz 95 W	C2x
81Y9300	A1SB / A1SR	Intel Xeon Processor E5-2680 8C 2.7 GHz 20 MB 1600 MHz 130 W	C5x
81Y9302	A1SD / A1ST	Intel Xeon Processor E5-2667 6C 2.9 GHz 15 MB 1600 MHz 130 W	CTO ^b
81Y9305	A1SG / A1SW	Intel Xeon Processor E5-2650L 8C 1.8 GHz 20 MB 1600 MHz 70 W	D1xy
94Y8562	A23R / A240	Intel Xeon Processor E5-2648L 8C 1.8 GHz 20 MB 1600 MHz 70 W	F1x
94Y8565	A23S / A23Z	Intel Xeon Processor E5-2658 8C 2.1 GHz 20 MB 1600 MHz 95 W	CTO ^b

Part number	Feature code ^a	Description	Supported model
94Y8570	A23T / A23W	Intel Xeon Processor E5-2637 2C 3.0 GHz 5 MB 1600 MHz 80 W	CTO ^b
94Y8589	A241 / A242	Intel Xeon Processor E5-2670 8C 2.6 GHz 20 MB 1600 MHz 115 W	C4x, G2x
94Y8671	A2MW / A2MX	Intel Xeon Processor E5-2665 8C 2.4 GHz 20 MB 1600 MHz 115 W	C3x

a. The first feature code is for the first processor. The second feature code is for the second processor.

b. The processor is not available in any of the standard models listed in Table 3-50 on page 286. The processor is, however, available through the configure-to-order (CTO) ordering process.

3.4.5 Memory

The HS23 has eight DIMM sockets per processor (16 DIMMs in total). The HS23 with Intel Xeon processor E5-2600 product family can support up to 512 GB of memory in total when using 32 GB DIMMs and with both processors installed. The HS23 uses Double Data Rate-3 (DDR-3) very-low-profile (VLP) registered DIMMs.

Figure 3-41 shows the location of the 16 memory DIMM sockets




Figure 3-41 DIMM layout of the HS23 system board

The Intel Xeon processor E5-2600 product family has an integrated memory controller. The memory DIMMs connect directly to one of the processors. The processor has four memory channels. With one processor installed, only eight DIMMs can be accessed (two DIMMs per channel). To access memory in the other eight DIMM connectors, the second processor must be installed.

Table 3-54 shows which memory DIMM connectors are connected to each processor.

Table 3-54 HS23 memory channel configuration

Processor	Memory channel	DIMM connector
1	Channel 0	DIMM connector 7 and 8
	Channel 1	DIMM connector 1 and 2
	Channel 2	DIMM connector 3 and 4
	Channel 3	DIMM connector 5 and 6
2	Channel 0	DIMM connector 15 and 16
	Channel 1	DIMM connector 13 and 14
	Channel 2	DIMM connector 9 and 10
	Channel 3	DIMM connector 11 and 12

Table 3-55 lists memory options for the HS23 (ordered by voltage and then DIMM capacity).

Table 3-55 List of memory modules for the HS23 type 7875 using the Xeon E5-2600 series processors

Part number	Feature code	Description	Models where used
Low-power DIMMs (1.35 V)			
46C0560	A0WX	2 GB (1x 2 GB, 1Rx8, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	-
46C0563	A0WY	4 GB (1x 4 GB, 1Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	-
46C0564	A0WZ	4 GB (1x 4 GB, 2Rx8, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	A1x, A2x, B1x, B2x, B3x, D1x, F1x, G1x
00D4981	A3BT	8GB (1x8GB, 1Rx4, 1.35V) PC3L-10600 CL9 ECC DDR3 1333MHz VLP	-
00D4985	A3BU	8GB (1x8GB, 2Rx8, 1.35V) PC3L-10600 CL9 ECC DDR3 1333MHz VLP	-
46C0599	2422	16 GB (1x 16 GB, 2Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	-
00D5008	A3KN	32GB (1x32GB, 4Rx4, 1.35V) PC3L-10600 CL9 ECC DDR3 1333MHz VLP	-
Standard DIMMs (1.5 V)			
90Y3147	A1S0	4 GB (2 Gb, 1Rx4, 1.5 V) PC3-12800 CL11 ECC DDR3 1600 MHz VLP	-
90Y3148	A1S1	4 GB (2 Gb, 2Rx8, 1.5 V) PC3-12800 CL11 ECC DDR3 1600 MHz VLP	C1x, C3x, C4x, G2x
00D4989	A3BV	8GB (1x8GB, 1Rx4, 1.5V) PC3-12800 CL11 ECC DDR3 1600MHz VLP	93x, C7x
00D4993	A3BW	8GB (1x8GB, 2Rx8, 1.5V) PC3-12800 CL11 ECC DDR3 1600MHz VLP	C6x, C8x
90Y3157	A3BS	16GB (1x16GB, 2Rx4, 1.5V) PC3-12800 CL11 ECC DDR3 1600MHz VLP	-

Memory DIMM installation suggestions

The HS23 server supports up to eight DIMMs when one processor is installed and up to 16 DIMMs when two processors are installed. Each processor has four memory channels, and there are two DIMMs per channel.

The following rules apply when selecting the memory configuration:

- ▶ Mixing 1.5 V and 1.35 V DIMMs in the same server is supported. In such a case, all DIMMs operate at 1.5 V.
- ▶ The maximum number of ranks supported per channel is eight.
- ▶ The maximum quantity of DIMMs that can be installed in the server depends on the number of processors, DIMM rank, and operating voltage, as shown in the “Maximum quantity supported” row in Table 3-56.
- ▶ All DIMMs in all processor memory channels operate at the same speed, which is determined as the lowest value of:
 - Memory speed supported by specific processor
 - Lowest maximum operating speed for the selected memory configuration that depends on rated speed, as shown under the “Maximum operating speed” section in Table 3-56.

Table 3-56 Maximum memory speeds for HS23

DIMM type	RDIMM							
DIMM specification								
Rank	Single rank			Dual rank			Quad rank	
Rated speed	1333 MHz		1600 MHz	1333 MHz		1600 MHz	1333 MHz	
Rated voltage	1.35 V		1.5 V	1.35 V		1.5 V	1.35 V	
Operating voltage	1.35 V	1.5 V	1.5 V	1.35 V	1.5 V	1.5 V	1.35 V	1.5 V
Maximum quantity supported	16	16	16	16	16	16	16	16
Maximum DIMM capacity	8 GB	8 GB	8 GB	16 GB	16 GB	16 GB	32 GB	32 GB
Maximum memory capacity	128 GB	128 GB	128 GB	256 GB	256 GB	256 GB	512 GB	512 GB
Max. memory at rated speed	128 GB	128 GB	128 GB	256 GB	256 GB	256 GB	N/A	N/A
Maximum operating speed								
1 DIMM per channel	1333 MHz	1333 MHz	1600 MHz	1333 MHz	1333 MHz	1600 MHz	800 MHz	1066 MHz
2 DIMMs per channel	1333 MHz	1333 MHz	1600 MHz	1333 MHz	1333 MHz	1600 MHz	800 MHz	1066 MHz

Terminology: The term *DIMMs per channel* is commonly referred to as the abbreviation DPC. One DIMM per channel can be referred to as 1DPC and, likewise, two DIMMs per channel can be referred to as 2DPC. The ability to run two DIMMs per channel at 1600 MHz is often referred to as the abbreviation 2DPC@1600MHz.

Memory modes

The HS23 type 7875 supports three memory modes. These modes can be selected in the UEFI BIOS setup:

- Independent channel (normal) mode (see Table 3-57)

Independent channel mode offers a maximum of 128 GB of usable memory with one microprocessor installed, and 256 GB of usable memory with two microprocessors installed (using 16 GB DIMMs). The DIMMs can be installed without matching sizes.

- Memory sparing (rank sparing) mode (see Table 3-57)

If memory sparing (rank sparing) is used, the DIMMs' population order is the same as that for independent channel mode (the DIMMs can be installed without matching sizes). In rank-sparing mode, one rank of a DIMM in each populated channel is reserved as spare memory. The spare rank is not available as active memory. When the active rank memory fails, its content is copied to the spare rank memory, which becomes active. The memory set aside for spare is one rank per channel. The size of a rank varies depending on the DIMMs installed.

Table 3-57 DIMM population sequence for independent channel mode and rank-sparing mode

Processors installed	DIMM connectors
One	1, 7, 3, 5, 2, 8, 4, and 6
Two	1, 16, 7, 14, 3, 9, 5, 11, 2, 15, 8, 13, 4, 10, 6, and 12

- Memory mirrored channel mode (see Table 3-58)

If memory mirroring is used, DIMMs must be installed in pairs (minimum of one pair per processor), and both DIMMs in a pair must be identical in type and size. In mirrored channel mode, the memory contents on channel 0 are duplicated in channel 2, and the memory contents of channel 1 are duplicated in channel 3. The effective memory that is available to the system is only half of the memory installed.

Table 3-58 DIMM population sequence for mirrored channel mode

DIMM pair	DIMM connectors	
	One processor installed	Two processors installed
First	1 and 7	1 and 7
Second	3 and 5	14 and 16
Third	2 and 8	3 and 5
Fourth	4 and 6	9 and 11
Fifth	None	2 and 8
Sixth	None	13 and 15
Seventh	None	4 and 6
Eighth	None	10 and 12
Note: The DIMM's pairs must be identical in size, type, and rank count.		

Important: DIMMs or DIMM fillers must occupy DIMM connectors 1, 2, 13, 14, 15, and 16 for the correct cooling.

3.4.6 Standard onboard features

The HS23 blade server has the standard onboard features shown in the following sections.

UEFI

The HS23 uses an integrated Unified Extensible Firmware Interface (UEFI) next-generation BIOS.

New capabilities include:

- ▶ Human readable event logs. No more beep codes!
- ▶ A complete setup solution by allowing the adapter configuration function to be moved into the UEFI.
- ▶ Complete out-of-band coverage by the Advance Settings Utility to simplify remote setup.

Using all of the features that UEFI offers requires a UEFI-aware operating system and adapters. UEFI is fully compatible with earlier versions of the BIOS.

For more information about UEFI, see *Introducing uEFI-Compliant Firmware on IBM System x and BladeCenter Servers*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5083207>

Onboard network adapter

The HS23 offers two integrated Gigabit Ethernet ports and two integrated 10 Gb Ethernet ports with the integrated Emulex BladeEngine 3 (BE3) controller. Two Gigabit Ethernet ports are routed to the chassis I/O bays 1 and 2, and two 10 Gb Ethernet ports are routed to the chassis I/O bays 7 and 9 using either a 10 Gb LOM Interposer Card or Emulex Virtual Fabric Adapter II for HS23.

The integrated BE3 4-port NIC has the following features:

- ▶ Two Gigabit Ethernet ports and two 10 Gb Ethernet ports (1 Gb and 10 Gb auto-negotiation)
- ▶ Full-duplex (FDX) capability
- ▶ 10 Gb ports operate in either a virtual NIC (vNIC) or physical NIC (pNIC) mode:
 - vNIC mode: Up to six vNICs (up to three vNICs per one 10 Gb port):
 - Virtual Fabric mode or Switch Independent operational mode.
 - Virtual port bandwidth allocation in 100 Mbps increments.
 - Up to two vNICs can be configured as an iSCSI or FCoE vNICs (one per port) with optional Advanced Upgrade (90Y9310).
 - pNIC mode: Dual-port 1/10 Gb Ethernet adapter
- ▶ IPv4/IPv6 offload:
 - TCP and UDP checksum offload
 - Large send offload (LSO)
 - Large receive offload (LRO)
 - Receive side scaling (RSS)
- ▶ IPv4 TCP Chimney Offload
- ▶ IEEE 802.1Q VLAN tagging
- ▶ VLAN insertion and extraction
- ▶ Jumbo frames up to 9000 bytes

- ▶ Load balancing and failover teaming support, including adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), and IEEE 802.3ad
- ▶ Enhanced Ethernet (draft) support:
 - Enhanced Transmission Selection (ETS) (P802.1Qaz)
 - Priority-based Flow Control (PFC) (P802.1Qbb)
 - Data Center Bridging eXchange Protocol (DCBX) (P802.1Qaz)
- ▶ Supports Serial over LAN (SoL) and concurrent KVM (cKVM)
- ▶ Preboot Execution Environment (PXE) support
- ▶ Wake On LAN support
- ▶ PCIe x8 Gen 2 host interface
- ▶ Message Signal Interrupt (MSI-X) support

Table 3-59 shows the ordering part number and feature code for Advanced Upgrade for integrated NIC to enable FCoE or iSCSI protocols in it.

Table 3-59 Integrated NIC upgrade ordering part number and feature code

Part number	Feature code	Description
90Y9310	A2TD	IBM Virtual Fabric Advanced Software Upgrade (LOM)

Important: A 10 Gb Interposer card for IBM BladeCenter HS23 (or 10 Gb LOM Interposer Card) part number 94Y8550 must be installed if you plan to use the 10 Gb ports of the onboard network adapter. For details, see “LOM Interposer Card” on page 300.

Integrated systems management processor

The HS23 type 7875 blade server includes an Integrated Management Module Version 2 (IMM V2) to monitor server availability, perform Predictive Failure Analysis, and trigger IBM Systems Director alerts. IMM V2 provides system management capabilities with IPMI 2.0 compliance and Serial over LAN.

For more information about the IMM, see *Transitioning to UEFI and IMM*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5079769>

Video controller

The video controller is a Matrox G200 with 16 MB integrated into IMM v2. The maximum resolution is 1600x1200 at 75 Hz with 16 M colors (32 bits per pixel).

3.4.7 Local storage

The HS23 blade server features an onboard LSI Logic LSISAS2004 SAS controller with two hot-swap drives bays that are accessible from the front of the blade server (Figure 3-42 on page 297). The LSI Logic LSISAS2004 SAS controller provides RAID 0 or RAID 1 capability and supports up to two internal hot-swap SAS or SATA HDDs or two internal hot-swap solid-state drives. The LSISAS2004 SAS controller ports operate at a burst rate of 6.0 Gbps.




Figure 3-42 The HS23 showing the front hot-swap disk drive bays

Local SAS or SATA disk drives

The HS23 type 7875 has support for up to two hot-swap SFF SAS or SATA hard disk drives. These two hot-swap hard disk drives are accessible from the front of the blade server without removing the blade server from the BladeCenter chassis. See Table 3-60 for a list of supported SAS and SATA HDDs.

Table 3-60 Supported SAS and SATA HDDs

Part number	Feature code	Description
10K SAS hard disk drives		
90Y8877	A2XC	IBM 300GB 10K 6Gbps SAS 2.5" SFF G2HS HDD
42D0637	5599	IBM 300GB 10K 6Gbps SAS 2.5" SFF Slim-HS HDD
90Y8872	A2XD	IBM 600GB 10K 6Gbps SAS 2.5" SFF G2HS HDD
49Y2003	5433	IBM 600GB 10K 6Gbps SAS 2.5" SFF Slim-HS HDD
00AJ300	A4VB	IBM 600GB 15K 6Gbps SAS 2.5" G2HS HDD
81Y9650	A282	IBM 900GB 10K 6Gbps SAS 2.5" SFF HS HDD
00AD075	A48S	IBM 1.2TB 10K 6Gbps SAS 2.5" G2HS HDD
15K SAS hard disk drives		
90Y8926	A2XB	IBM 146GB 15K 6Gbps SAS 2.5" SFF G2HS HDD
42D0677	5536	IBM 146GB 15K 6Gbps SAS 2.5" SFF Slim-HS HDD
81Y9670	A283	IBM 300GB 15K 6Gbps SAS 2.5" SFF HS HDD
SAS hybrid hard disk drives		
00AJ236	A4VD	IBM 300GB 15K 6Gbps SAS 2.5" G2HS Hybrid
00AD102	A4G7	IBM 600GB 10K 6Gbps SAS 2.5" G2HS Hybrid
00AJ246	A4VF	IBM 600GB 15K 6Gbps SAS 2.5" G2HS Hybrid
NL SAS hard disk drives		
90Y8953	A2XE	IBM 500GB 7.2K 6Gbps NL SAS 2.5" SFF G2HS HDD
42D0707	5409	IBM 500GB 7200 6Gbps NL SAS 2.5" SFF HS HDD
81Y9690	A1P3	IBM 1TB 7.2K 6Gbps NL SAS 2.5" SFF HS HDD

Part number	Feature code	Description
NL SATA hard disk drives		
81Y9722	A1NX	IBM 250GB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD
81Y9726	A1NZ	IBM 500GB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD
81Y9730	A1AV	IBM 1TB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD

Solid-state drives

SSDs are more reliable than spinning hard disk drives and they consume much less power than a standard SAS drive. Target Applications for SSDs include video surveillance, transaction-based databases, and other applications that have high performance but moderate space requirements.

Table 3-61 lists the supported SSDs.

Table 3-61 Supported solid-state drives (SSDs)

Part number	Feature code	Description
Solid-state drives - Enterprise		
00W1125	A3HR	IBM 100GB SATA 2.5" MLC HS Enterprise SSD
43W7718	A2FN	IBM 200GB SATA 2.5" MLC HS SSD
49Y6129	A3EW	IBM 200GB SAS 2.5" MLC HS Enterprise SSD
49Y6134	A3EY	IBM 400GB SAS 2.5" MLC HS Enterprise SSD
49Y6139	A3F0	IBM 800GB SAS 2.5" MLC HS Enterprise SSD
49Y6195	A4GH	IBM 1.6TB SAS 2.5" MLC HS Enterprise SSD
41Y8331	A4FL	S3700 200GB SATA 2.5" MLC HS Enterprise SSD
41Y8336	A4FN	S3700 400GB SATA 2.5" MLC HS Enterprise SSD
41Y8341	A4FQ	S3700 800GB SATA 2.5" MLC HS Enterprise SSD
Solid-state drives - Enterprise value		
49Y5839	A3AS	IBM 64GB SATA 2.5" MLC HS Enterprise Value SSD
90Y8648	A2U4	IBM 128GB SATA 2.5" MLC HS Enterprise Value SSD
90Y8643	A2U3	IBM 256GB SATA 2.5" MLC HS Enterprise Value SSD
49Y5844	A3AU	IBM 512GB SATA 2.5" MLC HS Enterprise Value SSD
00AJ000	A4KM	S3500 120GB SATA 2.5" MLC HS Enterprise Value SSD
00AJ005	A4KN	S3500 240GB SATA 2.5" MLC HS Enterprise Value SSD
00AJ010	A4KP	S3500 480GB SATA 2.5" MLC HS Enterprise Value SSD
00AJ015	A4KQ	S3500 800GB SATA 2.5" MLC HS Enterprise Value SSD

PCIe SSD adapters

The HS23 server supports the High IOPS SSD adapters listed in Table 3-81. The adapters must be installed in an IBM BladeCenter PCI Express Gen 2 Expansion Blade II. Up to eight

High IOPS adapters supported per one HS23 (two per PCI Express Gen 2 Expansion Blade II and up to four PCI Expansion Blades per HS23)

Table 3-62 Supported PCIe SSD adapters

Part number	Feature code	Description	Slots supported
46C9078	A3J3	IBM 365GB High IOPS MLC Mono Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
46C9081	A3J4	IBM 785GB High IOPS MLC Mono Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4377	A3DY	IBM 1.2TB High IOPS MLC Mono Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4397	A3DZ	IBM 2.4TB High IOPS MLC Duo Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4361	A3MZ	IBM 300GB High IOPS MLC Modular Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4365	A3N0	IBM 600GB High IOPS MLC Modular Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4369	A3N1	IBM 800GB High IOPS MLC Modular Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4373	A3N2	IBM 300GB High IOPS SLC Modular Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)

3.4.8 Integrated virtualization

The HS23 offers an IBM standard USB flash drive option preinstalled with VMware ESXi. This is an embedded version of VMware ESXi. It is fully contained on the flash drive, without requiring any disk space. The IBM USB Memory Key for VMware Hypervisor plugs into the USB port at the HS23 system board.

Table 3-63 lists the part number, feature code, and description of the IBM USB Memory Key for VMware Hypervisor that is supported in the HS23.

Table 3-63 IBM USB Memory Key for VMware Hypervisor

Part Number	Feature Code	Description
41Y8300	A2VC	IBM USB Memory Key for VMWare ESXi 5.0
41Y8307	A383	IBM USB Memory Key for VMWare ESXi 5.0 Update 1
41Y8311	A2R3	IBM USB Memory Key for VMWare ESXi 5.1
41Y8382	A4WZ	IBM USB Memory Key for VMWare ESXi 5.1 Update 1
41Y8298	A2G0	IBM Blank USB Memory Key for VMWare ESXi Downloads

For a complete description of the features and capabilities of VMware ESX Server, go to the following address:

<http://www.vmware.com/products/vi/esx/>

3.4.9 I/O expansion

The HS23 type 7875 is able to connect to a wide variety of networks and fabrics by installing the appropriate I/O expansion card. Supported networks and fabrics include 1 Gb and 10 Gb Ethernet, 4 Gb and 8 Gb Fibre Channel, Serial Attached SCSI (SAS), and InfiniBand.

An HS23 blade server with an installed I/O expansion card must be installed in a BladeCenter chassis with switch modules (or pass-through) that are compatible with the I/O expansion card.

The HS23 type 7875 blade server supports the following formats of I/O expansion cards: the CIOv format and the CFFh format, as described in the following sections.

LOM Interposer Card

The 10 Gb Interposer Card for IBM BladeCenter HS23 provides I/O connections between the 10 Gb ports of an integrated network adapter (see “Onboard network adapter” on page 295) and high speed switch modules in bays 7 and 9 of a BladeCenter H or BladeCenter HT chassis.

Figure 3-43 shows a 10 Gb LOM Interposer Card installed in the slot of the HS23.




Figure 3-43 The HS23 type 7875 showing the 10 Gb LOM Interposer Card position

Table 3-64 shows which HS23 models include the 10 GB LOM Interposer Cards.

Table 3-64 HS23 with LOM Interposer Cars included as standard

Blade	7875-A1x	7875-A2x	7875-B1x	7875-B2x	7875-B3x	7875-C1x	7875-C2x	7875-C3x	7875-C4x	7875-C5x	7875-D1x	7875-F1x	7875-G1x	7875-G2x	7875-91x	7875-92x
10 Gb LOM Interposer Card	N	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	Y	Y

Table 3-65 lists the 10 Gb Interposer Card for IBM BladeCenter HS23 ordering part number and feature code.

Table 3-65 Ordering part number and feature code

Part number	Feature code	Description
94Y8550	A244	10 Gb Interposer Card for IBM BladeCenter HS23

CIOv

The CIOv I/O expansion connector provides I/O connections to modules in bays 3 and 4 of a supported BladeCenter chassis. The CIOv slot is a third-generation PCI Express x8 (PCIe 3.0 x8) slot. A maximum of one CIOv I/O expansion card is supported per single-wide blade server. A CIOv I/O expansion card can be installed on a blade server at the same time as a CFFh I/O expansion card is installed. Table 3-66 lists the CIOv expansion cards that are supported in the HS23 type 7875.

Table 3-66 Supported CIOv expansion cards for the HS23 server

Part number	Feature code	Description
44W4475	1039	Ethernet Expansion Card (CIOv)
46M6065	3594	QLogic 4 Gb Fibre Channel Expansion Card (CIOv)
43W4068	1041	SAS Connectivity Card (CIOv)
46M6140	3598	Emulex 8 Gb Fibre Channel Expansion Card (CIOv)
44X1945	1462	QLogic 8 Gb Fibre Channel Expansion Card (CIOv)

For the latest information about the expansion cards supported by the HS23, see the IBM ServerProven website:

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

CIOv expansion cards are installed in the CIOv slot in the HS23 (Figure 3-44).



Figure 3-44 The HS23 type 7875 showing the CIOv I/O expansion card position

The CIOv expansion cards require a supported I/O module installed in bay 3 and bay 4 of the chassis where the cards and servers are installed.

CFFh

The CFFh I/O expansion connector provides I/O connections to high speed switch modules in bays 7, 8, 9, and 10 of a BladeCenter H or BladeCenter HT chassis or to switch bay 2 in a BladeCenter S chassis. The CFFh slot is a third-generation PCI Express x16 (PCIe 3.0 x16) slot. A maximum of one CFFh I/O expansion card is supported per blade server. A CFFh I/O expansion card can be installed on a blade server at the same time as a CIOv I/O expansion card is installed. Table 3-67 lists the CFFh expansion cards that are supported in the HS23 type 7875.

Table 3-67 Supported CFFh expansion cards for the HS23 type 7875 server

Part number	Feature code	Description
46M6001	0056	2-port 40 Gb InfiniBand Expansion Card (CFFh)
44W4479	5476	2/4 Port Ethernet Expansion Card (CFFh)
46M6164	0098	Broadcom 10 Gb Gen2 4-port Ethernet Exp Card (CFFh)
81Y3133	A1QR	Broadcom 2-port 10 Gb Virtual Fabric Adapter
81Y1650	5437	Brocade 2 port 10 GbE Converged Network Adapter
81Y3120	A287	Emulex 10 GbE VFA II for HS23
90Y9332	A2ZN	Emulex 10 GbE VFA Advanced II for HS23
2C1810	3593	Intel 10 Gb 2-port Ethernet Expansion Card (CFFh)
90Y3570	A1NW	Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh)
00Y3280 ^a	A3JB	QLogic 2-port 10 Gb Converged Network Adapter (CFFh)
00Y3332	A4AC	QLogic 10Gb Virtual Fabric Adapter (CFFh)
00Y5618	A4AD	QLogic 10 Gb Virtual Fabric Converged Network Adapter (CFFh)
00Y3270 ^b	A3JC	QLogic Eth and 8 Gb Fibre Channel Exp Card (CFFh)

a. Replaces 42C1830.

b. Replaces 44X1940.

For the latest information about the expansion cards supported by the HS23, see the IBM ServerProven website:

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

CFFh expansion cards are installed in the CFFh slot in the HS23 (Figure 3-45).




Figure 3-45 The HS23 type 7875 showing the CFFh I/O expansion card position

The CFFh I/O expansion card requires that a supported high speed I/O module is installed in bay 7, bay 8, bay 9, or bay 10 of the BladeCenter H or BladeCenter HT chassis. Examples of the types of modules supported in the high speed I/O module bays are the InfiniBand Pass-thru Module and the IBM Virtual Fabric 10Gb Switch Module.

A Multi-switch Interconnect Module can be used in the high speed switch module bays of the BladeCenter H (MSIM) or BladeCenter HT (MSIM-HT) to allow connectivity to standard switch modules.

In a BladeCenter S chassis, the CFFh I/O Expansion Card requires a supported switch module in bay 2. When used in a BladeCenter S chassis, a maximum of two ports are routed from the CFFh I/O Expansion Card to the switch module in bay 2.

For the latest information about the switch modules that are supported with each CFFh I/O expansion card, see the *IBM BladeCenter Interoperability Guide*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5073016>

Emulex 10 GbE VFA II and VFA II Advanced (CFFh) for HS23

The Emulex 10 GbE VFA II CFFh adapters for HS23 use a CFFh I/O expansion connector to provide I/O connections to high speed switch modules in bays 8 and 10 of a BladeCenter H or BladeCenter HT chassis. These adapters provide I/O connections between 10 Gb ports of an integrated network adapter (see “Onboard network adapter” on page 295) and high speed switch modules in bays 7 and 9 of a BladeCenter H or BladeCenter HT chassis.

An Emulex 10 GbE VFA II for HS23 CFFh I/O expansion card replaces a 10 Gb LOM Interposer Card if you need 2+2 10 Gb ports (total four ports). A maximum of one VFA II for HS23 is supported per blade server.

The Emulex 10GbE VFA II cards require that a supported high speed I/O module is installed in bay 7, bay 8, bay 9, or bay 10 of the BladeCenter H or BladeCenter HT chassis. Examples of the types of modules supported in the high speed I/O module bays are the IBM Virtual Fabric 10 Gb Switch Module and the Cisco Nexus 4001I Switch Module.

For more information about this card, see 3.1.9, “Emulex 10 GbE Virtual Fabric Adapter II/Advanced II for HS23” on page 231.

PCI Express Gen 2 Expansion Blade II

The HS23 supports the attachment of PCI Express Gen 2 Expansion Blades II. This unit allows certain standard PCI-Express adapters to be attached to the HS23. If there is a CFFh I/O expansion card already installed on the HS23, it must be removed to allow for the attachment of the PCI Express Gen 2 Expansion Blade II. The CFFh I/O expansion card can then be installed in the topmost PCI Express Gen 2 Expansion Blade. No other type of expansion blade can be mixed with the PCI Express Gen 2 Expansion Blade.

For more information about the BPE4, see 3.10.2, “PCI Express Gen 2 Expansion Blade/Expansion Blade II” on page 403.

IBM BladeCenter GPU Expansion Blade II

The HS23 supports the attachment of up to four IBM BladeCenter GPU Expansion Blades II. This unit is integrated with one NVIDIA Tesla M2075, NVIDIA Tesla M2070Q, or NVIDIA Tesla M2090 GPU adapter that can be attached to the HS23. If there is a CFFh I/O expansion card already installed on the HS23, it must be removed to allow for the attachment of the IBM BladeCenter GPU Expansion Blade. The CFFh I/O expansion card can then be installed in the topmost IBM BladeCenter GPU Expansion Blade. No other type of expansion blade can be mixed with the IBM BladeCenter GPU Expansion Blade.

For more information about the BGE, see 3.10.3, “IBM BladeCenter GPU Expansion Blade” on page 408.

3.5 BladeCenter HS23 (E5-2600 v2)

The IBM BladeCenter HS23 (E5-2600 v2) is a next-generation two-socket blade server running the Intel Xeon processor E5-2600 v2 product family. It is ideal for virtualization, hosted client, SAP, and enterprise applications. This blade server is compatible with the IBM BladeCenter E, H, HT, and S.

The topics covered in this section are:

- ▶ 3.5.1, “Overview”
- ▶ 3.5.2, “Model information” on page 306
- ▶ 3.5.3, “Chassis support” on page 307
- ▶ 3.5.4, “Processor” on page 308
- ▶ 3.5.5, “Memory” on page 310
- ▶ 3.5.6, “Standard onboard features” on page 314
- ▶ 3.5.7, “Local storage” on page 316
- ▶ 3.5.8, “Integrated virtualization” on page 318
- ▶ 3.5.9, “I/O expansion” on page 319

3.5.1 Overview

The HS23 (E5-2600 v2) delivers a robust set of features valued by clients who are focused on performance:

- ▶ Two Intel Xeon processor E5-2600 v2 product family processors
- ▶ 16 memory DIMMs
- ▶ Up to 1866 MHz memory speeds

- ▶ Two hot-swap drive bays
- ▶ Internal USB port for embedded hypervisor support
- ▶ Two ports of integrated 1 GbE ports and two ports of integrated 10 GbE with upgrade path to 4x 10 GbE ports
- ▶ Two PCI slots: CIOv and CFFh
- ▶ Up to 18 total I/O ports per blade server (two integrated GbE ports, up to 14 virtual ports and two physical ports with CIOv adapter)
- ▶ Run multiple protocols (FCoE/iSCSI)

Figure 3-46 shows the HS23.



Figure 3-46 The HS23 type 7875

Table 3-68 lists the features of the HS23 (E5-2600 v2).

Table 3-68 Features of the HS23 (7875, E5-2600 v2)

Feature	Description
Processor options	Intel Xeon processor E5-2600 v2 product family (4-, 6-, 8, 10-, 12-core)
Number of processors (std/max)	1 / 2
Memory speed	Up to 1866 MHz (dependent on processor model)
Cache	10 MB, 15 MB, 20 MB, 25 MB or 30 MB (shared between cores)
Memory slots	16 DIMM slots
Memory Maximums	512 GB using 32 GB DIMMs
Memory types	DDR 3 ECC Very Low Profile (VLP) Registered DIMM
Hot swap disk drives (std/max)	0 / 2 (SAS, SATA, or SSD)
Maximum internal storage	Optional: Up to 2.4 TB with two 2.5-inch SAS hot swap drives Optional: Up to 2 TB with two 2.5-inch NL SAS or NL SATA hot swap drives Optional: Up to 3.2 TB with two 2.5-inch SAS SSDs
Network	Emulex ServerEngines BE3 Dual 10 Gb and Dual 1 Gb
I/O upgrade	Two PCI Express connectors (CFFh and CIOv form factors)

Figure 3-47 shows the components on the system board of the HS23.

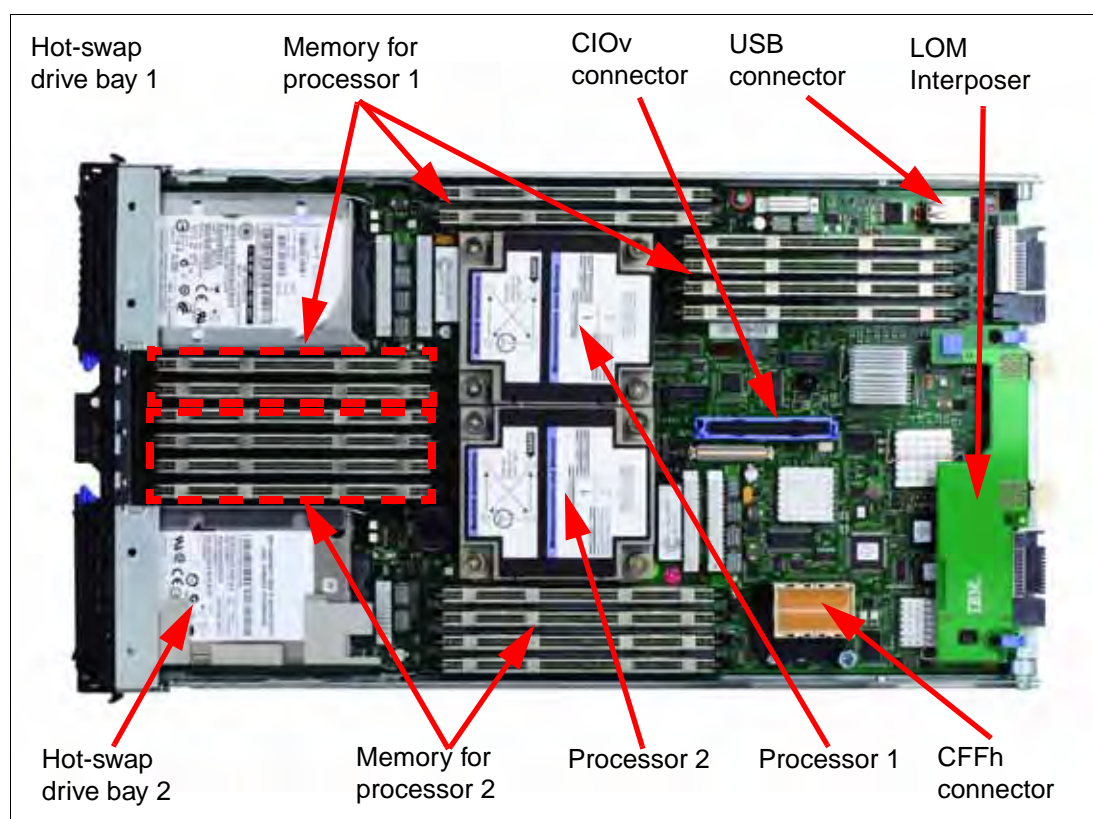


Figure 3-47 Layout of the HS23 system board

3.5.2 Model information

Table 3-69 lists the available models of the HS23 (E5-26xx v2).

Table 3-69 Models of the HS23 (E5-26xx v2) type 7875

Server models ^a	Intel Xeon processor (two maximum)	Memory (std/ max ^b)	Onboard Disk controller	Drive bays (used /free)	Network	I/O slots (used /max)
7875-A4x	1x E5-2603 v2 4C 1.8GHz 10MB 1333MHz 80W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE	0 / 2
7875-A3x	1x E5-2609 v2 4C 2.5GHz 10MB 1333MHz 80W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE	0 / 2
7875-B4x	1x E5-2620 v2 6C 2.1GHz 15MB 1600MHz 80W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE + 2x 10 GbE	1 / 2 ^c
7875-B5x	1x E5-2630 v2 6C 2.6GHz 15MB 1600MHz 80W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE + 2x 10 GbE	1 / 2 ^c
7875-B6x	1x E5-2640 v2 8C 2.0GHz 20MB 1600MHz 95W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE + 2x 10 GbE	1 / 2 ^c
7875-C9x	1x E5-2650 v2 8C 2.6GHz 20MB 1866MHz 95W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE + 2x 10 GbE	1 / 2 ^c

Server models ^a	Intel Xeon processor (two maximum)	Memory (std/ max ^b)	Onboard Disk controller	Drive bays (used /free)	Network	I/O slots (used /max)
7875-CAx	1x E5-2660 v2 10C 2.2GHz 25MB 1866MHz 95W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE + 2x 10 GbE	1 / 2 ^c
7875-CBx	1x E5-2670 v2 10C 2.5GHz 25MB 1866MHz 115W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE + 2x 10 GbE	1 / 2 ^c
7875-CCx	1x E5-2680 v2 10C 2.8GHz 25MB 1866MHz 115W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE + 2x 10 GbE	1 / 2 ^c
7875-D2x	1x E5-2650L v2 10C 1.7GHz 25MB 1600MHz 70W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE + 2x 10 GbE	1 / 2 ^c
7875-F2x	1x E5-2697 v2 12C 2.7GHz 30MB 1866MHz 130W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE + 2x 10 GbE	1 / 2 ^c
7875-G3x	1x E5-2628L v2 8C 1.9GHz 20MB 1600MHz 70W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE + 4x 10 GbE ^d	1 / 2 ^c
7875-G4x	1x E5-2658 v2 10C 2.4GHz 25MB 1866MHz 95W	1x 8 GB / 512 GB	LSI SAS2004	0 / 2	2x GbE + 4x 10 GbE ^d	1 / 2 ^c

a. Model numbers provided are worldwide generally available variant (GAV) model numbers that are not orderable as listed and must be modified by country. The US GAV model numbers use the following nomenclature: xxU. For example, the US orderable part number for 7875-A4x is 7875-A4U. See the product-specific IBM announcement letter for other country-specific GAV model numbers.

b. Maximum system memory capacity of 512 GB is based on using 32 GB memory DIMMs.

c. The 10 Gb Interposer Card in these models provides the same midplane connectors as the adapter installed in the CFFh slot (if any). To use the CFFh slot, you must first remove the Interposer Card.

d. These models come standard with Emulex 10GbE VFA Advanced II for IBM BladeCenter HS23, 90Y9332 and IBM Virtual Fabric Advanced Software Upgrade (LOM), 90Y9310.

3.5.3 Chassis support

The HS23 (7875, E5-2600 v2) chassis support depends on the type and models of the chassis and the power supplies and blowers in the chassis, as shown in Table 3-70 on page 308.

BladeCenter E considerations:

- ▶ The HS23 is not supported in BladeCenter E chassis with power supplies that have a capacity less than 2000 W.
- ▶ The HS23 requires an advanced management module to be installed in the BladeCenter E.

Table 3-70 HS23 (E5-2600 v2) chassis compatibility

CPU TDP ^a	Maximum number of HS23 supported in each chassis								
	BC-E with AMM (8677) (14 bays)		BC-S (8886) (6 bays)	BC-H (14 bays)				BC-HT DC (8740) (12 bays) ^b	BC-HT AC (8750) (12 bays) ^b
	2000 W power supplies	2320 W power supplies		2900 W power supplies		2980 W power supplies ^c			
				Std. blower	Enh. blower ^d	Std. blower	Enh. blower ^d		
Intel Xeon processors									
130 W	None ^f	None ^f	6	None ^f	14 ^e	None ^f	14 ^e	5+5	5+5
115 W	None ^f	None ^f	6	None ^f	14	None ^f	14	5+5	5+5
95 W	None ^f	None ^f	6	None ^f	14	None ^f	14	12	12
80 W	6+7 ^g	14 ^g	6	14 ^g	14	14 ^g	14	12	12
70 W	None ^f	None ^f	6	None ^f	14	None ^f	14	12	12
60 W	None ^f	None ^f	6	None ^f	14	None ^f	14	12	12
Intel Xeon robust thermal profile processors ^h									
95 W	5+7 ^{g,i}	14 ^{g,i}	6	14 ^{g,i}	14	14 ^{g,i}	14	12	12
70 W	14 ^g	14 ^g	6	14 ^g	14	14 ^g	14	12	12
50 W	14	14	6	14	14	14	14	12	12

a. TDP = Thermal Design Power.

b. Support shown is for non-NEBS (Enterprise) environments.

c. IBM BladeCenter H 2980 W AC Power Modules, 68Y6601 (standard in 4Tx, 5Tx, and 9xx BC-H chassis models; optional with all other BC-H chassis models).

d. IBM BladeCenter H Enhanced Cooling Modules, 68Y6650 (standard in 4Tx, 5Tx, and 9xx BC-H chassis models; optional with all other BC-H chassis models).

e. E5-2697 v2 and E5-2690 v2 processors throttle at Steady State at ambient temperature of 31 °C in the BladeCenter H.

f. Not supported.

g. When one blower fails, the HS23 (7875, E5-2600 v2) with specified processor TDP only supports ambient temperature of up to 28 °C when installed in the BladeCenter H with the standard blowers or BladeCenter E.

h. Intel Xeon E5-2618L v2 (50 W), E5-2628L v2 (70 W), E5-2648L v2 (70 W), and E5-2658 v2 (95 W) are robust thermal profile processors used in HS23 (7875, E5-2600 v2).

i. The HS23 (7875, E5-2600 v2) with the Intel Xeon processor E5-2658 v2 (95 W) only supports one DIMM per channel when installed in the BladeCenter H with the standard blowers or BladeCenter E.

3.5.4 Processor

The HS23 (E5-2600 v2) features the Intel Xeon processor E5-2600 v2 product family. The new Intel Xeon processor E5-2600 v2 product family has four, six, eight, ten, or twelve cores per processor, and the memory controller is built into the processor. The memory is directly connected to the processor through four memory controller buses.

The two Intel Xeon processors in the HS23 (E5-2600 v2) are connected through a serial coherency link called QuickPath Interconnect (QPI), as shown in Figure 3-48 on page 309. The QPI link is capable of up to 8.0 GTps, depending on the processor model.

Figure 3-48 shows the major components of the HS23 (E5-2600 v2) system.

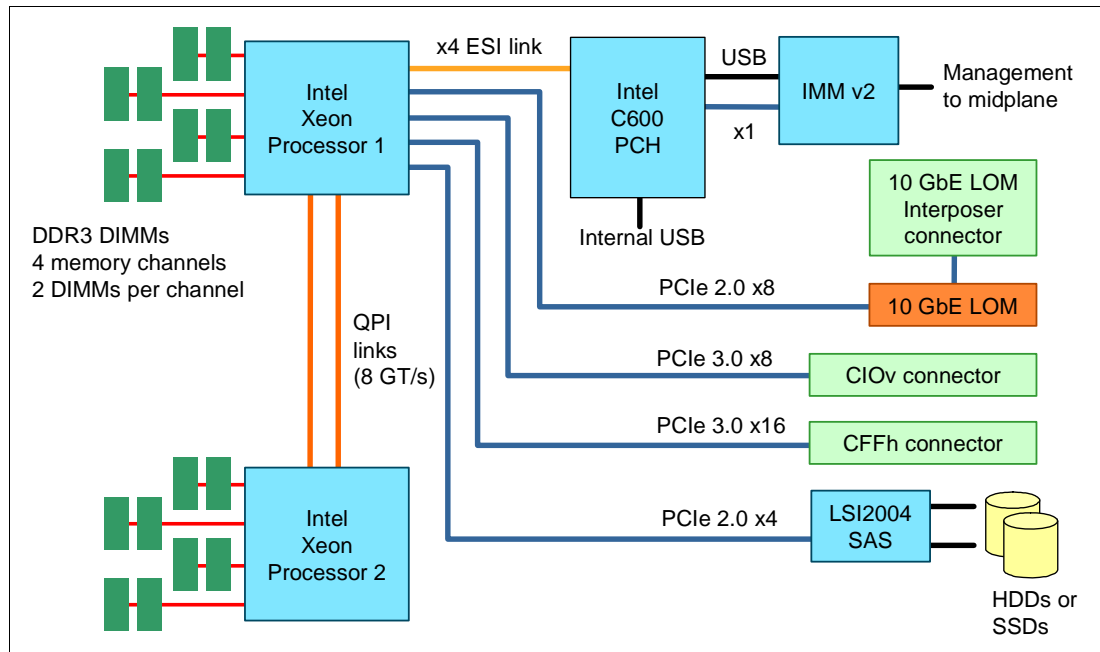


Figure 3-48 HS23 processor block diagram

The processors must be identical in a two-socket configuration. The Intel Xeon processor E5-2600 v2 product family are available in various clock speeds and have standard and lower power offerings.

Table 3-71 lists the processor options for the HS23 (E5-2600 v2) and the models that include them (if one exists).

Table 3-71 Available processor options

Part number	Feature code ^a	Description	Supported model
94Y5260	A4T5 / A4UR	Intel Xeon Processor E5-2603 v2 4C 1.8GHz 10MB 1333MHz 80W	A4x
94Y5261	A4SU / A4US	Intel Xeon Processor E5-2609 v2 4C 2.5GHz 10MB 1333MHz 80W	A3x
94Y5278	A4SQ / A4UT	Intel Xeon Processor E5-2618L v2 6C 2.0GHz 15MB 1333MHz 50W ^b	CTO ^c
94Y5262	A4SV / A4UU	Intel Xeon Processor E5-2620 v2 6C 2.1GHz 15MB 1600MHz 80W	B4x
94Y5279	A4SR / A4UV	Intel Xeon Processor E5-2628L v2 8C 1.9GHz 20MB 1600MHz 70W ^b	G3x
94Y5263	A4SW / A4UW	Intel Xeon Processor E5-2630 v2 6C 2.6GHz 15MB 1600MHz 80W	B5x
94Y5274	A4SL / A4UX	Intel Xeon Processor E5-2630L v2 6C 2.4GHz 15MB 1600MHz 60W	CTO ^c
94Y5271	A4T3 / A4UY	Intel Xeon Processor E5-2637 v2 4C 3.5GHz 15MB 1866MHz 130W	CTO ^c
94Y5264	A4SX / A4UZ	Intel Xeon Processor E5-2640 v2 8C 2.0GHz 20MB 1600MHz 95W	B6x
94Y5281	A4SS / A4V1	Intel Xeon Processor E5-2648L v2 10C 1.9GHz 25MB 1866MHz 70W ^b	CTO ^c
94Y5265	A4SY / A4V2	Intel Xeon Processor E5-2650 v2 8C 2.6GHz 20MB 1866MHz 95W	C9x
94Y5275	A4SM / A4V3	Intel Xeon Processor E5-2650L v2 10C 1.7GHz 25MB 1600MHz 70W	D2x
94Y5282	A4ST / A4V4	Intel Xeon Processor E5-2658 v2 10C 2.4GHz 25MB 1866MHz 95W ^b	G4x

Part number	Feature code ^a	Description	Supported model
94Y5266	A4SZ / A4V5	Intel Xeon Processor E5-2660 v2 10C 2.2GHz 25MB 1866MHz 95W	CAx
94Y5267	A4T0 / A4V6	Intel Xeon Processor E5-2670 v2 10C 2.5GHz 25MB 1866MHz 115W	CBx
94Y5268	A4T1 / A4V7	Intel Xeon Processor E5-2680 v2 10C 2.8GHz 25MB 1866MHz 115W	CCx
94Y5269	A4T2 / A4V8	Intel Xeon Processor E5-2690 v2 10C 3.0GHz 25MB 1866MHz 130W	CTO ^c
94Y5276	A4SN / A4V9	Intel Xeon Processor E5-2695 v2 12C 2.4GHz 30MB 1866MHz 115W	CTO ^c
94Y5277	A4SP / A4VA	Intel Xeon Processor E5-2697 v2 12C 2.7GHz 30MB 1866MHz 130W	F2x

a. The first feature code is for the first processor. The second feature code is for the second processor.

b. Intel Xeon robust thermal profile processors.

c. The processor is not available in any of the standard models listed in Table 3-69 on page 306. The processor is, however, available through the configure-to-order (CTO) ordering process.

3.5.5 Memory

The HS23 has eight DIMM sockets per processor (16 DIMMs in total). The HS23 with Intel Xeon processor E5-2600 v2 product family can support up to 512 GB of memory in total when using 32 GB DIMMs and with both processors installed. The HS23 (E5-2600 v2) uses Double Data Rate-3 (DDR-3) very-low-profile (VLP) registered DIMMs.

Figure 3-49 shows the location of the 16 memory DIMM sockets




Figure 3-49 DIMM layout of the HS23 system board

The Intel Xeon processor E5-2600 v2 product family has an integrated memory controller. The memory DIMMs connect directly to one of the processors. The processor has four memory channels. With one processor installed, only eight DIMMs can be accessed (two DIMMs per channel). To access memory in the other eight DIMM connectors, the second processor must be installed.

Table 3-72 shows which memory DIMM connectors are connected to each processor.

Table 3-72 HS23 (E5-2600 v2) memory channel configuration

Processor	Memory channel	DIMM connector
1	Channel 0	DIMM connector 7 and 8
	Channel 1	DIMM connector 1 and 2
	Channel 2	DIMM connector 3 and 4
	Channel 3	DIMM connector 5 and 6
2	Channel 0	DIMM connector 15 and 16
	Channel 1	DIMM connector 13 and 14
	Channel 2	DIMM connector 9 and 10
	Channel 3	DIMM connector 11 and 12

Table 3-73 lists the memory options for the HS23 (7875, E5-2600 v2).

Table 3-73 List of memory modules for the HS23 (7875, E5-2600 v2)

Part number	Feature code	Description	Models where used
1866 MHz DIMMs			
46W0688	A4G0	4GB (1x4GB, 2Rx8, 1.5V) PC3-14900 CL13 ECC DDR3 1866MHz VLP RDIMM	-
46W0696	A4G1	8GB (1x8GB, 1Rx4, 1.5V) PC3-14900 CL13 ECC DDR3 1866MHz VLP RDIMM	-
46W0704	A4G2	8GB (1x8GB, 2Rx8, 1.5V) PC3-14900 CL13 ECC DDR3 1866MHz VLP RDIMM	C9x, CAx, CBx, CCx, F2x, G3x, G4x
46W0712	A4G3	16GB (1x16GB, 2Rx4, 1.5V) PC3-14900 CL13 ECC DDR3 1866MHz VLP RDIMM	-
1600 MHz DIMMs			
46W0684	A4FV	4GB (1x4GB, 1Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600MHz VLP RDIMM	-
46W0692	A4FW	4GB (1x4GB, 2Rx8, 1.35V) PC3L-12800 CL11 ECC DDR3 1600MHz VLP RDIMM	-
46W0700	A4FX	8GB (1x8GB, 1Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600MHz VLP RDIMM	-
46W0708	A4FY	8GB (1x8GB, 2Rx8, 1.35V) PC3L-12800 CL11 ECC DDR3 1600MHz VLP RDIMM	A3x, A4x, B4x, B5x, B6x, D2x
46W0716	A4G9	16GB (1x16GB, 2Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600MHz VLP RDIMM	-
1333 MHz DIMMs			
00D5008	A3KN	32GB (1x32GB, 4Rx4, 1.35V) PC3L-10600 CL9 ECC DDR3 1333MHz VLP RDIMM	-

Memory DIMM installation suggestions

The HS23 server supports up to eight DIMMs when one processor is installed and up to 16 DIMMs when two processors are installed. Each processor has four memory channels, and there are two DIMMs per channel.

The following rules apply when selecting the memory configuration:

- ▶ Mixing 1.5 V and 1.35 V DIMMs in the same server is supported. In such a case, all DIMMs operate at 1.5 V.
- ▶ The maximum number of ranks supported per channel is eight.
- ▶ The maximum quantity of DIMMs that can be installed in the server depends on the number of processors, DIMM rank, and operating voltage, as shown in the “Maximum quantity supported” row in Table 3-74.
- ▶ All DIMMs in all processor memory channels operate at the same speed, which is determined as the lowest value of:
 - Memory speed supported by specific processor
 - Lowest maximum operating speed for the selected memory configuration that depends on rated speed, as shown under the “Maximum operating speed” section in Table 3-56.

Table 3-74 Maximum memory speeds for HS23 (7875, E5-2600 v2)

DIMM type	RDIMM							
DIMM specification								
Rank	Single rank			Dual rank			Quad rank	
Rated speed	1600 MHz		1866 MHz	1600 MHz		1866 MHz	1333 MHz	
Rated voltage	1.35 V		1.5 V	1.35 V		1.5 V	1.35 V	
Operating voltage	1.35 V	1.5 V	1.5 V	1.35 V	1.5 V	1.5 V	1.35 V	1.5 V
Maximum quantity supported	16	16	16	16	16	16	16	16
Maximum DIMM capacity	8 GB	8 GB	8 GB	16 GB	16 GB	16 GB	32 GB	32 GB
Maximum memory capacity	128 GB	128 GB	128 GB	256 GB	256 GB	256 GB	512 GB	512 GB
Max. memory at rated speed	N/A	128 GB	64 GB	N/A	256 GB	128 GB	N/A	N/A
Maximum operating speed								
1 DIMM per channel	1333 MHz	1600 MHz	1866 MHz	1333 MHz	1600 MHz	1866 MHz	800 MHz	1066 MHz
2 DIMMs per channel	1333 MHz	1600 MHz	1600 MHz	1333 MHz	1600 MHz	1600 MHz	800 MHz	800 MHz

Terminology: The term *DIMMs per channel* is commonly referred to as the abbreviation DPC. One DIMM per channel can be referred to as 1DPC and, likewise, two DIMMs per channel can be referred to as 2DPC. The ability to run two DIMMs per channel at 1600 MHz is often referred to as the abbreviation 2DPC@1600MHz.

Memory modes

The HS23 type 7875 supports three different memory modes.

These modes can be selected in the UEFI BIOS setup:

► Independent channel (normal) mode

Independent channel mode offers a maximum of 256 GB of usable memory with one microprocessor installed, and 512 GB of usable memory with two microprocessors installed (using 32 GB DIMMs). The DIMMs can be installed without matching sizes.

► Memory sparing (rank sparing) mode

If memory sparing (rank sparing) is used, the DIMMs' population order is the same as that for independent channel mode (the DIMMs can be installed without matching sizes). In rank-sparing mode, one rank of a DIMM in each populated channel is reserved as spare memory. The spare rank is not available as active memory. When the active rank memory fails, its content is copied to the spare rank memory, which becomes active. The memory set aside for spare is one rank per channel. The size of a rank varies depending on the DIMMs installed.

Table 3-75 DIMM population sequence for independent channel mode and rank-sparing mode

Processors installed	DIMM connectors
One	1, 7, 3, 5, 2, 8, 4, and 6
Two	1, 16, 7, 14, 3, 9, 5, 11, 2, 15, 8, 13, 4, 10, 6, and 12

► Memory mirrored channel mode

If memory mirroring is used, DIMMs must be installed in pairs (minimum of one pair per processor), and both DIMMs in a pair must be identical in type and size. In mirrored channel mode, the memory contents on channel 0 are duplicated in channel 2, and the memory contents of channel 1 are duplicated in channel 3. The effective memory that is available to the system is only half of the memory installed.

Table 3-76 DIMM population sequence for mirrored channel mode

	DIMM connectors	
DIMM pair	One processor installed	Two processors installed
First	1 and 7	1 and 7
Second	3 and 5	14 and 16
Third	2 and 8	3 and 5
Fourth	4 and 6	9 and 11
Fifth	None	2 and 8
Sixth	None	13 and 15
Seventh	None	4 and 6
Eighth	None	10 and 12
Note: The DIMM's pairs must be identical in size, type, and rank count.		

Important: DIMMs or DIMM fillers must occupy DIMM connectors 1, 2, 13, 14, 15, and 16 for the correct cooling.

3.5.6 Standard onboard features

The HS23 lvB blade server has the standard onboard features shown in the following sections.

UEFI

The HS23 lvB uses an integrated Unified Extensible Firmware Interface (UEFI) next-generation BIOS.

New capabilities include:

- ▶ Human readable event logs. No more beep codes!
- ▶ A complete setup solution by allowing the adapter configuration function to be moved into the UEFI.
- ▶ Complete out-of-band coverage by the Advance Settings Utility to simplify remote setup.

Using all of the features that UEFI offers requires a UEFI-aware operating system and adapters. UEFI is fully compatible with earlier versions of the BIOS.

For more information about UEFI, see *Introducing uEFI-Compliant Firmware on IBM System x and BladeCenter Servers*, found at:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5083207>

Onboard network adapter

The HS23 offers two integrated Gigabit Ethernet ports and two integrated 10 Gb Ethernet ports with the integrated Emulex BladeEngine 3 (BE3) controller. Two Gigabit Ethernet ports are routed to the chassis I/O bays 1 and 2, and two 10 Gb Ethernet ports are routed to the chassis I/O bays 7 and 9 using either a 10 Gb LOM Interposer Card or Emulex Virtual Fabric Adapter II for HS23.

The integrated BE3 4-port NIC has the following features:

- ▶ Two Gigabit Ethernet ports and two 10 Gb Ethernet ports (1 Gb and 10 Gb auto-negotiation)
- ▶ Full-duplex (FDX) capability
- ▶ 10 Gb ports operate in either a virtual NIC (vNIC) or physical NIC (pNIC) mode:
 - vNIC mode: Up to six vNICs (up to three vNICs per one 10 Gb port):
 - Virtual Fabric mode or Switch Independent operational mode.
 - Virtual port bandwidth allocation in 100 Mbps increments.
 - Up to two vNICs can be configured as an iSCSI or FCoE vNICs (one per port) with optional Advanced Upgrade (90Y9310).
 - pNIC mode: Dual-port 1/10 Gb Ethernet adapter
- ▶ IPv4/IPv6 offload:
 - TCP and UDP checksum offload
 - Large send offload (LSO)
 - Large receive offload (LRO)
 - Receive side scaling (RSS)

- ▶ IPv4 TCP Chimney Offload
- ▶ IEEE 802.1Q VLAN tagging
- ▶ VLAN insertion and extraction
- ▶ Jumbo frames up to 9000 bytes
- ▶ Load balancing and failover teaming support, including adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), and IEEE 802.3ad
- ▶ Enhanced Ethernet (draft) support:
 - Enhanced Transmission Selection (ETS) (P802.1Qaz)
 - Priority-based Flow Control (PFC) (P802.1Qbb)
 - Data Center Bridging eXchange Protocol (DCBX) (P802.1Qaz)
- ▶ Supports Serial over LAN (SoL) and concurrent KVM (cKVM)
- ▶ Preboot Execution Environment (PXE) support
- ▶ Wake On LAN support
- ▶ PCIe x8 Gen 2 host interface
- ▶ Message Signal Interrupt (MSI-X) support

Table 3-59 shows the ordering part number and feature code for Advanced Upgrade for integrated NIC to enable FCoE or iSCSI protocols in it.

Table 3-77 Integrated NIC upgrade ordering part number and feature code

Part number	Feature code	Description
90Y9310	A2TD	IBM Virtual Fabric Advanced Software Upgrade (LOM)

Important: A 10 Gb Interposer card for IBM BladeCenter HS23 (or 10 Gb LOM Interposer Card) part number 94Y8550 must be installed if you plan to use the 10 Gb ports of the onboard network adapter. For details, see “LOM Interposer Card” on page 300.

Integrated systems management processor

The HS23 type 7875 blade server includes an Integrated Management Module Version 2 (IMM V2) to monitor server availability, perform Predictive Failure Analysis, and trigger IBM Systems Director alerts. IMM V2 provides system management capabilities with IPMI 2.0 compliance and Serial over LAN.

For more information about the IMM, see *Transitioning to UEFI and IMM*, found at:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5079769>

Video controller

The video controller is a Matrox G200 with 16 MB integrated into IMM v2. The maximum resolution is 1600x1200 at 75 Hz with 16 M colors (32 bits per pixel).

Trusted Platform Module

Trusted computing is an industry initiative that provides a combination of secure software and secure hardware to create a trusted platform. It is a specification that increases network security by building unique hardware IDs into computing devices.

The HS23 implements Trusted Platform Module (TPM) Version 1.2 support.

The trusted platform module (TPM) in the HS23 is one of the three layers of the trusted computing initiative (Table 3-78).

Table 3-78 Trusted computing layers

Layer	Implementation
Level 1: Tamper-proof hardware, used to generate trustable keys	Trusted Platform Module
Level 2: Trustable platform	<ul style="list-style-type: none"> ▶ UEFI or BIOS ▶ Intel processor
Level 3: Trustable execution	<ul style="list-style-type: none"> ▶ Operating system ▶ Drivers

3.5.7 Local storage

The HS23 blade server features an onboard LSI Logic SAS2004 SAS controller with two hot-swap drives bays that are accessible from the front of the blade server (Figure 3-42). The LSI Logic SAS2004 SAS controller provides RAID 0 or RAID 1 capability and supports up to two internal hot-swap SAS or SATA HDDs or two internal hot-swap solid-state drives. The LSI SAS2004 SAS controller ports operate at a burst rate of 6.0 Gbps.




Figure 3-50 The HS23 showing the front hot-swap disk drive bays

Local SAS or SATA disk drives

The HS23 (E5-2600 v2) has support for up to two hot-swap SFF SAS or SATA hard disk drives. These two hot-swap hard disk drives are accessible from the front of the blade server without removing the blade server from the BladeCenter chassis. See Table 3-79 for a list of supported SAS and SATA HDDs.

Table 3-79 Supported SAS and SATA HDDs

Part number	Feature code	Description
10K SAS hard disk drives		
90Y8877	A2XC	IBM 300GB 10K 6Gbps SAS 2.5" SFF G2HS HDD
90Y8872	A2XD	IBM 600GB 10K 6Gbps SAS 2.5" SFF G2HS HDD
00AJ300	A4VB	IBM 600GB 15K 6Gbps SAS 2.5" G2HS HDD

Part number	Feature code	Description
81Y9650	A282	IBM 900GB 10K 6Gbps SAS 2.5" SFF HS HDD
00AD075	A48S	IBM 1.2TB 10K 6Gbps SAS 2.5" G2HS HDD
15K SAS hard disk drives		
90Y8926	A2XB	IBM 146GB 15K 6Gbps SAS 2.5" SFF G2HS HDD
81Y9670	A283	IBM 300GB 15K 6Gbps SAS 2.5" SFF HS HDD
SAS hybrid hard disk drives		
00AJ236	A4VD	IBM 300GB 15K 6Gbps SAS 2.5" G2HS Hybrid
00AD102	A4G7	IBM 600GB 10K 6Gbps SAS 2.5" G2HS Hybrid
00AJ246	A4VF	IBM 600GB 15K 6Gbps SAS 2.5" G2HS Hybrid
NL SAS hard disk drives		
90Y8953	A2XE	IBM 500GB 7.2K 6Gbps NL SAS 2.5" SFF G2HS HDD
81Y9690	A1P3	IBM 1TB 7.2K 6Gbps NL SAS 2.5" SFF HS HDD
NL SATA hard disk drives		
81Y9722	A1NX	IBM 250GB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD
81Y9726	A1NZ	IBM 500GB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD
81Y9730	A1AV	IBM 1TB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD

Solid-state drives

SSDs are more reliable than spinning hard disk drives and they consume much less power than a standard SAS drive. Target Applications for SSDs include video surveillance, transaction-based databases, and other applications that have high performance but moderate space requirements.

Table 3-80 lists the supported SSDs.

Table 3-80 Supported solid-state drives (SSDs)

Part number	Feature code	Description
Solid-state drives - Enterprise		
00W1125	A3HR	IBM 100GB SATA 2.5" MLC HS Enterprise SSD
49Y6129	A3EW	IBM 200GB SAS 2.5" MLC HS Enterprise SSD
49Y6134	A3EY	IBM 400GB SAS 2.5" MLC HS Enterprise SSD
49Y6139	A3F0	IBM 800GB SAS 2.5" MLC HS Enterprise SSD
49Y6195	A4GH	IBM 1.6TB SAS 2.5" MLC HS Enterprise SSD
41Y8331	A4FL	S3700 200GB SATA 2.5" MLC HS Enterprise SSD
41Y8336	A4FN	S3700 400GB SATA 2.5" MLC HS Enterprise SSD
41Y8341	A4FQ	S3700 800GB SATA 2.5" MLC HS Enterprise SSD

Part number	Feature code	Description
Solid-state drives - Enterprise value		
49Y5839	A3AS	IBM 64GB SATA 2.5" MLC HS Enterprise Value SSD
90Y8648	A2U4	IBM 128GB SATA 2.5" MLC HS Enterprise Value SSD
90Y8643	A2U3	IBM 256GB SATA 2.5" MLC HS Enterprise Value SSD
49Y5844	A3AU	IBM 512GB SATA 2.5" MLC HS Enterprise Value SSD
00AJ000	A4KM	S3500 120GB SATA 2.5" MLC HS Enterprise Value SSD
00AJ005	A4KN	S3500 240GB SATA 2.5" MLC HS Enterprise Value SSD
00AJ010	A4KP	S3500 480GB SATA 2.5" MLC HS Enterprise Value SSD
00AJ015	A4KQ	S3500 800GB SATA 2.5" MLC HS Enterprise Value SSD

PCIe SSD adapters

The HS23 server supports the High IOPS SSD adapters listed in Table 3-81. The adapters must be installed in an IBM BladeCenter PCI Express Gen 2 Expansion Blade II. Up to eight High IOPS adapters supported per one HS23 (two per PCI Express Gen 2 Expansion Blade II and up to four PCI Expansion Blades per HS23)

Table 3-81 Supported PCIe SSD adapters

Part number	Feature code	Description	Slots supported
46C9078	A3J3	IBM 365GB High IOPS MLC Mono Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
46C9081	A3J4	IBM 785GB High IOPS MLC Mono Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4377	A3DY	IBM 1.2TB High IOPS MLC Mono Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4397	A3DZ	IBM 2.4TB High IOPS MLC Duo Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4361	A3MZ	IBM 300GB High IOPS MLC Modular Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4365	A3N0	IBM 600GB High IOPS MLC Modular Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4369	A3N1	IBM 800GB High IOPS MLC Modular Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4373	A3N2	IBM 300GB High IOPS SLC Modular Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)

3.5.8 Integrated virtualization

The HS23 offers an IBM standard USB flash drive option preinstalled with VMware ESXi. This is an embedded version of VMware ESXi. It is fully contained on the flash drive, without requiring any disk space. The IBM USB Memory Key for VMware Hypervisor plugs into the USB port at the HS23 system board.

Table 3-82 lists the part number, feature code, and description of the IBM USB Memory Key for VMware Hypervisor that is supported in the HS23.

Table 3-82 IBM USB Memory Key for VMware Hypervisor

Part Number	Feature Code	Description
41Y8300	A2VC	IBM USB Memory Key for VMWare ESXi 5.0
41Y8307	A383	IBM USB Memory Key for VMWare ESXi 5.0 Update 1
41Y8311	A2R3	IBM USB Memory Key for VMWare ESXi 5.1
41Y8382	A4WZ	IBM USB Memory Key for VMWare ESXi 5.1 Update 1
41Y8298	A2G0	IBM Blank USB Memory Key for VMWare ESXi Downloads

For a complete description of the features and capabilities of VMware ESX Server, go to the following address:

<http://www.vmware.com/products/vi/esx/>

3.5.9 I/O expansion

The HS23 (7875, E5-2600 v2) is able to connect to a wide variety of networks and fabrics by installing the appropriate I/O expansion card. Supported networks and fabrics include 1 Gb and 10 Gb Ethernet, 4 Gb and 8 Gb Fibre Channel, Serial Attached SCSI (SAS), and InfiniBand.

An HS23 blade server with an installed I/O expansion card must be installed in a BladeCenter chassis with switch modules (or pass-through) that are compatible with the I/O expansion card.

The HS23 blade server supports the following formats of I/O expansion cards: the CIOv format and the CFFh format, as described in the following sections.

LOM Interposer Card

The 10 Gb Interposer Card for IBM BladeCenter HS23 provides I/O connections between the 10 Gb ports of an integrated network adapter (see “Onboard network adapter” on page 295) and high speed switch modules in bays 7 and 9 of a BladeCenter H or BladeCenter HT chassis.

Figure 3-51 shows a 10 Gb LOM Interposer Card installed in the slot of the HS23.




Figure 3-51 The HS23 showing the 10 Gb LOM Interposer Card position

Table 3-83 shows which HS23 (E5-2600 v2) models ship with 10 Gb LOM Interposer Cards.

Table 3-83 HS23 with LOM Interposer Cars included as standard

Blade	7875-A3x	7875-A4x	7875-B4x	7875-B5x	7875-B6x	7875-C9x	7875-CAx	7875-CBx	7875-CCx	7875-C5x	7875-D2x	7875-F2x	7875-G3x	7875-G4x
10 Gb LOM Interposer Card	N	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N

Table 3-84 lists the 10 Gb Interposer Card for IBM BladeCenter HS23 ordering part number and feature code.

Table 3-84 Ordering part number and feature code

Part number	Feature code	Description
94Y8550	A244	10 Gb Interposer Card for IBM BladeCenter HS23

CIOv

The CIOv I/O expansion connector provides I/O connections to modules in bays 3 and 4 of a supported BladeCenter chassis. The CIOv slot is a third-generation PCI Express x8 (PCIe 3.0 x8) slot. A maximum of one CIOv I/O expansion card is supported per single-wide blade server. A CIOv I/O expansion card can be installed on a blade server at the same time as a CFFh I/O expansion card is installed.

Table 3-85 on page 321 lists the CIOv expansion cards that are supported in the HS23.

Table 3-85 Supported CIOv expansion cards for the HS23 server

Part number	Feature code	Description
44W4475	1039	Ethernet Expansion Card (CIOv)
46M6065	3594	QLogic 4 Gb Fibre Channel Expansion Card (CIOv)
43W4068	1041	SAS Connectivity Card (CIOv)
46M6140	3598	Emulex 8 Gb Fibre Channel Expansion Card (CIOv)
44X1945	1462	QLogic 8 Gb Fibre Channel Expansion Card (CIOv)

For the latest information about the expansion cards supported by the HS23, see the IBM ServerProven website at:

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

CIOv expansion cards are installed in the CIOv slot in the HS23 (Figure 3-52).



Figure 3-52 The HS23 tlvB type 7875 showing the CIOv I/O expansion card position

The CIOv expansion cards require a supported I/O module installed in bay 3 and bay 4 of the chassis where the cards and servers are installed.

CFFh

The CFFh I/O expansion connector provides I/O connections to high speed switch modules in bays 7, 8, 9, and 10 of a BladeCenter H or BladeCenter HT chassis or to switch bay 2 in a BladeCenter S chassis. The CFFh slot is a third-generation PCI Express x16 (PCIe 3.0 x16) slot. A maximum of one CFFh I/O expansion card is supported per blade server. A CFFh I/O expansion card can be installed on a blade server at the same time as a CIOv I/O expansion card is installed.

Table 3-86 lists the CFFh expansion cards that are supported in the HS23 (E5-2600 v2).

Table 3-86 Supported CFFh expansion cards for the HS23 (E5-2600 v2)

Part number	Feature code	Description
46M6001	0056	2-port 40 Gb InfiniBand Expansion Card (CFFh)
44W4479	5476	2/4 Port Ethernet Expansion Card (CFFh)
46M6164	0098	Broadcom 10 Gb Gen2 4-port Ethernet Exp Card (CFFh)
81Y3133	A1QR	Broadcom 2-port 10 Gb Virtual Fabric Adapter
81Y1650	5437	Brocade 2 port 10 GbE Converged Network Adapter
81Y3120	A287	Emulex 10 GbE VFA II for HS23
90Y9332	A2ZN	Emulex 10 GbE VFA Advanced II for HS23
2C1810	3593	Intel 10 Gb 2-port Ethernet Expansion Card (CFFh)
90Y3570	A1NW	Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh)
00Y3280 ^a	A3JB	QLogic 2-port 10 Gb Converged Network Adapter (CFFh)
00Y3332	A4AC	QLogic 10Gb Virtual Fabric Adapter (CFFh)
00Y5618	A4AD	QLogic 10 Gb Virtual Fabric Converged Network Adapter (CFFh)
00Y3270	A3JC	QLogic Eth and 8 Gb Fibre Channel Exp Card (CFFh)

a. Replaces 42C1830.

For the latest information about the expansion cards supported by the HS23, see the IBM ServerProven website at:

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

CFFh expansion cards are installed in the CFFh slot in supported servers, such as the HS23 (Figure 3-53).




Figure 3-53 The HS23 showing the CFFh I/O expansion card position

The CFFh I/O expansion card requires that a supported high speed I/O module is installed in bay 7, bay 8, bay 9, or bay 10 of the BladeCenter H or BladeCenter HT chassis. Examples of the types of modules supported in the high speed I/O module bays are the InfiniBand Pass-thru Module and the IBM Virtual Fabric 10Gb Switch Module.

A Multi-switch Interconnect Module can be used in the high speed switch module bays of the BladeCenter H (MSIM) or BladeCenter HT (MSIM-HT) to allow connectivity to standard switch modules.

In a BladeCenter S chassis, the CFFh I/O Expansion Card requires a supported switch module in bay 2. When used in a BladeCenter S chassis, a maximum of two ports are routed from the CFFh I/O Expansion Card to the switch module in bay 2.

For the latest information about the switch modules that are supported with each CFFh I/O expansion card, see *IBM BladeCenter Interoperability Guide*, found at:

<http://www.redbooks.ibm.com/big>

Emulex 10 GbE VFA II and VFA II Advanced (CFFh) for HS23

The Emulex 10 GbE VFA II CFFh adapters for HS23 use a CFFh I/O expansion connector to provide I/O connections to high speed switch modules in bays 8 and 10 of a BladeCenter H or BladeCenter HT chassis. These adapters provide I/O connections between 10 Gb ports of an integrated network adapter (see “Onboard network adapter” on page 295) and high speed switch modules in bays 7 and 9 of a BladeCenter H or BladeCenter HT chassis.

An Emulex 10 GbE VFA II for HS23 CFFh I/O expansion card replaces a 10 Gb LOM Interposer Card if you need 2+2 10 Gb ports (total four ports). A maximum of one VFA II for HS23 is supported per blade server.

The Emulex 10GbE VFA II cards require that a supported high speed I/O module is installed in bay 7, bay 8, bay 9, or bay 10 of the BladeCenter H or BladeCenter HT chassis. Examples of the types of modules supported in the high speed I/O module bays are the IBM Virtual Fabric 10 Gb Switch Module and the Cisco Nexus 4001I Switch Module.

For more information about this card, see 3.1.9, “Emulex 10 GbE Virtual Fabric Adapter II/Advanced II for HS23” on page 231.

PCI Express Gen 2 Expansion Blade II

The HS23 supports the attachment of PCI Express Gen 2 Expansion Blades II. This unit allows certain standard PCI Express adapters to be attached to the HS23. If there is a CFFh I/O expansion card already installed on the HS23, it must be removed to allow for the attachment of the PCI Express Gen 2 Expansion Blade II. The CFFh I/O expansion card can then be installed in the topmost PCI Express Gen 2 Expansion Blade. No other type of expansion blade can be mixed with the PCI Express Gen 2 Expansion Blade.

For more information about the BPE4, see 3.10.2, “PCI Express Gen 2 Expansion Blade/Expansion Blade II” on page 403.

IBM BladeCenter GPU Expansion Blade II

The HS23 supports the attachment of up to four IBM BladeCenter GPU Expansion Blades II. This unit ships integrated with one NVIDIA Tesla M2075, NVIDIA Tesla M2070Q, or NVIDIA Tesla M2090 GPU adapter that can be attached to the HS23. If there is a CFFh I/O expansion card already installed on the HS23, it must be removed to allow for the attachment of the IBM BladeCenter GPU Expansion Blade. The CFFh I/O expansion card can then be installed in the topmost IBM BladeCenter GPU Expansion Blade. No other type of expansion blade can be mixed with the IBM BladeCenter GPU Expansion Blade.

For more information about the BGE, see 3.10.3, “IBM BladeCenter GPU Expansion Blade” on page 408.

3.6 BladeCenter HS23E

The IBM BladeCenter HS23E is a versatile, dual-socket blade server running the Intel Xeon processor E5-2400 product family. The server offers performance for value with new levels of memory capacity, processor performance, and flexible configuration options. A standard 30 mm single-wide form factor protects your investments by providing compatibility with the IBM BladeCenter H, E, S, and HT chassis.

The following topics are covered in this section:

- ▶ 3.6.1, “Overview” on page 324
- ▶ 3.6.2, “Model information” on page 326
- ▶ 3.6.3, “Chassis support” on page 327
- ▶ 3.6.4, “Processor” on page 328
- ▶ 3.6.5, “Memory” on page 330
- ▶ 3.6.6, “Standard onboard features” on page 334
- ▶ 3.6.7, “Local storage” on page 336
- ▶ 3.6.8, “Integrated virtualization” on page 338
- ▶ 3.6.9, “I/O expansion” on page 339

3.6.1 Overview

The HS23E has the following features:

- ▶ Supports up to two Intel Xeon E5-2400 family processors, or one Intel Xeon E5-1410, or one Intel Pentium 1400 family processor (Intel Pentium 1403 or Intel Pentium 1407)
- ▶ Twelve DIMM slots for up to 192 GB of memory (using 16 GB DIMMs)
- ▶ Supports up to two hot-swap HDDs
- ▶ Internal USB port for embedded hypervisor support
- ▶ Integrated dual-port Gigabit Ethernet
- ▶ Two PCI slots: CIOv and CFFh

Figure 3-54 shows the HS23E.




Figure 3-54 The HS23E Type 8038

Table 3-87 lists the features of the HS23E.

Table 3-87 Features of the HS23E type 8038

Feature	Description
Processor options and number of processors (std/max)	Intel Xeon E5-2400 family processors: - 1/2 Intel Xeon E5-1410 - 1/1 Intel Pentium 1400 family processors - 1/1
Cache	10 MB, 15 MB, or 20 MB 10 MB 5 MB
Memory speed	Up to 1600 MHz (dependent on processor model) 1333 MHz 1066 MHz
Memory slots	12 DIMM slots
Memory maximum	192 GB using 16 GB DIMMs
Memory types	DDR 3 ECC Very Low Profile (VLP) Registered DIMM
Hot swap disk drives (std/max)	0/2 (SAS, SATA, or SSD)
Maximum internal storage	Optional: Up to 1800 GB with two 2.5-inch SAS hot swap drives Optional: Up to 2 TB with two 2.5-inch NL SAS or NL SATA hot swap drives Optional: Up to 400 GB with two 2.5-inch solid-state drives (SSDs) SATA hot swap drives
Network	Two Gigabit Ethernet ports with an integrated Broadcom BCM5718 controller
I/O upgrade	Two PCI Express connectors (CFFh and CIOv form factors)

Figure 3-55 shows the components on the system board of the HS23E.




Figure 3-55 Layout of the HS23E system board

3.6.2 Model information

Table 3-88 lists available models of the HS23E.

Table 3-88 Models of the HS23E type 8038

Server models ^a	Processor (std/max)	Memory (std/max ^b)	Onboard disk controller	Drive bays (used/free)	Network	I/O slots (used/max)
8038-B1x	1x E5-2403 4C 1.8 GHz 10 MB 1066 MHz 80 W / 2	1 x 2 GB / 192 GB	C105 ^c	0/2	2x 1 GbE	0/2
8038-B3x	1x E5-2407 4C 2.2 GHz 10 MB 1066 MHz 80 W / 2	3 x 4 GB / 192 GB	C105 ^c	0/2	2x 1 GbE	0/2
8038-C2x	1x E5-2430 6C 2.2 GHz 15 MB 1333 MHz 95 W / 2	3 x 4 GB / 192 GB	C105 ^c	0/2	2x 1 GbE	0/2
8038-C4x	1x E5-2430 6C 2.2 GHz 15 MB 1333 MHz 95 W / 2	3 x 8 GB / 192 GB	C105 ^c	0/2	2x 1 GbE	0/2
8038-C3x	1x E5-2440 6C 2.4 GHz 15 MB 1333 MHz 95 W / 2	3 x 4 GB / 192 GB	C105 ^c	0/2	2x 1 GbE	0/2
8038-D1x	1x E5-2450 8C 2.1 GHz 20 MB 1600 MHz 95 W / 2	3 x 8 GB / 192 GB	C105 ^c	0/2	2x 1 GbE	0/2
8038-D2x	1x E5-2450 8C 2.1 GHz 20 MB 1600 MHz 95 W / 2	3 x 8 GB / 192 GB	C105 ^c	0/2	2x 1 GbE	0/2

Server models ^a	Processor (std/max)	Memory (std/max ^b)	Onboard disk controller	Drive bays (used/free)	Network	I/O slots (used/max)
8038-D3x	1x E5-2470 8C 2.3 GHz 20 MB 1600 MHz 95 W / 2	3 x 8GB / 192 GB	C105 ^c	0/2	2x 1 GbE	0/2
8038-D4x	1x E5-2470 8C 2.3 GHz 20 MB 1600 MHz 95 W / 2	3 x 8 GB / 192 GB	C105 ^c	0/2	2x 1 GbE	0/2
8038-F1x	1x Pentium 1403 2C 2.6 GHz 5 MB 1066 MHz 80 W / 1	1 x 2 GB / 96 GB ^d	C105 ^c	0/2	2x 1 GbE	0/2
8038-F2x	1x Pentium 1407 2C 2.8 GHz 5 MB 1066 MHz 80 W / 1	3 x 4 GB / 96 GB ^d	C105 ^c	0/2	2x 1 GbE	0/2

a. The model numbers provided are worldwide generally available variant (GAV) model numbers that are not orderable as listed and must be modified by country. The US GAV model numbers use the following nomenclature: xxU. For example, the US orderable part number for 8038-B1x is 8038-B1U. See the product-specific official IBM Announcement Letter for other country-specific GAV model numbers.

b. Maximum system memory capacity of 192 GB is based on using 16 GB memory DIMMs.

c. Support for RAID 0 and 1 for SATA HDDs only.

d. Maximum system memory capacity of 96 GB is based on using 6 x 16 GB memory DIMMs.

3.6.3 Chassis support

The HS23E type 8038 is supported in the BladeCenter chassis, as listed in Table 3-89 on page 327. See the following subsections for details about this support.

Table 3-89 HS23E chassis support

Blade	TDP ^a	BC E 8677	BC T 8720/8730	BC S 8886	BC H 8852	BC HT DC 8740	BC HT AC 8750
HS23E	80 W or 95 W	Some limits ^b	No	Full	Full	No	No
HS23E	Up to 95 W	Full	No	Full	Full	Some limits ^b	Some limits ^b

a. Thermal design power (TDP) of the processor in the server; see Table 3-88 to determine to which HS23 server models TDP applies.

b. See Table 3-90 for details.

The specifics of the supported chassis depend on the models of the chassis and the power supplies and cooling modules (blowers) in the chassis, as shown in Table 3-90.

- ▶ A green cell (■) means that the chassis can be filled with HS23E blade servers up to the maximum number of blade bays in the chassis (14 blades in the BladeCenter H).
- ▶ A yellow cell (■) means that the maximum number of HS23E blades that the chassis can hold is fewer than the total available blade bays. *All other bays must remain empty.* Empty bays must be distributed evenly between the two power domains of the chassis (bays 1 - 6 and bays 7 - 14).

Table 3-90 HS23E chassis compatibility

CPU TDP ^a	Maximum number of HS23E supported in each chassis									
	BC-E with AMM (8677) (14 bays)		BC-S (8886) (6 bays)	BC-H (models other than 4Tx)				BC-H (-4Tx)	BC-HT DC (8740) (12 bays)	BC-HT AC (8750) (12 bays)
	2000 W power supplies	2320 W power supplies		2900 W supplies		2980 W supplies ^b		2980 W ^c		
				Std. blower	Enh. blower ^c	Std. blower	Enh. blower ^d	Enh. blower ^d		
Intel Xeon processors										
95 W	5+6	14	6	14	14	14	14	14	None ^d	None ^d
80 W	5+7	14	6	14	14	14	14	14	None ^d	None ^d
70 W	14	14	6	14	14	14	14	14	None ^d	None ^d
60 W	14	14	6	14	14	14	14	14	None ^d	None ^d
Intel Xeon robust thermal profile processors ^e										
70 W	14	14	6	14	14	14	14	14	12	12
60 W	14	14	6	14	14	14	14	14	12	12
50 W	14	14	6	14	14	14	14	14	12	12

a. Thermal design power (TDP)

b. IBM BladeCenter H 2980 W AC Power Modules, 68Y6601 (standard in 4Tx, optional with all other BC-H chassis models).

c. IBM BladeCenter H Enhanced Cooling Modules, 68Y6650 (standard in 4Tx, optional with all other BC-H chassis models).

d. Not supported.

e. Intel Xeon E5-2418L (50 W), E5-2428L (60 W), and E5-2448L (70 W) are robust thermal profile processors that are used in HS23E.

3.6.4 Processor

The server supports up to two Intel Xeon processor E5-2400 product family processors, or one Intel Xeon processor E5-1410, or one Intel Pentium processor 1403, or one Intel Pentium processor 1407.

The Intel Xeon processor E5-2400 product family has models with four, six, or eight cores per processor. The Intel Xeon processor E5-1410 has four cores per processor. The Intel Pentium processor 1403 and Intel Pentium processor 1407 have two cores per processor. The memory controller is built into the processor. The memory is directly connected to the processor through three memory controller buses.

In two-socket configurations, the two Intel Xeon processors (CPU1 and CPU2) in the HS23E are connected through a serial coherency link called QuickPath Interconnect (QPI), as shown in Figure 3-56. The QPI link is capable of up to 8.0 GTps, depending on the processor model.

In the single processor configurations, only one processor (CPU1) is used.

Figure 3-56 shows the major components of the HS23E system.

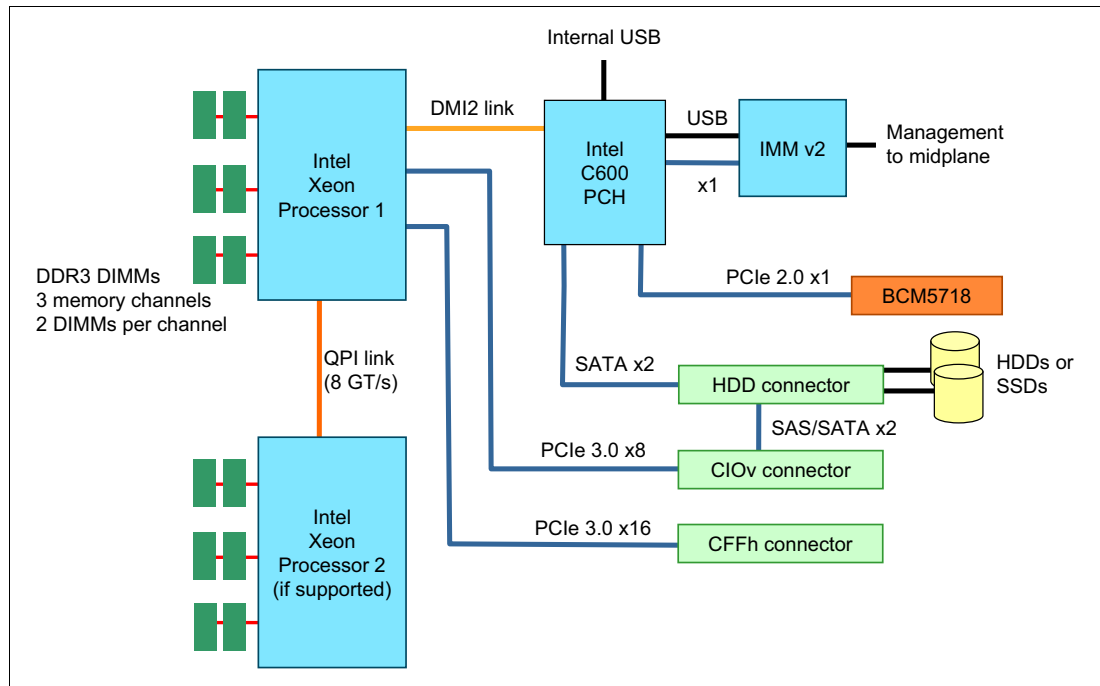


Figure 3-56 HS23E processor block diagram

The processors must be identical in a two-socket configuration. The Intel Xeon processors are available in various clock speeds and have standard and lower power offerings.

Table 3-91 lists the processor options for the HS23E and the models that include them (if one exists).

Table 3-91 Available processor options for HS23E

Part number	Feature code ^a	Description	Supported model
Single or dual processor support: Intel Xeon processor E5-2400 product family			
90Y5292	A2D7 / A2DJ	Intel Xeon Processor E5-2403 4C 1.8 GHz 10 MB Cache 1066 MHz 80 W	B1x
90Y5291	A2D6 / A2DH	Intel Xeon Processor E5-2407 4C 2.2 GHz 10 MB Cache 1066 MHz 80 W	B3x
94Y6288	A34J / A34M	Intel Xeon Processor E5-2418L 4C 2.0 GHz 10 MB Cache 1333 MHz 50 W	CTO ^b
90Y5290	A2D5 / A2DG	Intel Xeon Processor E5-2420 6C 1.9 GHz 15 MB Cache 1333 MHz 95 W	CTO ^b
94Y6289	A34K / A34N	Intel Xeon Processor E5-2428L 6C 1.8 GHz 15 MB Cache 1333 MHz 60 W	CTO ^b
90Y5288	A2D4 / A2DF	Intel Xeon Processor E5-2430 6C 2.2 GHz 15 MB Cache 1333 MHz 95 W	C2x, >>C4x
90Y5294	A2D9 / A2DL	Intel Xeon Processor E5-2430L 6C 2.0 GHz 15 MB Cache 1333 MHz 60 W	CTO ^b
90Y5287	A2D3 / A2DE	Intel Xeon Processor E5-2440 6C 2.4 GHz 15 MB Cache 1333 MHz 95 W	C3x
94Y6290	A34L / A34P	Intel Xeon Processor E5-2448L 8C 1.8 GHz 20 MB Cache 1600 MHz 70 W	CTO ^b

Part number	Feature code ^a	Description	Supported model
90Y5286	A2D2 / A2DD	Intel Xeon Processor E5-2450 8C 2.1 GHz 20 MB Cache 1600 MHz 95 W	D1x, >>D2x
90Y5293	A2D8 / A2DK	Intel Xeon Processor E5-2450L 8C 1.8 GHz 20 MB Cache 1600 MHz 70 W	CTO ^b
90Y5284	A2D1 / A2DC	Intel Xeon Processor E5-2470 8C 2.3 GHz 20 MB Cache 1600 MHz 95 W	D3x, >>D4x
Single processor support only: Intel Xeon processor E5-1410 and Intel Pentium processor 1400 product family			
None ^c	A2DA	Intel Pentium Processor 1403 2C 2.6 GHz 5 MB Cache 1066 MHz 80 W	F1x
None ^c	A2DB	Intel Pentium Processor 1407 2C 2.8 GHz 5 MB Cache 1066 MHz 80 W	F2x
None ^c	A35P	Intel Xeon Processor E5-1410 4C 2.8 GHz 10 MB Cache 1333 MHz 80 W	CTO ^b

- a. The first feature code is for the first processor. The second feature code is for the second processor.
- b. The processor is not available in any of the standard models listed in Table 3-88 on page 326. The processor is, however, available through the configure-to-order (CTO) ordering process.
- c. These processors support only single processor configurations.

3.6.5 Memory

The HS23E supports up to six DIMMs when one processor is installed and up to 12 DIMMs when two processors are installed. The HS23E with Intel Xeon processor E5-2400 product family can support up to 192 GB of memory in total when using 16 GB DIMMs and with both processors installed.

Figure 3-57 shows the location of the 12 memory DIMM sockets.




Figure 3-57 DIMM layout of the HS23E system board

With one processor installed, only six DIMMs can be accessed (two DIMMs per channel). To access memory in the other six DIMM connectors, the second processor must be installed (if supported).

Table 3-92 shows which memory DIMM connectors are connected to each processor.

Table 3-92 HS23E memory channel configuration

Processor	Memory channel	DIMM connector
1	Channel 1	DIMM connector 1 and 2
	Channel 2	DIMM connector 3 and 4
	Channel 3	DIMM connector 5 and 6
2	Channel 1	DIMM connector 7 and 8
	Channel 2	DIMM connector 9 and 10
	Channel 3	DIMM connector 11 and 12

Table 3-93 lists memory options for the HS23E (ordered by voltage and then DIMM capacity).

Table 3-93 List of memory modules for the HS23E type 8038

Part number	Feature code	Description	Models where used
1066 MHz RDIMMs			
90Y3221	A2QP	16 GB (2 Gb, 4Rx4, 1.35 V) PC3L-8500 DDR3-1066 Planar-X VLP	-
1333 MHz RDIMMs			
46C0560	A0WX	2 GB (1x 2 GB, 1Rx8, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	B1x, F1x
46C0563	A0WY	4 GB (1x 4 GB, 1Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	-
46C0564	A0WZ	4 GB (1x 4 GB, 2Rx8, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	B3x, C2x, C3x, F2x
00D4981	A3BT	8 GB (1x 8 GB, 1Rx4, 1.35V) PC3L-10600 CL9 ECC DDR3 1333MHz VLP	C4x
46C0568	8644	8 GB (1x 8 GB, 2Rx4, 1.35 V) PC3-10600 CL9 ECC DDR3 1333 MHz VLP	-
00D4985	A3BU	8 GB (1x 8 GB, 2Rx8, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	-
46C0599	2422	16 GB (1x 16 GB, 2Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz VLP	-
1600 MHz RDIMMs			
90Y3147	A1S0	4 GB (1x 4 Gb, 1Rx4, 1.5 V) PC3-12800 CL11 ECC DDR3 1600 MHz VLP	-
90Y3148	A1S1	4 GB (1x 4 Gb, 2Rx8, 1.5 V) PC3-12800 CL11 ECC DDR3 1600 MHz VLP	-
00D4989	A3BV	8 GB (1x 8 GB, 1Rx4, 1.5V) PC3-12800 CL11 ECC DDR3 1600 MHz VLP	-
90Y3149	A1S2	8 GB (1x 8 GB, 2Rx4, 1.5 V) PC3-12800 CL11 ECC DDR3 1600 MHz VLP	D1x, D3x
00D4993	A3BW	8 GB (1x 8 GB, 2Rx8, 1.5 V) PC3-12800 CL11 ECC DDR3 1600 MHz VLP	D2x, D4x
90Y3157	A3BS	16 GB (1x 16 GB, 2Rx4, 1.5 V) PC3-12800 CL11 ECC DDR3 1600 MHz VLP	-

Memory DIMM installation suggestions

Each processor has three memory channels, and there are two DIMMs per channel. The BladeCenter HS23E supports Very Low Profile (VLP) DDR3 memory RDIMMs.

The following rules apply when selecting the memory configuration:

- ▶ The server supports RDIMMs.
- ▶ Mixing 1.5 V and 1.35 V DIMMs in the same server is supported. In this case, all DIMMs operate at 1.5 V.
- ▶ The maximum number of ranks supported per channel is eight.
- ▶ The maximum quantity of DIMMs that can be installed in a server depends on the number of processors (six DIMMs with one processor installed, 12 DIMMs with two processors installed).
- ▶ All DIMMs in all processor memory channels operate at the same speed, which is determined as the lowest value of:
 - Memory speed that is supported by specific processor.
 - Lowest maximum operating speed for the selected memory configuration that depends on rated speed, as shown under the “Maximum operating speed” section in Table 3-94.

Table 3-94 Maximum memory speeds for HS23E

DIMM type →	RDIMM							
DIMM specification	Single rank			Dual rank			Quad rank	
Rank								
Rated speed	1333 MHz		1600 MHz	1333 MHz		1600 MHz	1066 MHz	
Rated voltage	1.35 V		1.5 V	1.35 V		1.5 V	1.35 V	
Operating voltage	1.35 V	1.5 V	1.5 V	1.35 V	1.5 V	1.5 V	1.35 V	1.5 V
Maximum quantity supported	12	12	12	12	12	12	12	12
Maximum DIMM capacity	8 GB	8 GB	8 GB	16 GB	16 GB	16 GB	16 GB	16 GB
Maximum memory capacity	96 GB	96 GB	96 GB	192 GB	192 GB	192 GB	192 GB	192 GB
Maximum operating speed								
1 DIMM per channel	1333 MHz	1333 MHz	1600 MHz	1333 MHz	1333 MHz	1600 MHz	800 MHz	800 MHz
2 DIMMs per channel	1333 MHz	1333 MHz	1600 MHz	1333 MHz	1333 MHz	1600 MHz	800 MHz	800 MHz

Terminology: The term *DIMMs per channel* is commonly abbreviated as DPC. One DIMM per channel can be referred to as 1DPC and, likewise, two DIMMs per channel can be referred to as 2DPC. The ability to run two DIMMs per channel at 1333 MHz is often referred to as the abbreviation 2DPC@1333MHz.

Memory modes

Depending on the memory mode that is set in the Setup utility, the blade server can support a minimum of 2 GB and a maximum of 96 GB of system memory on the system board in a blade server with one microprocessor. If two microprocessors are installed, the blade server can support a minimum of 4 GB and a maximum of 192 GB of system memory.

There are three memory modes that can be selected in the UEFI BIOS setup:

► Independent channel mode

Independent channel mode provides a maximum of 96 GB of usable memory with one microprocessor installed, and 192 GB of usable memory with two microprocessors installed (using 16 GB DIMMs). The DIMMs can be installed without matching sizes. Table 3-95 shows the memory installation order.

Table 3-95 DIMM population sequence for independent channel mode

Processors installed	DIMM connectors
One	5, 3, 1, 6, 4, and 2
Two	5, 11, 3, 9, 1, 7, 6, 12, 4, 10, 2, and 8

► Rank-sparing mode

In rank-sparing mode, one memory DIMM rank serves as a spare of the other ranks on the same channel. The spare rank is held in reserve and is not used as active memory. The spare rank must have identical or larger memory capacity than all the other active memory ranks on the same channel. After an error threshold is surpassed, 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.

Rank-sparing mode memory requirement: Rank-sparing mode is supported if the blade server meets one of the following memory requirements:

- One quad-rank DIMM
- More than one DIMM per channel
- An even number of single-rank or dual-rank DIMMs

For the rank-sparing mode memory installation order on a blade server with quad-rank DIMMs, see Table 3-95 on page 333.

Table 3-96 shows the order that single-rank or dual-rank DIMMs are installed to use rank-sparing mode.

Table 3-96 Rank-sparing mode DIMM installation sequence for single-rank or dual-rank DIMMs

DIMM pair	DIMM connectors	
	One processor installed	Two processors installed
First	5 and 6	5 and 6, 11 and 12
Second	3 and 4	3 and 4
Third	1 and 2	9 and 10
Fourth	None	1 and 2
Fifth	None	7 and 8

Important: In rank-sparing mode, if an installed DIMM does not meet the requirements listed, the system is executed in independent channel mode.

► Mirrored channel mode

In mirrored channel mode, memory is installed in pairs. Each DIMM in a pair must be identical in capacity, type, and rank count. The channels are grouped in pairs with each channel receiving the same data. One channel is used as a backup of the other, which provides redundancy. The memory contents on channel 2 are duplicated on channel 3. Channel 1 DIMM connectors 1, 2, 7, and 8 are not used in mirrored channel mode. The maximum available memory (with 16 GB DIMMs) is 32 GB for a single-microprocessor system and 64 GB for a dual-microprocessor system.

Important: The memory configuration of channel 2 must match the memory configuration of channel 3. For example, if a 4 GB dual-rank DIMM is installed into DIMM connector 3 (channel 2), a 4 GB dual-rank DIMM must also be installed into DIMM connector 5 (channel 3). Table 3-92 on page 331 lists each channel and which DIMM connectors belong to the channel.

Table 3-97 shows the order that memory DIMMs are installed to use mirrored channel mode.

Table 3-97 Mirrored channel mode DIMM installation sequence - single-rank or dual-rank DIMMs

DIMM pair	DIMM connectors	
	One processor installed	Two processors installed
First	3 and 5	3 and 5, 9 and 11
Second	4 and 6	4 and 6
Third	None	10 and 12

Important: In mirrored channel mode, if any installed DIMM does not meet the requirements listed, the system is executed in independent channel mode.

3.6.6 Standard onboard features

The HS23E blade server has the standard onboard features that are shown in the following sections.

UEFI

The HS23E uses an integrated Unified Extensible Firmware Interface (UEFI) next-generation BIOS.

UEFI capabilities include the following items:

- Human-readable event logs. *No more beep codes.*
- A complete setup solution by allowing the adapter configuration function to be moved into the UEFI.
- Complete out-of-band coverage by the Advance Settings Utility to simplify remote setup.

Using all of the features that UEFI offers requires a UEFI-aware operating system and adapters. UEFI is fully compatible with earlier versions of the BIOS.

For more information about UEFI, see *Introducing uEFI-Compliant Firmware on IBM System x and BladeCenter Servers*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5083207>

Onboard network adapter

The HS23E offers two integrated Gigabit Ethernet ports with the onboard Broadcom BCM5718 controller. These ports are routed to the chassis I/O bays 1 and 2.

The integrated dual-port NIC has the following features:

- ▶ Two Gigabit Ethernet ports with BCM5718 ASIC
- ▶ Full-duplex (FDX) capability
- ▶ IEEE 802.3x FDX flow control
- ▶ I/O virtualization support for VMware NetQueue and Microsoft VMQ
- ▶ Function Level Reset (FLR)
- ▶ IPv4/IPv6 offload:
 - TCP, IP, and UDP checksum offload
 - Large send offload (LSO)
 - TCP segmentation offload (TSO)
 - Receive-side scaling (RSS)
 - Transmit-side scaling (TSS)
- ▶ IEEE 802.1Q VLAN tagging
- ▶ VLAN insertion and extraction
- ▶ Jumbo frames up to 9,600 bytes
- ▶ Load balancing and failover teaming support, including adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), and IEEE 802.3ad
- ▶ Support for Serial over LAN (SoL) and concurrent KVM (cKVM)
- ▶ Preboot Execution Environment (PXE) support
- ▶ Wake On LAN support
- ▶ PCIe 2.0 x1 host interface
- ▶ Message Signal Interrupt (MSI-X) support

Integrated systems management processor

The HS23E type 8038 blade server includes an Integrated Management Module Version 2 (IMM V2) to monitor server availability, perform Predictive Failure Analysis, and trigger IBM Systems Director alerts. IMM V2 provides system management capabilities with IPMI 2.0 compliance and Serial over LAN.

For more information about the IMM, see *Transitioning to UEFI and IMM*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5079769>

Video controller

The video controller is Matrox G200e R2 video core with 16 MB video memory integrated into the IMM2. The maximum resolution is 1600x1200 at 75 Hz with 16 M colors (32 bits per pixel).

3.6.7 Local storage

The HS23E server has two hot-swap drive bays that are accessible from the front of the blade server. These bays are connected to the integrated ServeRAID C105 controller supplying basic software RAID capabilities for SATA hard drives (Figure 3-58 on page 336).




Figure 3-58 The HS23E showing the front hot-swap disk drive bays

The ServeRAID C105 onboard controller has the following specifications:

- ▶ Support for up to two internal hot-swap SATA HDDs (SAS HDDs and SSDs are not supported)
- ▶ Support for RAID 0 and RAID 1 (non-RAID is not supported)
- ▶ Three Gbps throughput per port
- ▶ Support for up to eight volumes
- ▶ Support for virtual drive sizes greater than 2 TB
- ▶ Fixed stripe unit size of 64 KB
- ▶ Support for MegaRAID Storage Manager management software

Important: ServeRAID C105 has no native (in-box) driver for Windows and Linux; the drivers must be downloaded separately. ServeRAID C105 is not supported by VMware, Hyper-V, Xen, and KVM hypervisors.

An optional ServeRAID H1135 Controller with hardware RAID capabilities is required to support SAS HDDs, SATA SSDs, and external SAS-based storage attachments.

For more information about the ServeRAID H1135, see 3.1.21, “ServeRAID H1135 (CIOv) Controller” on page 250.

Local SAS or SATA disk drives

The HS23E type 8038 has support for up to two hot-swap SATA hard disk drives with the ServeRAID C105 onboard controller, or support for up to two hot-swap SFF SAS or SATA

hard disk drives with the optional ServeRAID H1135 controller. These two hot-swap hard disk drives are accessible from the front of the blade server without removing the blade server from the BladeCenter chassis.

Table 3-98 lists supported SAS and SATA HDDs.

Table 3-98 Supported SAS and SATA HDDs

Part number	Feature code	Description
10K SAS hard disk drives		
90Y8877	A2XC	IBM 300GB 10K 6Gbps SAS 2.5" SFF G2HS HDD
90Y8872	A2XD	IBM 600GB 10K 6Gbps SAS 2.5" SFF G2HS HDD
00AJ300	A4VB	IBM 600GB 15K 6Gbps SAS 2.5" G2HS HDD
81Y9650	A282	IBM 900GB 10K 6Gbps SAS 2.5" SFF HS HDD
00AD075	A48S	IBM 1.2TB 10K 6Gbps SAS 2.5" G2HS HDD
15K SAS hard disk drives		
90Y8926	A2XB	IBM 146GB 15K 6Gbps SAS 2.5" SFF G2HS HDD
81Y9670	A283	IBM 300GB 15K 6Gbps SAS 2.5" SFF HS HDD
SAS hybrid hard disk drives		
00AJ236	A4VD	IBM 300GB 15K 6Gbps SAS 2.5" G2HS Hybrid
00AD102	A4G7	IBM 600GB 10K 6Gbps SAS 2.5" G2HS Hybrid
00AJ246	A4VF	IBM 600GB 15K 6Gbps SAS 2.5" G2HS Hybrid
NL SAS hard disk drives		
90Y8953	A2XE	IBM 500GB 7.2K 6Gbps NL SAS 2.5" SFF G2HS HDD
81Y9690	A1P3	IBM 1TB 7.2K 6Gbps NL SAS 2.5" SFF HS HDD
NL SATA hard disk drives		
81Y9722	A1NX	IBM 250GB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD
81Y9726	A1NZ	IBM 500GB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD
81Y9730	A1AV	IBM 1TB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD

Solid-state drives

Solid-state drives (SSDs) are more reliable than spinning hard disk drives, and they consume much less power than a standard SAS drive. Target applications for SSDs include video surveillance, transaction-based databases, and other applications that have high performance but moderate space requirements.

Table 3-99 lists the supported SSDs.

Table 3-99 Supported solid-state drives (SSDs)

Part number	Feature code	Description
Solid-state drives - Enterprise		
00W1125	A3HR	IBM 100GB SATA 2.5" MLC HS Enterprise SSD

Part number	Feature code	Description
49Y6129	A3EW	IBM 200GB SAS 2.5" MLC HS Enterprise SSD
49Y6134	A3EY	IBM 400GB SAS 2.5" MLC HS Enterprise SSD
49Y6139	A3F0	IBM 800GB SAS 2.5" MLC HS Enterprise SSD
49Y6195	A4GH	IBM 1.6TB SAS 2.5" MLC HS Enterprise SSD
41Y8331	A4FL	S3700 200GB SATA 2.5" MLC HS Enterprise SSD
41Y8336	A4FN	S3700 400GB SATA 2.5" MLC HS Enterprise SSD
41Y8341	A4FQ	S3700 800GB SATA 2.5" MLC HS Enterprise SSD
Solid-state drives - Enterprise value		
49Y5839	A3AS	IBM 64GB SATA 2.5" MLC HS Enterprise Value SSD
90Y8648	A2U4	IBM 128GB SATA 2.5" MLC HS Enterprise Value SSD
90Y8643	A2U3	IBM 256GB SATA 2.5" MLC HS Enterprise Value SSD
49Y5844	A3AU	IBM 512GB SATA 2.5" MLC HS Enterprise Value SSD
00AJ000	A4KM	S3500 120GB SATA 2.5" MLC HS Enterprise Value SSD
00AJ005	A4KN	S3500 240GB SATA 2.5" MLC HS Enterprise Value SSD
00AJ010	A4KP	S3500 480GB SATA 2.5" MLC HS Enterprise Value SSD
00AJ015	A4KQ	S3500 800GB SATA 2.5" MLC HS Enterprise Value SSD

Important: SSDs are supported with the optional Server RAID H1135 controller only.

PCIe SSD adapters

The HS23E server supports the High IOPS SSD adapters listed in Table 3-100. The adapters must be installed in an IBM BladeCenter PCI Express Gen 2 Expansion Blade II. Up to eight High IOPS adapters supported per one HS23E (two per PCI Express Gen 2 Expansion Blade II and up to four PCI Expansion Blades per HS23E)

Table 3-100 Supported PCIe SSD adapters

Part number	Feature code	Description	Slots supported
46C9078	A3J3	IBM 365GB High IOPS MLC Mono Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
46C9081	A3J4	IBM 785GB High IOPS MLC Mono Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4377	A3DY	IBM 1.2TB High IOPS MLC Mono Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)
90Y4397	A3DZ	IBM 2.4TB High IOPS MLC Duo Adapter	PCI Express Gen 2 Expansion Blade II (68Y7484)

3.6.8 Integrated virtualization

The HS23E offers an IBM standard USB flash drive option preinstalled with VMware ESXi. This version is an embedded version of VMware ESXi. It is fully contained on the flash drive, without requiring any disk space. The IBM USB Memory Key for VMware Hypervisor plugs into the USB port at the HS23E system board.

Table 3-101 lists the part number, feature code, and description of the IBM USB Memory Key for VMware Hypervisor supported in the HS23E.

Table 3-101 IBM USB Memory Key for VMware Hypervisor

Part Number	Feature Code	Description
41Y8287	3033	IBM USB Memory Key for VMware ESXi 4.1
41Y8300	A2VC	IBM USB Memory Key for VMWare ESXi 5.0
41Y8311	A2R3	IBM USB Memory Key for VMWare ESXi 5.1
41Y8382	A4WZ	IBM USB Memory Key for VMWare ESXi 5.1 Update 1
41Y8298	A2G0	IBM Blank USB Memory Key for VMWare ESXi Downloads

For a complete description of the features and capabilities of the VMware ESX Server, see the following address:

<http://www.vmware.com/products/vi/esx/>

3.6.9 I/O expansion

The HS23E type 8038 is able to connect to a wide variety of networks and fabrics by the installation of the appropriate I/O expansion card. Supported networks and fabrics include 1 Gb and 10 Gb Ethernet, 4 Gb and 8 Gb Fibre Channel, serial-attached SCSI (SAS), and InfiniBand.

An HS23E blade server with an installed I/O expansion card must be installed in a BladeCenter chassis with switch (or pass-through) modules that are compatible with the I/O expansion card.

The HS23E type 8038 blade server supports the following formats of I/O expansion cards: the CIOv format and the CFFh format, as described in the following sections.

CIOv

The CIOv I/O expansion connector provides I/O connections to modules in bays 3 and 4 of a supported BladeCenter chassis. The CIOv slot is a third-generation PCI Express x8 (PCIe 3.0 x8) slot. A maximum of one CIOv I/O expansion card is supported per single-wide blade server. A CIOv I/O expansion card can be installed on a blade server at the same time that a CFFh I/O expansion card is installed. Table 3-102 lists the CIOv expansion cards that are supported in the HS23E type 8038.

Table 3-102 Supported CIOv expansion cards for the HS23E server

Part number	Feature code	Description
44W4475	1039	Ethernet Expansion Card (CIOv)
46M6065	3594	QLogic 4 Gb Fibre Channel Expansion Card (CIOv)
90Y4750	A1XJ	ServeRAID H1135 Controller (CIOv)
46M6140	3598	Emulex 8 Gb Fibre Channel Expansion Card (CIOv)
44X1945	1462	QLogic 8 Gb Fibre Channel Expansion Card (CIOv)

For the latest information about the expansion cards that are supported by the HS23E, see the IBM ServerProven website:

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

CIOv expansion cards are installed in the CIOv slot in the HS23E. The CIOv expansion cards require a supported I/O module installed in bay 3 and bay 4 of the chassis where the cards and servers are installed.

CFFh

The CFFh I/O expansion connector provides I/O connections to high-speed switch modules in bays 7, 8, 9, and 10 of a BladeCenter H or BladeCenter HT chassis or to switch bay 2 in a BladeCenter S chassis. The CFFh slot is a third-generation PCI Express x16 (PCIe 3.0 x16) slot. A maximum of one CFFh I/O expansion card is supported per blade server. A CFFh I/O expansion card can be installed on a blade server at the same time that a CIOv I/O expansion card is installed. Table 3-103 lists the CFFh expansion cards that are supported in the HS23E type 8038.

Table 3-103 Supported CFFh expansion cards for the HS23E type 8038

Part number	Feature code	Description
44W4479	5476	2/4 Port Ethernet Expansion Card (CFFh)
46M6001	0056	Two port 40 Gb InfiniBand Expansion Card (CFFh)
46M6164	0098	Broadcom 10 Gb Gen2 4-port Ethernet Expansion Card (CFFh)
81Y3133	A1QR	Broadcom 2-port 10 Gb Virtual Fabric Adapter
81Y1650	5437	Brocade 2-port 10 GbE Converged Network Adapter
00y3264	A3NW	Emulex 10 GbE Virtual Fabric Adapter Advanced II
00Y3266	A3NV	Emulex 10 GbE Virtual Fabric Adapter II
42C1810	3593	Intel 10 Gb 2-port Ethernet Expansion Card (CFFh)
00Y3280 ^a	A3JB	QLogic 2-port 10 Gb Converged Network Adapter (CFFh)
00Y3332	A4AC	QLogic 10Gb Virtual Fabric Adapter (CFFh)
00Y5618	A4AD	QLogic 10 Gb Virtual Fabric Converged Network Adapter (CFFh)
00Y3270 ^b	A3JC	QLogic Eth and 8 Gb Fibre Channel Expansion Card (CFFh)

a. Replaces 42C1830.

b. Replaces 44X1940.

For the latest information about the expansion cards that are supported by the HS23E, see the IBM ServerProven website:

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

PCI Express Gen 2 Expansion Blade/Expansion Blade II

The HS23E supports the attachment of PCI Express Gen 2 Expansion Blades and PCI Express Gen 2 Expansion Blades II. The HS23E allows certain standard PCI-Express adapters to be attached to the HS23E. If there is a CFFh I/O expansion card already installed on the HS23E, it must be removed to allow for the attachment of the PCI Express Gen 2

Expansion Blade/Expansion Blade II. The CFFh I/O expansion card can then be installed in the topmost PCI Express Gen 2 Expansion Blade/Expansion Blade II. No other type of expansion blade can be mixed with the PCI Express Gen 2 Expansion Blade/Expansion Blade II.

For more information about the PCI Express Gen 2 Expansion Blades, see 3.10.2, “PCI Express Gen 2 Expansion Blade/Expansion Blade II” on page 403.

IBM BladeCenter GPU Expansion Blade/Expansion Blade II

The HS23E supports the attachment of up to four IBM BladeCenter GPU Expansion (BGE) Blades or IBM BladeCenter GPU Expansion Blades II. The HS23E is integrated with one NVIDIA Tesla M2075, NVIDIA Tesla M2070Q, or NVIDIA Tesla M2090 GPU adapter that can be attached to the HS23E. If there is a CFFh I/O expansion card already installed on the HS23E, it must be removed to allow for the attachment of the IBM BladeCenter GPU Expansion Blade/Expansion Blade II. The CFFh I/O expansion card can then be installed in the topmost IBM BladeCenter GPU Expansion Blade/Expansion Blade II. No other type of expansion blade can be mixed with the IBM BladeCenter GPU Expansion Blade/Expansion Blade II.

For more information, see 3.10.3, “IBM BladeCenter GPU Expansion Blade” on page 408.

3.7 IBM BladeCenter HX5

The IBM BladeCenter HX5 blade server showcases the eX5 architecture and technology in a blade form factor. This section introduces the server and describes the features and options that are currently announced.

The following topics are covered in this section:

- ▶ 3.7.1, “Introduction” on page 341
- ▶ 3.7.2, “Models” on page 344
- ▶ 3.7.3, “System architecture” on page 346
- ▶ 3.7.4, “Chassis support” on page 349
- ▶ 3.7.5, “IBM MAX5 V2 for BladeCenter” on page 350
- ▶ 3.7.6, “Scalability” on page 351
- ▶ 3.7.7, “Processors” on page 355
- ▶ 3.7.8, “Memory” on page 358
- ▶ 3.7.9, “Storage” on page 364
- ▶ 3.7.10, “I/O expansion cards” on page 367
- ▶ 3.7.11, “Standard onboard features” on page 370
- ▶ 3.7.12, “Integrated virtualization” on page 371
- ▶ 3.7.13, “Partitioning capabilities” on page 372

3.7.1 Introduction

The IBM BladeCenter HX5 type 7873 supports up to two of Intel Xeon E7 family of processors. The HX5 supports up to 40 DIMM modules with the addition of the MAX5 memory expansion blade.

Figure 3-59 shows the following three configurations:

- ▶ Single-wide HX5 with two processor sockets and 16 DIMM sockets
- ▶ Double-wide HX5 with four processor sockets and 32 DIMM sockets

- Double-wide HX5 with two processors and 40 DIMM sockets: 16 in the HX5 server and 24 in the attached MAX5 memory expansion blade

Important: The MAX5 can connect only to a single HX5 server. Only one MAX5 can be connected per HX5 server.



Figure 3-59 IBM BladeCenter HX5 blade server configurations

Table 3-104 lists the features of the HX5 7873.

Table 3-104 Features of the HX5 type 7873

Features	HX5 2-socket	HX5 4-socket	HX5 2-socket with MAX5
Form Factor	30 mm (1-wide)	60 mm (2-wide)	60 mm (2-wide)
Maximum number of processors	Two	Four	Two
Processor options	Intel Xeon E7-8800, E7-4800 and E7-2800: 6-core, 8-core, and 10-core	Intel Xeon E7-8800 and E7-4800: 6-core, 8-core, and 10-core	Intel Xeon E7-8800, E7-4800 and E7-2800: 6-core, 8-core, and 10-core
Cache	18 MB or 24 MB or 30 MB (shared between cores) (processor dependent)		
Memory Speed	1066, 978, or 800 MHz (processor SMI link speed dependant)		
Memory (DIMM slots / maximum)	16 DIMM slots / 512 GB, maximum of 32 GB DIMMs	32 DIMM slots / 1024 GB, maximum of 32 GB DIMMs	40 DIMM slots / 1.28 TB, maximum of 32 GB DIMMs
Memory Type	DDR 3 ECC Very Low Profile (VLP) Registered DIMMs		
Internal storage	Optional 1.8-inch solid-state drives (SSDs); non-hot-swap (require an additional SSD carrier)		
Maximum number of drives	Two	Four	Two
Maximum Internal Storage	Up to 800 GB using two 400 GB SSDs	Up to 1.6 TB using four 400 GB SSDs	Up to 800 GB using two 400 GB SSDs
I/O expansion	<ul style="list-style-type: none"> ► One CIOv ► One CFFh 	<ul style="list-style-type: none"> ► Two CIOv ► Two CFFh 	<ul style="list-style-type: none"> ► One CIOv ► One CFFh

Figure 3-60 shows the components on the system board of the HX5.




Figure 3-60 Layout of HX5 (two-node four-socket configuration shown)

The MAX5 memory expansion blade is a device with the same dimensions as the HX5 (Figure 3-61). When the MAX5 is attached to the HX5, the combined unit occupies two blade bays in the BladeCenter chassis.




Figure 3-61 Layout of MAX5 system board

3.7.2 Models

Table 3-105 shows the base models of the BladeCenter HX5 type 7873, with and without the MAX5 memory expansion blade.

Table 3-105 Models of the HX5 type 7873

Model ^a	Intel Xeon model and cores/max	TDP per CPU	HX5 max memory speed	MAX5 memory speed	Scalable to four socket	MA X5 V2	10 GbE card	Standard memory ^b
Standard HX5 models								
7873-F5x	1x E7-4870 10C 2.40 GHz	130 W	1066 MHz	1066 MHz	Yes	Opt ^c	Std ^d	2x 4 GB
7873-D1x	1x E7-8867L 10C 2.13 GHz	105 W	1066 MHz	1066 MHz	Yes	Opt	Opt	2x 4 GB
7873-B2x	1x E7-4830 8C 2.13 GHz	105 W	1066 MHz	1066 MHz	Yes	Opt	Opt	2x 4 GB
7873-F4x	1x E7-4830 8C 2.13 GHz	105 W	1066 MHz	1066 MHz	Yes	Opt	Std ^d	2x 4 GB
7873-C1x	1x E7-8837 8C 2.67 GHz	130 W	978 MHz	978 MHz	Yes	Opt	Opt	2x 4 GB
7873-B1x	1x E7-4807 6C 1.86 GHz	95 W	800 MHz	800 MHz	Yes	Opt	Opt	2x 4 GB
HX5 models with MAX 5 V2 included as standard								
7873-A3x	2x E7-2870 10C 2.40 GHz	130 W	1066 MHz	1066 MHz	No	Std	Opt	HX5: 4x 4 GB MAX5: None
7873-A2x	2x E7-2860 10C 2.26 GHz	130 W	1066 MHz	1066 MHz	No	Std	Opt	HX5: 4x 4 GB MAX5: None
7873-A1x	2x E7-2830 8C 2.13 GHz	105 W	1066 MHz	1066 MHz	No	Std	Opt	HX5: 4x 4 GB MAX5: None

a. This column lists worldwide generally available variant (GAV) model numbers. They are not orderable as listed and must be modified by country. The US GAV model numbers use the following nomenclature: xxU. For example, the US orderable part number for 7873-A2x is 7873-A2U. See the product-specific official IBM announcement letter for other country-specific GAV model numbers.

b. The HX5 type 7873 has 16 DIMM sockets and can hold 512 GB using 32 GB memory DIMMs. The MAX5 V2 has 24 DIMM sockets and can hold 768 GB using 32 GB memory DIMMs. A 1-node HX5 + MAX5 V2 supports 1.28 TB total using 32GB DIMMs.

c. Option (Opt)

d. Emulex Virtual Fabric Adapter Advanced II (CFFh); standard (Std)

Workload optimized models

There are also workload-optimized models of the HX5 available. These preconfigured and pre-tested models are targeted at specific workloads.

Workload optimized models for IBM zEnterprise BladeCenter Extension

The 7873-AAx, -ABx, -ACx, and -ADx are the models for IBM zEnterprise® BladeCenter Extension (zBX) and include the following features in addition to standard HX5 features:

- ▶ Two E7-2830 processors
- ▶ A total of 64, 128, 192 or 256 GB of available memory (depends on the model)
- ▶ Broadcom 2-port 10Gb Virtual Fabric Adapter for IBM BladeCenter (CFFh)
- ▶ QLogic 8 Gb Fibre Channel Expansion Card (CIOv)
- ▶ SSD Expansion Card for IBM BladeCenter HX5
- ▶ Two IBM 100 GB SATA 1.8-in. MLC solid-state drives (SSDs)

Table 3-106 describes the models for zBX.

Table 3-106 HX5 type 7873 models for zBX

Model	Intel Xeon processor	Scalable to four socket	MAX5 V2	Standard memory	CFFh card ^a	CIOv card ^b	SSDs (std/max) ^c
Models for zBX							
7873-AAx	2x Xeon E7-2830 8C 2.13 GHz	No	Opt	8x 8 GB	Std	Std	2/2
7873-ABx	2x Xeon E7-2830 8C 2.13 GHz	No	Opt	16x 8 GB	Std	Std	2/2
7873-ACx	2x Xeon E7-2830 8C 2.13 GHz	No	Opt	8x 8 GB + 16x 8 GB	Std	Std	2/2
7873-ADx	2x Xeon E7-2830 8C 2.13 GHz	No	Opt	16x 16 GB	Std	Std	2/2

a. Broadcom 2-port 10Gb Virtual Fabric Adapter for IBM BladeCenter (CFFh).

b. QLogic 8 Gb Fibre Channel Expansion Card (CIOv).

c. Two BM 100GB SATA 1.8" MLC Enterprise solid state drives.

Workload optimized models for IBM BladeCenter Foundation for Cloud

The 7873-91x, and -92x models are workload optimized models for IBM BladeCenter Foundation for Cloud and include the following features in addition to standard HX5 features:

- ▶ Two E7-2830 processors
- ▶ Eight 8 GB memory DIMMs for a total of 64 GB of available memory
- ▶ One dual port Emulex Virtual Fabric Adapter Advanced II (CFFh) or QLogic 2-port 10 Gb Converged Network Adapter (CFFh) (depends on the model)
- ▶ VMware ESXi 4.1 Update 1
- ▶ IBM Systems Director Standard Edition V6 with 1-year support and subscription

Table 3-107 describes the Foundation for Cloud models.

Table 3-107 Foundation for Cloud models of the HX5 type 7873

Model	Intel Xeon processor	Scalable to four socket	MAX5 V2	Standard memory	CFFh card	CIOv card	SSDs (std/max)
Foundation for Cloud models							
7873-91x	2x Xeon E7-2830 8C 2.13 GHz	No	No	8x 8 GB	Std ^a	Opt	0/2
7873-92x	2x Xeon E7-2830 8C 2.13 GHz	No	No	8x 8 GB	Std ^b	Opt	0/2

a. Emulex Virtual Fabric Adapter Advanced II.

b. QLogic 2-port 10 Gb Converged Network Adapter.

Models for 2-node (4-socket) HX5 solutions

When configuring a 2-node (4-socket) HX5 using the Standalone Solution Configuration Tool (SSCT), you must choose the first 2-socket node of 4-socket server. The second 2-socket node is added automatically by SSCT. The second node has the same processor, memory modules, and adapters as the first one and includes IBM HX5 2-Node Scalability Kit as a standard.

Table 3-108 describes the HX5 models for a 4-socket configurations.

Table 3-108 Models of the HX5 type 7873 for 2-node configurations

Model	Intel Xeon processor	Scalability kit as standard	Standard memory	CFFh card	CIOv card	SSDs (std/max)
Models for 2-node configurations						
7873-BAx	1x Xeon E7-4807 6C 1.86 GHz	Yes	2x 4 GB	Opt	Opt	0/2
7873-BHx	1x Xeon E7-4807 6C 1.86 GHz	Connect to model BAx	2x 4 GB	Opt	Opt	0/2
7873-BBx	1x Xeon E7-4830 8C 2.13 GHz	Yes	2x 4 GB	Opt	Opt	0/2
7873-BJx	1x Xeon E7-4830 8C 2.13 GHz	Connect to model BBx	2x 4 GB	Opt	Opt	0/2
7873-CAx	1x Xeon E7-8837 8C 2.67 GHz	Yes	2x 4 GB	Opt	Opt	0/2
7873-CHx	1x Xeon E7-8837 8C 2.67 GHz	Connect to model CAx	2x 4 GB	Opt	Opt	0/2
7873-DAx	1x Xeon E7-8867L 10C 2.13 GHz	Yes	2x 4 GB	Opt	Opt	0/2
7873-DHx	1x Xeon E7-8867L 10C 2.13 GHz	Connect to model DAx	2x 4 GB	Opt	Opt	0/2
7873-FEx	1x Xeon E7-4870 10C 2.4 GHz	Yes	2x 4 GB	Std ^a	Opt	0/2
7873-FNx	1x Xeon E7-4870 10C 2.4 GHz	Connect to model FBx	2x 4 GB	Std ^a	Opt	0/2

a. Emulex 10 GbE Virtual Fabric Adapter Advanced II.

3.7.3 System architecture

The Intel Xeon E7-8800, E7-4800, and E7-2800 processors in the HX5 have up to 10 cores with 20 threads per socket. The processors have up to 30 MB of shared L3 cache, Hyper-Threading, several with Turbo Boost, four QPI Links, one integrated memory controller, and up to four buffered SMI channels.

The HX5 type 7873 2-socket has the following system architecture features as standard:

- ▶ Two 1567-pin LGA processor sockets
- ▶ An Intel 7500 *Boxboro-EX* chip set
- ▶ An Intel ICH10 south bridge
- ▶ Eight second-generation Intel Scalable Memory Buffers, each with two memory channels
- ▶ One DIMM per memory channel
- ▶ Sixteen DDR3 DIMM sockets
- ▶ One Broadcom BCM5709S dual-port Gigabit Ethernet controller

- ▶ One Integrated Management Module
- ▶ One Trusted Platform Module 1.2 Controller
- ▶ One PCI Express x16 CFFh I/O expansion connector
- ▶ One PCI Express x16 CFFh-style connector for use with the SSD Expansion Card for and solid-state drives, or a future I/O card type
- ▶ One CIOv I/O expansion connector
- ▶ Scalability connector
- ▶ One internal USB port for embedded virtualization

Figure 3-62 shows the HX5 block diagram.

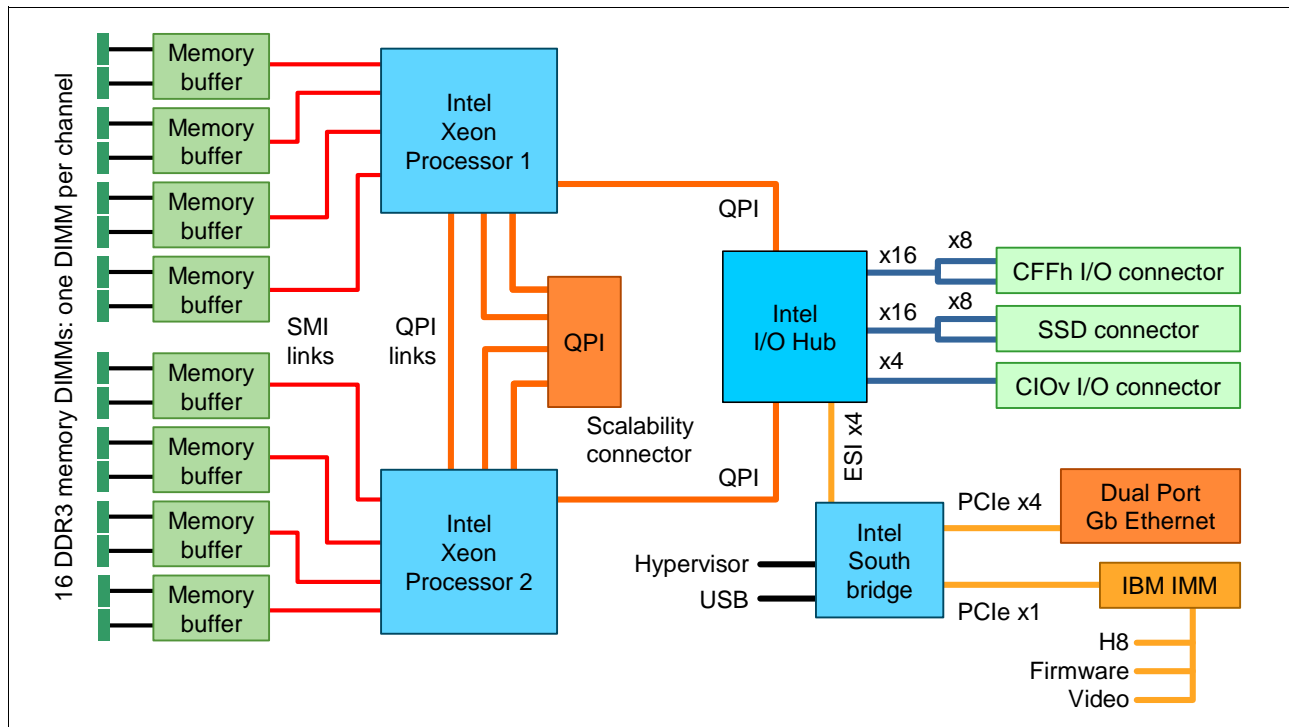


Figure 3-62 HX5 block diagram

IBM HX5 1-Node Speed Burst Card

To increase performance in a 2-socket HX5 server (with two processors installed), install the IBM HX5 1-Node Speed Burst Card. The HX5 1-Node Speed Burst Card takes the QPI links that would be used for scaling two HX5 2-socket blades and routes them back to the processors on the same blade. Table 3-109 lists the ordering information.

Table 3-109 HX5 1-Node Speed Burst Card

Part number	Feature code	Description
59Y5889	1741	IBM HX5 1-Node Speed Burst Card

Figure 3-63 shows the block diagram of how the Speed Burst Card attaches to the system.

Important: The Speed Burst Card is not required for an HX5 with only one processor installed. It is also not needed for a 2-node configuration (a separate card is available for a 2-node configuration, as described in 3.7.6, “Scalability” on page 351).

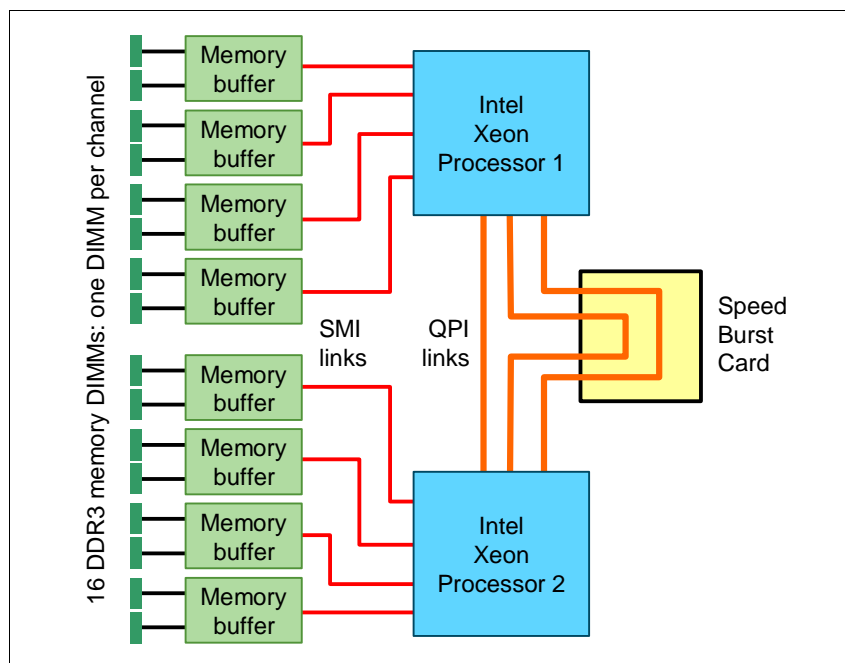


Figure 3-63 HX5 1-Node Speed Burst Card block diagram

Figure 3-64 shows where the Speed Burst Card is installed in the HX5.




Figure 3-64 Installing the Speed Burst Card

3.7.4 Chassis support

The HX5 is supported in BladeCenter chassis S, H, and HT (Table 3-111).

Table 3-110 HX5 chassis compatibility

Description	BC-E 8677	BC-S 8886	BC-H 8852	BC-HT AC 8750	BC-HT DC 8740
HX5 server	No	Yes	Yes ^a	Yes	Yes
HX5+MAX5 server	No	Yes	Yes ^a	Yes	Yes

a. See Table 3-111.

The number of HX5 servers supported in each chassis depends on the thermal design power of the processors used in the HX5 servers (Table 3-111):

- ▶ A green cell (■) means that the chassis can be filled with HX5 blade servers up to the maximum number of blade bays in the chassis (for example, 14 blades in the BladeCenter H).
- ▶ A yellow cell (■) means that the maximum number of HX5 blades that the chassis can hold is fewer than the total available blade bays (for example, 12 in a BladeCenter H). *All other bays must remain empty.* Empty bays must be distributed evenly between the two power domains of the chassis (bays 1 - 6 and bays 7 - 14).

Table 3-111 HX5 chassis compatibility

Server	Thermal design power (TDP) of the CPUs	Maximum number of servers supported in each chassis							
		BC-S (8886)	BC-H (models other than 4Tx)				BC-H (-4Tx)	BC-HT DC (8740)	BC-HT AC (8750)
			2900 W supplies		2980 W supplies ^a				
			Std. blower	Enh. blower ^b	Std. blower	Enh. blower ^b	Enh. blower ^b		
HX5 1-node (30 mm)	95 W, 105 W	5	14	14	14	14	14	10	10
	130 W	4	None ^c	10	None ^c	12	12	6	8
HX5 2-node (60 mm)	95 W, 105 W	2	7	7	7	7	7	5	5
	130 W	2	None ^c	5	None ^c	6	6	3	4
HX5 1-socket + MAX5 (60 mm)	95 W, 105 W	2	7	7	7	7	7	5	5
	130 W	2	6	6	7	7	7	4	5

a. IBM BladeCenter H 2980W AC Power Modules, 68Y6601 (standard in 4Tx, optional with all other BC-H chassis models)

b. IBM BladeCenter H Enhanced Cooling Modules, 68Y6650 (standard in 4Tx, optional with all other BC-H chassis models)

c. Not supported

NEBS compliance: The HX5 is currently not a NEBS-compliant offering in the BC-HT.

3.7.5 IBM MAX5 V2 for BladeCenter

The IBM MAX5 V2 for BladeCenter (Figure 3-61 on page 343) is a memory expansion blade that attaches to a single HX5 2-socket blade server.

The MAX5 V2 have the following system architecture features:

- ▶ IBM EXA memory controller.
- ▶ 24 DIMM slots/six memory buffers.
- ▶ Four DIMM slots per memory buffer (two per channel).
- ▶ Supports VLP DDR3 memory DIMMs.
- ▶ Attaches to a single HX5 using the IBM HX5 MAX5 1-node Scalability kit, 59Y5877, as described in “1-node HX5 with MAX5” on page 354.
- ▶ Communicates with the processors on the HX5 using high-speed QPI links.
- ▶ Supports low voltage (1.35 V) DIMM modules.

The MAX5 V2 is standard with certain models (Table 3-105 on page 344), and for other models it is available as an option (Table 3-112).

Table 3-112 IBM MAX5 for BladeCenter

Part number	Feature code	Description
88Y6128	A16N	IBM MAX5 V2 for BladeCenter
59Y5877	1742	IBM HX5 MAX5 1-node Scalability kit

Table 3-113 lists the available memory DIMM for the MAX5 V2.

Table 3-113 Memory options for the MAX5 V2

Part number	Feature code	Description
46C0560	A0WX	2 GB (1x 2 GB, 1Rx8, 1.35 V) PC3-10600 CL9 ECC DDR3 VLP
46C0564	A0WZ	4 GB (1x 4 GB, 2Rx8, 1.35 V) PC3-10600 CL9 ECC DDR3 VLP
46C0570	A17Q	8 GB (1x 8 GB, 4Rx8, 1.35 V) PC3-8500 CL7 ECC DDR3 VLP
00D4985	A3BU	8GB (1x8GB, 2Rx8, 1.35V) PC3L-10600 CL9 ECC DDR3 VLP
46C0599	2422	16 GB (1x 16 GB, 4Rx8, 1.35 V) PC3-10600 CL9 ECC DDR3 VLP
00D5008	A3KN	32GB (1x32GB, 4Rx4, 1.35V) PC3L-10600 CL9 ECC DDR3 VLP

The MAX5 V2 is composed of the EX5 node controller chip, six memory buffers, and 24 DIMM sockets. The MAX5 has three power domains: A, B, and C. Each power domain includes two memory controllers and eight DIMM sockets. Figure 3-65 shows the layout of the MAX5 V2.




Figure 3-65 MAX5 memory expansion blade

How the MAX5 connects to the HX5 is described in 3.7.6, “Scalability” on page 351. Memory options and rules are described in 3.7.8, “Memory” on page 358.

3.7.6 Scalability

In this section, we explain how the HX5 can be expanded to increase the number of processors and the number of memory DIMMs.

The HX5 blade architecture allows for several scalable configurations, including the use of a MAX5 memory expansion blade, but the blade currently supports only three configurations (Figure 3-66 on page 352):

- ▶ A single HX5 server with two processor sockets: This blade is a standard 30 mm blade, also known as *single wide* server or *single-node* server.
- ▶ Two HX5 servers connected together to form a single image four-socket server: This server is a 60 mm blade, also known as *double wide* server or *2-node* server.
- ▶ A single HX5 server with two processor sockets plus a MAX5 memory expansion blade attached to it, resulting in a 60 mm blade configuration. This configuration is sometimes referred to as a *1-node+MAX5* configuration.




Figure 3-66 IBM BladeCenter HX5 scalability

We describe each of these configurations in the following subsections. Supported BladeCenter chassis for each configuration are listed in 3.7.4, “Chassis support” on page 349.

Single HX5 configuration

This is the base configuration and supports one or two processors installed in the single-wide 30 mm server.

When the server has two processors installed, ensure that the server has the HX5 1-Node Speed Burst Card installed, as described in “IBM HX5 1-Node Speed Burst Card” on page 347. Table 3-114 provides the ordering information.

Table 3-114 HX5 1-Node Speed Burst Card

Part number	Feature code	Description
59Y5889	1741	IBM HX5 1-Node Speed Burst Card

Important: The HX5 1-Node Speed Burst Card is not required for an HX5 with only one processor installed. It is also not needed for a 2-node configuration (a separate card is available for 2-node).

2-node HX5 configuration

In the 2-node configuration, the two HX5 servers are physically connected together and a 2-node scalability card is attached to the side of the blades, which provides the path for the QPI scaling.

You can join two HX5 servers in a 2-node configuration, with each HX5 server having only one processor and using only eight DIMM slots per node. However, this configuration is not the best usage of the full HX5 server capabilities, so for performance reasons, consider having two processors installed in each node.

The two servers are connected together using a 2-node scalability card (Figure 3-67). The scalability card is immediately next to the processors and provides a direct connection between the processors in the two nodes.

Figure 3-67 shows the double-wide configuration, which is two HX5 servers connected together.




Figure 3-67 Two-node HX5 with the 2-node scalability card indicated

This configuration consumes two blade slots and has the 2-node scalability card attached. The scaling is performed through QPI scaling. The 2-node scalability card is not included with the server and must be ordered separately as part of the IBM HX5 2-Node Scalability Kit listed in Table 3-115.

Table 3-115 HX5 2-Node Scalability Kit

Part number	Feature code	Description
46M6975	1737	IBM HX5 2-Node Scalability Kit

The HX5 2-Node Scalability Kit contains the 2-node scalability card plus the necessary hardware to physically attach the two HX5 servers together.

Figure 3-68 shows the block diagram of a 2-node HX5 and the location of the HX5 2-node scalability card.

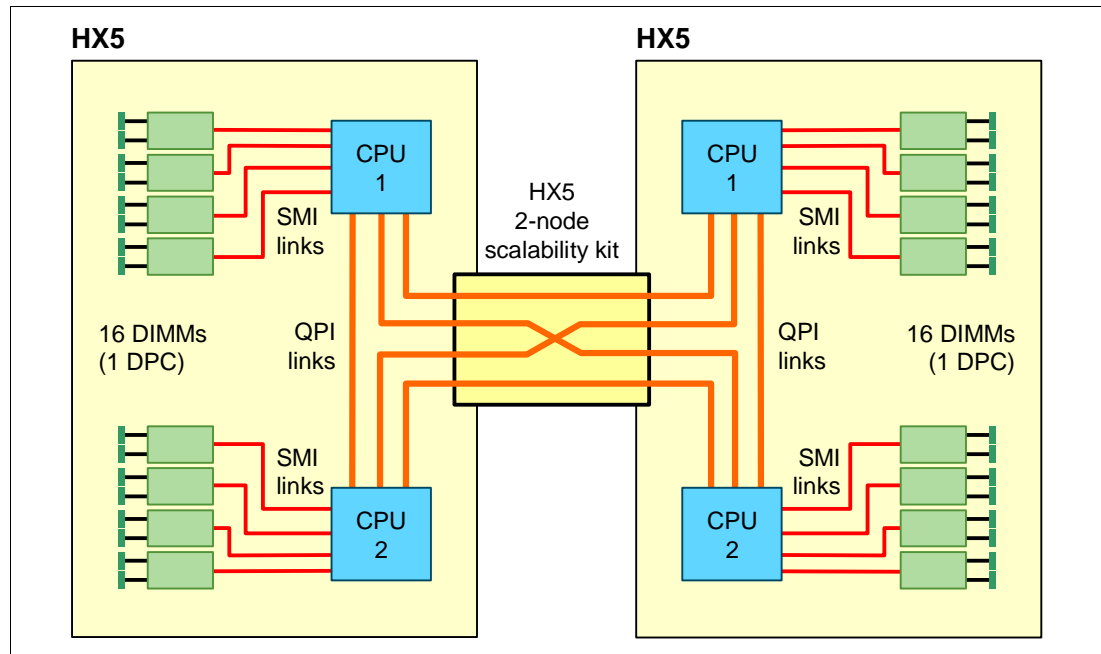


Figure 3-68 Block diagram of a 2-node HX5

1-node HX5 with MAX5

In the HX5+MAX5 configuration, the HX5 and MAX5 units are connected together through a 1-node MAX5 scalability card, which provides QPI scaling. This configuration gives the HX5 with the MAX5 the ability to have 40 DIMM modules installed, when the HX5 has two processors installed.

You can also attach the MAX5 to an HX5, where the HX5 has only one processor installed, which means that eight DIMMs modules in the HX5 and 24 DIMMs modules can be installed, for a total of 32 DIMM modules.



Figure 3-69 1-node HX5 + MAX5

The card used to connect the MAX5 to the HX5 is the IBM HX5 MAX5 1-node Scalability Kit, which is similar in physical appearance to the 2-node Scalability Kit (Figure 3-67 on page 353). Table 3-116 lists the ordering information.

Table 3-116 HX5 2-Node Scalability Kit

Part number	Feature code	Description
59Y5877	1742	IBM HX5 MAX5 1-node Scalability Kit

Figure 3-70 shows a block diagram of a single-node HX5 with MAX5.

Important: MAX5 can only connect to a single HX5 server. A configuration of two MAX5 units connected to a 2-node HX5 is not supported.

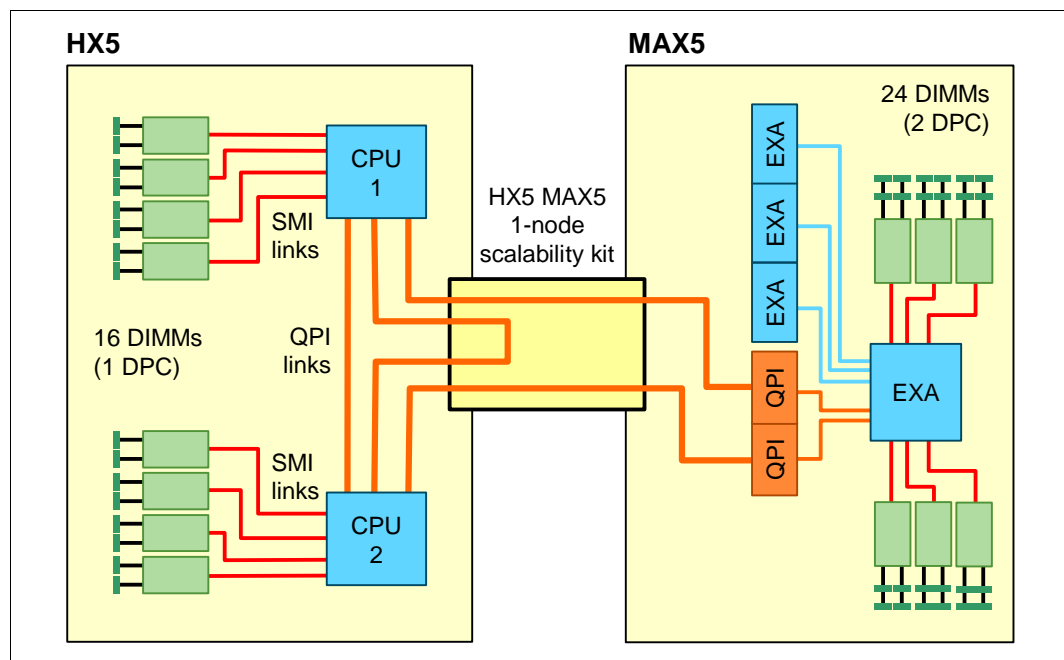


Figure 3-70 HX5 1-node with MAX5 block diagram

3.7.7 Processors

The HX5 type 7873 supports Intel Xeon E7-8800, E7-4800, and E7-2800 six-core, eight-core, or 10-core processors. The processors must be identical in both two-socket and four-socket configurations. The Intel Xeon processors are available in various clock speeds and have standard and lower power offerings.

Table 3-117 lists the processor options for the HX5 and the models that include them (if one exists).

Table 3-117 Available processor options

Part number	Feature code	Description	Can scale to two-node	Supported model
88Y6124	A17P	Xeon E7-8867L 10C 2.13 GHz 30 MB 1066 MHz 105 W	Yes	D1x
88Y6112	A17M	Xeon E7-8837 8C 2.67 GHz 24 MB 1066 MHz 130 W	Yes	C1x
88Y6160	A18W	Xeon E7-4870 10C 2.40 GHz 30 MB 1066 MHz 130 W	Yes	F5x
88Y6102	A17K	Xeon E7-4860 10C 2.26 GHz 24 MB 1066MHz 130 W	Yes	CTO ^a
88Y6092	A17H	Xeon E7-4850 10C 2.00 GHz 24 MB 1066 MHz 130 W	Yes	CTO ^a
88Y6082	A17G	Xeon E7-4830 8C 2.13 GHz 24 MB 1066 MHz 105 W	Yes	B2x, F4x
88Y6076	A17F	Xeon E7-4820 8C 2.00 GHz 18 MB 978 MHz 105 W	Yes	CTO ^a
88Y6070	A17E	Xeon E7-4807 6C 1.86 GHz 18 MB 800 MHz 95 W	Yes	B1x
88Y6150	A18U	Xeon E7-2870 10C 2.40 GHz 30 MB 1066 MHz 130 W	No	A3x
69Y3094	A17C	Xeon E7-2860 10C 2.26 GHz 24 MB 1066 MHz 130 W	No	A2x
69Y3084	A17A	Xeon E7-2850 10C 2.00 GHz 24 MB 1066 MHz 130 W	No	CTO ^a
69Y3074	A179	Xeon E7-2830 8C 2.13 GHz 24 MB 1066 MHz 105 W	No	A1x
69Y3068	A178	Xeon E7-2820 8C 2.00 GHz 18 MB 978 MHz 105 W	No	CTO ^a
69Y3062	A177	Xeon E7-2803 6C 1.73 GHz 18 MB 800 MHz 105 W	No	CTO ^a

a. The processor is not available in any of the standard models listed in Table 3-105 on page 344. The processor is, however, available through the configure-to-order (CTO) ordering process.

Table 3-118 lists the capabilities of each processor option available for the HX5.

Table 3-118 Intel Xeon E7-8800, E7-4800, and E7-2800 features

Processor model/cores	Scalable to 4-socket	Processor frequency	Turbo ^a	HT ^b	L3 cache	Power	QPI speed ^c	HX5 memory speed	MAX5 V2 memory speed
E7-8800 Series processors									
E7-8867L 10C	Yes	2.13 GHz	Yes	Yes	30 MB	105 W	6.4 GTps	1066 MHz	1066 MHz
E7-8837 8C	Yes	2.67 GHz	Yes	No	24 MB	130 W	6.4 GTps	1066 MHz	1066 MHz
E7-4800 Series processors									
E7-4870 10C	Yes	2.40 GHz	Yes	Yes	30 MB	130 W	6.4 GTps	1066 MHz	1066 MHz
E7-4860 10C	Yes	2.26 GHz	Yes	Yes	24 MB	130 W	6.4 GTps	1066 MHz	1066 MHz
E7-4850 10C	Yes	2.00 GHz	Yes	Yes	24 MB	130 W	6.4 GTps	1066 MHz	1066 MHz
E7-4830 8C	Yes	2.13 GHz	Yes	Yes	24 MB	105 W	6.4 GTps	1066 MHz	1066 MHz
E7-4820 8C	Yes	2.00 GHz	Yes	Yes	18 MB	105 W	5.86 GTps	978 MHz	978 MHz
E7-4807 6C	Yes	1.86 GHz	No	Yes	18 MB	95 W	4.8 GTps	800 MHz	800 MHz

Processor model/cores	Scalable to 4-socket	Processor frequency	Turbo ^a	HT ^b	L3 cache	Power	QPI speed ^c	HX5 memory speed	MAX5 V2 memory speed
E7-2800 Series processors									
E7-2870 10C	No	2.40 GHz	Yes	Yes	30 MB	130 W	6.4 GTps	1066 MHz	1066 MHz
E7-2860 10C	No	2.26 GHz	Yes	Yes	24 MB	130 W	6.4 GTps	1066 MHz	1066 MHz
E7-2850 10C	No	2.00 GHz	Yes	Yes	24 MB	130 W	6.4 GTps	1066 MHz	1066 MHz
E7-2830 8C	No	2.13 GHz	Yes	Yes	24 MB	105 W	6.4 GTps	1066 MHz	1066 MHz
E7-2820 8C	No	2.00 GHz	Yes	Yes	18 MB	105 W	5.86 GTps	978 MHz	N/A ^d
E7-2803 6C	No	1.73 GHz	No	Yes	18 MB	105 W	4.8 GTps	800 MHz	N/A ^d

a. Intel Turbo Boost Technology.

b. Intel Hyper-Threading Technology.

c. GTps: Giga-transfers per second.

d. Connecting a MAX5 is not supported with servers that have the E7-2820 or E7-2803 processors installed.

Socket support: As shown in Table 3-118, the Xeon E7-2800 series processor range does not support scaling to four sockets. The E7-2820 and the E7-2803 also do not support the attachment of the MAX5. This is a technical limitation of these particular processors.

Memory speed in the HX5 is dependent on the SMI link of the processor (the SMI link speed is listed in Table 3-118 on page 356 as a GTps value):

- ▶ If the SMI link speed is 6.4 GTps, the memory in the HX5 can operate at a maximum of 1066 MHz.
- ▶ If the SMI link speed is 5.86 GTps, the memory in the HX5 can operate at a maximum of 978 MHz.
- ▶ If the SMI link speed is 4.8 GTps, the memory in the HX5 can operate at a maximum of 800 MHz.

Table 3-118 on page 356 shows these memory speeds.

If a MAX5 V2 memory expansion blade is installed, the memory in the MAX5 can operate as high as 1066 MHz, depending on the DIMMs installed. The MAX5 memory speed is independent of the HX5 memory speed.

When installing two processors in a single HX5 server (without MAX5), add the IBM HX5 1-Node Speed Burst Card, 59Y588. For details about this feature, see “IBM HX5 1-Node Speed Burst Card” on page 347. Table 3-119 lists the option information.

Table 3-119 HX5 1-Node Speed Burst Card

Part number	Feature code	Description
59Y5889	1741	IBM HX5 1-Node Speed Burst Card

3.7.8 Memory

The HX5 2-socket has eight DIMM sockets per processor (a total of 16 DIMM sockets), supporting up to 256 GB of memory when using 16 GB DIMMs. With the addition of the MAX5 V2 memory expansion blade, a single HX5 7873 blade has access to a total of 40 DIMM sockets supporting up to 640 GB of memory when using 16 GB DIMMs.

Tip: The HX5 7873 supports DIMM capacities up to 16 GB; however, the HX5 7872 supports only up to 8 GB DIMMs.

The HX5 and MAX5 V2 use registered Double Data Rate-3 (DDR-3) very-low-profile (VLP) DIMMs and provides ECC and advanced Chipkill memory protection. The HX5 type 7873 and the MAX5 V2 support low voltage 1.35 V DIMM modules.

Memory options

Table 3-120 lists the memory options for the HX5 and MAX5 V2, available for order at time of writing.

Tip: Memory must be installed in pairs of two identical DIMMs, or in quads if memory mirroring is enabled. The options in Table 3-120, however, are for single DIMMs.

Table 3-120 Memory options for HX5 type 7873, and MAX5 V2

Part number	Feature code	HX5 Support, 7873	MAX5 V2 Support, 88Y6128	Description
46C0560	A0WX	Yes	Yes	2 GB (1x2 GB, 1Rx8, 1.35 V) PC3-10600 CL9 ECC VLP
46C0564	A0WZ	Yes	Yes	4 GB (1x4 GB, 2Rx8, 1.35 V) PC3-10600 CL9 ECC VLP
46C0570	A17Q	Yes	Yes	8 GB (1x8 GB, 4Rx8, 1.35 V) PC3-8500 CL7 ECC VLP
00D4985	A3BU	Yes	Yes	8GB (1x8GB, 2Rx8, 1.35V) PC3L-10600 CL9 ECC VLP
46C0599	2422	Yes	Yes	16 GB (1x16 GB, 4Rx8, 1.35 V) PC3-10600 CL9 ECC VLP
90Y3221	A2QP	Yes	No	16GB (1x16GB, 4Rx4, 1.35V) PC3L-8500 CL7 ECC VLP
00D5008	A3KN	Yes	Yes	32GB (1x32GB, 4Rx4, 1.35V) PC3L-10600 CL9 ECC VLP

For optimal performance, all DIMM slots must be populated.

Each processor controls eight DIMMs and four memory buffers in the server (Figure 3-71). To use all 16 DIMM sockets, both processors must be installed. If only one processor is installed, only eight DIMM sockets can be used.

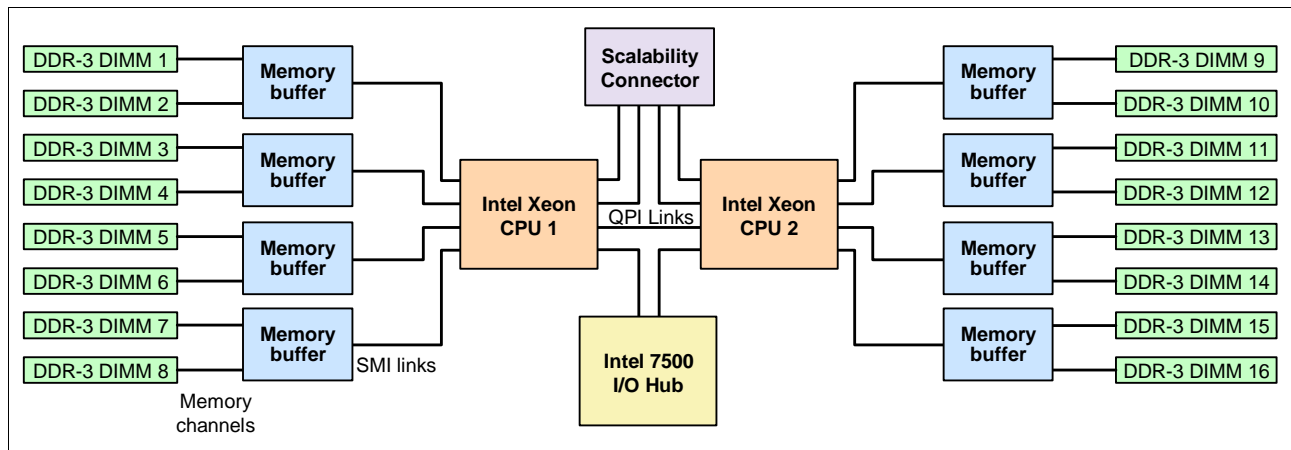


Figure 3-71 A portion of the HX5 block diagram showing the processors, memory buffers, and DIMMs

Figure 3-72 shows the physical locations of the 16 memory DIMM sockets.




Figure 3-72 DIMM layout on the HX5 system board

The MAX5 memory expansion blade has 24 memory DIMM sockets (Figure 3-73). The MAX5, which must be connected to an HX5 system (only the 1-node HX5 supports the MAX5), has one memory controller and six SMI connected memory buffers.




Figure 3-73 DIMM layout on the MAX5 system board

MAX5 memory runs at 1066, 978, or 800 MHz DDR3 speeds. The memory speed is dependent on the processor QPI speed in the HX5:

- ▶ A QPI speed of 6.4 GHz means that the speed of the MAX5 memory is 1066 MHz.
- ▶ A QPI speed of 5.8 GHz means that the speed of the MAX5 memory is 978 MHz.
- ▶ A QPI speed of 4.8 GHz means that the speed of the MAX5 memory is 800 MHz.

Memory features

This section describes the features of the HX5 memory subsystem.

Memory mirroring

Memory mirroring works much like disk mirroring and provides high server availability in the event of a memory DIMM failure. It is roughly equivalent to RAID 1 in disk arrays, in that memory is divided into two halves, and one half mirrors the other half. If 8 GB is installed, the operating system sees 4 GB after memory mirroring is enabled on the system.

Mirroring is handled at the hardware level through UEFI, and no operating system support is required. When memory mirroring is enabled, the memory of a DIMM quadrant (a grouping of four DIMMs, such as DIMMs 1 - 4, 5 - 8, and so on) of a processor is duplicated onto the memory of the other DIMM quadrant of the processor.

Identical memory: Memory must be identical across the two processors to enable memory mirroring.

Chipkill

Chipkill memory technology, an advanced form of error checking and correcting (ECC) from IBM, is available for the HX5 blade. Chipkill protects the memory in the system from any single memory chip failure. It also protects against multi-bit errors from any portion of a single memory chip.

The HX5 does not support redundant bit steering, because the integrated memory controller of the Intel Xeon 7500 processors does not support the feature.

Lock step

HX5 memory operates in lock step mode. Lock step is a memory protection feature that involves the pairing of two memory DIMMs. The paired DIMMs can perform the same operations and the results are compared. If any discrepancies exist between the results, a memory error is signaled. Lock step mode gives a maximum of 64 GB of usable memory with one processor installed, and 128 GB of usable memory with two processors installed (using 8 GB DIMMs).

Memory must be installed in pairs of two identical DIMMs per processors. Although the DIMM pairs installed can be of different sizes, the pairs must be of the same speed.

Memory sparing

With memory sparing, when a DIMM fails, inactive memory is automatically turned on until the failing DIMM is replaced.

Sparing provides a degree of redundancy in the memory subsystem, but not to the extent that mirroring does. In contrast to mirroring, sparing leaves more memory for the operating system. In sparing mode, the trigger for failover is a preset threshold of correctable errors. When this threshold is reached, the content is copied to its spare. The failed DIMM or rank is then taken offline, with the spare counterpart activated for use.

The HX5 supports DIMM sparing, and this feature is enabled in UEFI.

Rank sparing: Rank sparing is not supported on the HX5. MAX5 does not support rank sparing or DIMM sparing.

With DIMM sparing, a DIMM pair within the quadrant can be used as a spare of the other DIMM pair within the quadrant. The DIMM pairs must be identical so that memory sparing can be enabled. The size of the DIMMs that is used for sparing is subtracted from the usable capacity presented to the operating system.

Redundant bit steering: Redundant bit steering (RBS) or Double Device Data Correction (DDDC) is supported by the E7 processors, *but only with x4 memory* and not x8 memory. As shown in Table 3-120 on page 358, all memory currently available for the HX5 or MAX5 for BladeCenter is x8 technology, so RBS or DDDC is currently not supported on the HX5 or MAX5.

Memory population order

Installing DIMMs in the HX5 and MAX5 in the correct order is essential for system performance.

HX5 memory population order

As shown in Figure 3-71 on page 359, the HX5 design has two DIMMs per memory buffer and one DIMM socket per memory channel.

For best performance, install the DIMMs in the sockets listed in Table 3-121. This sequence spreads the DIMMs across as many memory buffers as possible. For example, looking at row three of the table, if you have one processor installed and have six DIMMs, you then install those DIMMs in sockets 1, 2, 3, 4, 5, and 8.

Table 3-121 DIMM installation for the HX5

Number of processors	Number of DIMMs	Processor 1								Processor 2							
		Buffer		Buffer		Buffer		Buffer		Buffer		Buffer		Buffer		Buffer	
		DIMM 1	DIMM 2	DIMM 3	DIMM 4	DIMM 5	DIMM 6	DIMM 7	DIMM 8	DIMM 9	DIMM 10	DIMM 11	DIMM 12	DIMM 13	DIMM 14	DIMM 15	DIMM 16
1	2	x			x												
1	4	x			x	x			x								
1	6	x	x	x	x	x			x								
1	8	x	x	x	x	x	x	x	x								
2	2	x			x												
2	4	x			x					x			x				
2	6	x			x	x			x	x			x				
2	8	x			x	x			x	x			x	x			x
2	10	x	x	x	x	x			x	x			x	x			x
2	12	x	x	x	x	x			x	x	x	x	x	x			x
2	14	x	x	x	x	x	x	x	x	x	x	x	x	x			x
2	16	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

In a two-node (four-socket) configuration with two HX5 servers, follow the memory installation sequence in both nodes. Memory must be populated to have a balance for each processor in the configuration.

Use the following guidelines:

- Install as many DIMMs as possible, preferably in every socket.
- Each processor must have identical amounts of memory.
- Spread out the memory DIMMs to all memory buffers, that is, install one DIMM to a memory buffer before installing a second DIMM to that same buffer. See Table 3-121 for DIMM placement.
- Memory DIMMs must be installed in order of DIMM size, with the largest DIMMs first, then next largest DIMMs, and so on. Placement must follow the DIMM socket installation shown in Table 3-121.
- For maximum memory performance, select a processor with the highest memory bus speed (as listed in Table 3-118 on page 356).

The lower value of the processor's memory bus speed and the DIMM speed determine how fast the memory bus can operate. Every memory bus operates at this speed.

MAX5 V2 memory population order

When using a configuration where an HX5 is connected to a MAX5, follow these rules:

- ▶ At least two DIMMs should be installed in the HX5 (four DIMMs if the HX5 has two processors installed).
- ▶ Fully populate the HX5 using the sequence listed in Table 3-121 on page 362, then populate the MAX using the sequence listed in Table 3-122 on page 364.
- ▶ The data widths for the following quads must match. For example, DIMMs in each quad must be all 4Rx8 or all 2Rx8. See Figure 3-70 on page 355 for the block diagram and Figure 3-73 on page 360 for the physical location of these DIMMs:
 - DIMMs 1, 2, 7, and 8
 - DIMMs 3, 4, 5, and 6
 - DIMMs 13, 14, 17, and 18
 - DIMMs 15, 16, 19, and 20
 - DIMMs 9, 10, 21, and 22
 - DIMMs 11, 12, 23, and 24

Based on the DIMM options that are currently supported in the MAX5 (listed in Table 3-120 on page 358), all DIMMs in each of the quads must be either 4 GB or 8 GB. You cannot mix 4 GB and 8 GB DIMMs in the same quad.

- ▶ Memory must be installed in matched pairs of DIMMs in the MAX5.
- ▶ Memory DIMMs must be installed in order of DIMM size, with the largest DIMMs first. For example, if you are installing both 4 GB and 8 GB DIMMs into the MAX5, use the population order listed in Table 3-122 on page 364 and install all 8 GB DIMMs first, then install the 4 GB DIMMs.

The DIMM sockets in the MAX5 are arranged in three power domains (A, B, C) (Figure 3-74). Each power domain includes two memory controllers and eight DIMM sockets.

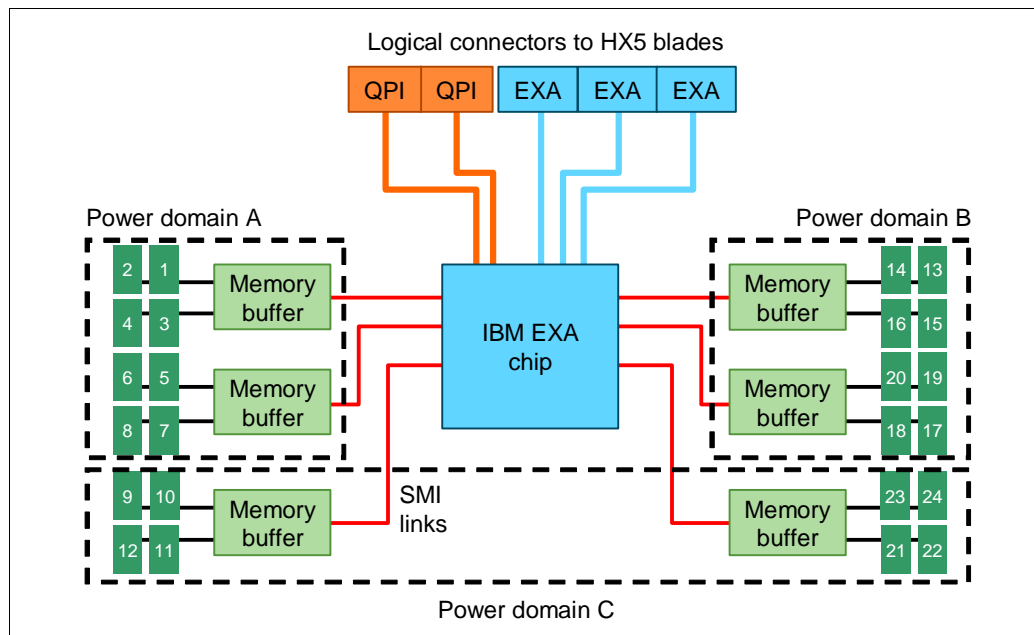


Figure 3-74 Power domains in the MAX5 memory expansion blade

The DIMMs in each power domain are as follows:

- ▶ Power domain A: 1 - 4 and 5 - 8
- ▶ Power domain B: 13 - 16 and 17 - 20
- ▶ Power domain C: 9 - 12 and 21 - 24

DIMMs should be installed spread out among all six memory buffers and all three power domains. Table 3-122 shows the installation order.

Table 3-122 DIMM installation for the MAX5 for IBM BladeCenter

Number of DIMMs	Power domain A								Domain C (½)				Power domain B								Domain C (½)			
	Buffer				Buffer				Buffer				Buffer				Buffer				Buffer			
	DIMM 1	DIMM 2	DIMM 3	DIMM 4	DIMM 5	DIMM 6	DIMM 7	DIMM 8	DIMM 9	DIMM 10	DIMM 11	DIMM 12	DIMM 13	DIMM 14	DIMM 15	DIMM 16	DIMM 17	DIMM 18	DIMM 19	DIMM 20	DIMM 21	DIMM 22	DIMM 23	DIMM 24
2	x							x																
4	x							x					x				x							
6	x							x				x	x				x							x
8	x		x			x		x				x	x				x							x
10	x		x			x		x				x	x		x		x		x					x
12	x		x			x		x		x		x	x		x		x		x			x		x
14	x	x	x			x	x	x		x		x	x		x		x		x			x		x
16	x	x	x			x	x	x		x		x	x	x	x		x	x	x			x		x
18	x	x	x			x	x	x		x	x	x	x	x	x		x	x	x			x	x	x
20	x	x	x	x	x	x	x	x		x	x	x	x	x	x		x	x	x			x	x	x
22	x	x	x	x	x	x	x	x		x	x	x	x	x	x	x	x	x	x	x		x	x	x
24	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

3.7.9 Storage

The storage system on the HX5 blade is based on the usage of the SSD Expansion Card for IBM BladeCenter HX5, which contains an LSI 1064E SAS Controller and two 1.8-inch micro SATA drive connectors. The SSD Expansion Card allows the attachment of two 1.8-inch solid-state drives (SSDs). If two SSDs are installed, the HX5 supports RAID 0 or RAID 1 capability.

Installation of the SSDs in the HX5 requires the SSD Expansion Card for IBM BladeCenter HX5. Only one SSD Expansion Card is needed for either one or two SSDs. Table 3-123 has ordering details.

Table 3-123 SSD Expansion Card for IBM BladeCenter HX5

Part number	Feature code	Description
46M6908	5765	SSD Expansion Card for IBM BladeCenter HX5

Figure 3-75 shows the SSD Expansion Card. The top view is on the left and the underside view on the right.




Figure 3-75 SSD Expansion Card for the HX5 (left: top view, right: underside view)

The SSD Expansion Card can be installed in the HX5 in combination with a CIOv I/O expansion card and a CFFh I/O expansion card (Figure 3-76).




Figure 3-76 Placement of SSD expansion card in combination with a CIOv card and CFFh card

Solid-state drives

Solid-state drives (SSDs) are a relatively new technology in the server world. SSDs are more reliable than spinning hard disk drives and they consume much less power than a standard SAS drive.

The use of any SSD in the HX5 type 7873 requires the installation of the SSD Expansion Card. See Table 3-123 on page 364 for ordering details.

Target applications for SSDs include video surveillance, transaction-based DB, and other applications that have high performance but moderate space requirements.

Table 3-124 lists the supported SSDs.

Table 3-124 Supported solid-state drives

Description	Part number	Feature code
IBM 100GB SATA 1.8" MLC Enterprise SSD	00W1120	A3HQ
IBM 200GB SATA 1.8" MLC Enterprise SSD	49Y6119	A3AN
IBM 400GB SATA 1.8" MLC Enterprise SSD	49Y6124	A3AP

Connecting to external SAS storage devices

The SAS Connectivity Card (CIOv) for IBM BladeCenter is an expansion card that offers the ideal way to connect the supported BladeCenter servers to a wide variety of SAS storage devices. The SAS Connectivity Card connects to two SAS Controller Modules in the BladeCenter chassis. These modules can then be attached to the IBM System Storage DS3500 from the BladeCenter H or HT chassis, or to Disk Storage Modules in the BladeCenter S.

SAS signals are routed from the LSI 1064E controller on the SSD Expansion Card to the SAS Connectivity Card, as shown in Figure 3-77.

Two of the SAS ports (SAS 0, SAS 1) from the LSI 1064E on the SSD Expansion Card are routed to the 1.8-inch SSD connectors. The other SAS ports (SAS 2, SAS 3) are routed from the LSI 1064E controller through the server system board to the CIOv connector where the SAS Connectivity Card (CIOv) is installed.

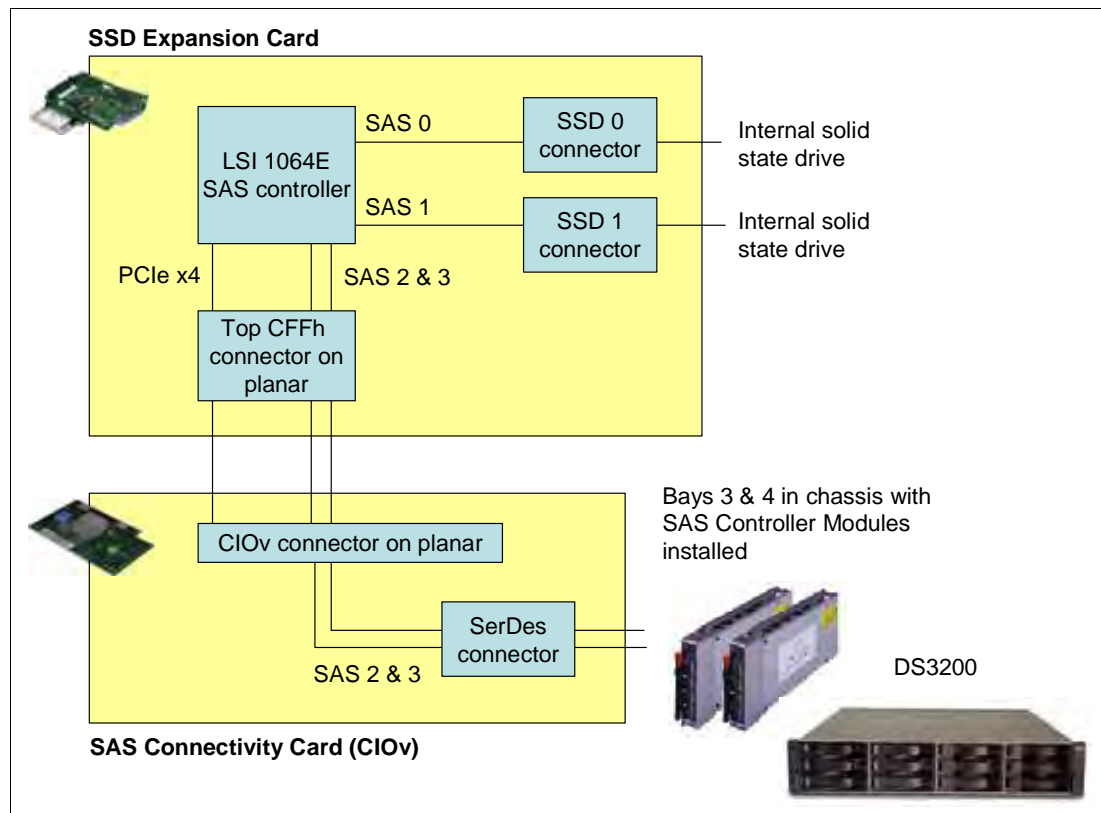


Figure 3-77 Connecting a SAS Connectivity Card and external SAS solution

3.7.10 I/O expansion cards

The HX5 type 7873 is able to connect to a wide variety of networks and fabrics by installing the appropriate I/O expansion card. Supported networks and fabrics include 1 Gb and 10 Gb Ethernet, 4 Gb and 8 Gb Fibre Channel, Serial Attached SCSI (SAS), and InfiniBand.

The HX5 blade server with an I/O expansion card is installed in a supported BladeCenter chassis complete with switch modules (or pass-through) that are compatible with the I/O expansion card in each blade. The HX5 supports two types of I/O expansion cards:

- ▶ The CIOv
- ▶ The CFFh form factor cards

CIOv

The CIOv I/O expansion connector provides I/O connections through the midplane of the chassis to modules in bays 3 and 4 of a supported BladeCenter chassis. The CIOv slot is a second-generation PCI Express 2.0 x8 slot. A maximum of one CIOv I/O expansion card is supported per HX5. A CIOv I/O expansion card can be installed on a blade server at the same time as a CFFh I/O expansion card is installed in the blade.

Table 3-125 lists the CIOv expansion cards that are supported in the HX5.

Table 3-125 Supported CIOv expansion cards

Part number	Feature code	Description
43W4068	1041	SAS Connectivity Card (CIOv)
44W4475	1039	Ethernet Expansion Card (CIOv)
44X1945	1462	QLogic 8 Gb Fibre Channel Expansion Card (CIOv)
46M6065	3594	QLogic 4 Gb Fibre Channel Expansion Card (CIOv)
46M6140	3598	Emulex 8 Gb Fibre Channel Expansion Card (CIOv)

For the latest information about the expansion cards supported by the HX5, see the IBM ServerProven website:

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

CIOv expansion cards are installed in the CIOv slot in the HX5 2-socket (Figure 3-78).



Figure 3-78 The HX5 type 7873 showing the CIOv I/O Expansion card position

CFFh

The CFFh I/O expansion connector provides I/O connections to high-speed switch modules in bays 7, 8, 9, and 10 of a BladeCenter H or BladeCenter HT chassis or to switch bay 2 in a BladeCenter S chassis.

The CFFh slot is a second-generation PCI Express x16 (PCIe 2.0 x16) slot. A maximum of one CFFh I/O expansion card is supported per blade server. A CFFh I/O expansion card can be installed on a blade server at the same time that a CIOv I/O expansion card is installed on the server.

Table 3-126 lists the supported CFFh I/O expansion cards.

Table 3-126 Supported CFFh expansion cards

Part number	Feature code	Description
44W4479	5476	2/4 Port Ethernet Expansion Card (CFFh)
46M6001	0056	2-port 40 Gb InfiniBand Expansion Card (CFFh)
46M6164	0098	Broadcom 10 Gb Gen2 4-port Ethernet Exp Cd (CFFh)
81Y3133	A1QR	Broadcom 2-port 10Gb Virtual Fabric Adapter
81Y1650	5437	Brocade 2-port 10 GbE Converged Network Adapter
00y3264	A3NW	Emulex 10 GbE Virtual Fabric Adapter Advanced II
00Y3266	A3NV	Emulex 10 GbE Virtual Fabric Adapter II
42C1810	3593	Intel 10 Gb 2-port Ethernet Expansion Card (CFFh)
90Y3570	A1NW	Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh)
00Y3280 ^a	A3JB	QLogic 2-port 10 Gb Converged Network Adapter (CFFh)

Part number	Feature code	Description
ooY3332	A4AC	QLogic 10Gb Virtual Fabric Adapter (CFFh)
00Y5618	A4AD	QLogic 10 Gb Virtual Fabric Converged Network Adapter (CFFh)
00Y3270 ^b	A3JC	QLogic Eth and 8 Gb Fibre Channel Exp Card (CFFh)

a. Replaces 42C1830.

b. Replaces 44X1940.

For the latest information about the expansion cards supported by the HX5, see the IBM ServerProven website:

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

CFFh expansion cards are installed in the CFFh slot in the HX5 (Figure 3-79).




Figure 3-79 The HX5 type 7873 showing the CFFh I/O Expansion card position

A CFFh I/O expansion card requires that a supported high-speed I/O module or a Multi-switch Interconnect Module be installed in bay 7, 8, 9, or 10 of the BladeCenter H or BladeCenter HT chassis.

In a BladeCenter S chassis, the CFFh I/O Expansion Card requires a supported switch module in bay 2. When used in a BladeCenter S chassis, a maximum of two ports are routed from the CFFh I/O Expansion Card to the switch module in bay 2.

For the latest information about the switch modules that are supported with each CFFh I/O expansion card, see the *IBM BladeCenter Interoperability Guide* website:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5073016>

PCI Express Gen 2 Expansion Blade (BPE4)

The HX5 type 7873 supports the attachment of up to three PCI Express Gen 2 Expansion Blades on a single HX5 or up to two on a 2-node HX5. The BPE4 allows certain standard PCI-Express adapters to be attached to the HX5. If there is a CFFh I/O Expansion Card already installed on the HX5, it must be removed to allow for the attachment of the PCI Express Gen 2 Expansion Blade. The CFFh I/O Expansion Card can then be installed in the

topmost PCI Express Gen 2 Expansion Blade. No other type of expansion blade can be mixed with the PCI Express Gen 2 Expansion Blade. The PCI Express Gen 2 Expansion Blade cannot be attached to the HX5 if the MAX5 memory expansion is already attached.

For more information about the BPE4, see 3.10.2, “PCI Express Gen 2 Expansion Blade/Expansion Blade II” on page 403.

3.7.11 Standard onboard features

This section describes the standard onboard features of the HX5 blade server:

- ▶ UEFI
- ▶ Onboard network adapters
- ▶ Integrated Management Module
- ▶ Video controller
- ▶ Trusted platform module

UEFI

The HX5 2-socket uses an integrated Unified Extensible Firmware Interface (UEFI) next-generation BIOS. Its capabilities include:

- ▶ Human readable event logs (no more beep codes)
- ▶ Complete setup solution by allowing the adapter configuration function to be moved into the UEFI
- ▶ Complete out-of-band coverage by the Advance Settings Utility to simplify remote setup

Using all of the features that UEFI offers requires a UEFI-aware operating system and adapters. UEFI is fully compatible with earlier versions of the BIOS.

For more information about UEFI, see *Introducing uEFI-Compliant Firmware on IBM System x and BladeCenter Servers*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5083207>

Onboard network adapters

The HX5 2-socket includes a dual-port Gigabit Ethernet controller with the following specifications:

- ▶ Broadcom BCM5709S dual-port Gigabit Ethernet controller
- ▶ Supports TCP offload engine (TOE)
- ▶ Supports failover and load balancing for better throughput and system availability
- ▶ Supports highly secure remote power management using IPMI 2.0
- ▶ Supports Wake on LAN and Preboot Execution Environment (PXE)
- ▶ Supports IPv4 and IPv6

Integrated Management Module

The HX5 blade server includes an Integrated Management Module to monitor server availability, perform predictive failure analysis, and trigger IBM Systems Director alerts. The IMM performs the functions of the Baseboard Management Controller (BMC) of earlier blade servers, and adds the features of the Remote Supervisor Adapter (RSA) found in System x servers, and remote control and remote media. A separate card, like the cKVM card, is no longer required to implement these functions.

For more information about the IMM, see *Transitioning to UEFI and IMM*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5079769>

The IMM controls the service processor LEDs and the light path diagnostic capability. The IMM controls the LEDs, indicating that there is an error and the physical location of the error. To enable illumination of the LEDs after the blade is removed from the chassis, the LEDs have a back-up power system. The LEDs are related to DIMMs, processors, battery, CIOv connector, CFFh connector, scalability, the system board, non-maskable interrupt (NMI), processor mismatch, and the SAS connector.

Video controller

The video subsystem in the HX5 supports an SVGA video display. The video subsystem is a component of the IMM and is based on a Matrox video controller. The HX5 has 16 MB of video memory. Table 3-127 lists the supported video resolutions.

Table 3-127 Supported video resolutions

Resolution	Maximum refresh rate
640 x 480	85 Hz
800 x 600	85 Hz
1024 x 768	75 Hz

3.7.12 Integrated virtualization

The HX5 offers an IBM USB flash drive option preinstalled with VMware ESXi. This is an embedded version of VMware ESX and the hypervisor is fully contained on the flash drive.

Table 3-128 lists the ordering information for the IBM USB Memory Key for VMware Hypervisor.

Table 3-128 USB Hypervisor option

Part number	Feature code	Description
41Y8287	3033	IBM USB Memory Key for VMware ESXi 4.1
41Y8296	A1NP	IBM USB Memory Key for VMware ESXi 4.1 Update 1
41Y8300	A2VC	IBM USB Memory Key for VMware ESXi 5.0
41Y8307	A383	IBM USB Memory Key for VMware ESXi 5.0 Update 1
41Y8311	A2R3	IBM USB Memory Key for VMware ESXi 5.1
41Y8298	A2G0	IBM Blank USB Memory Key for VMware ESXi Downloads

The flash drive plugs into the Hypervisor Interposer, which in turn is attached to the system board near the processors (Figure 3-80). The Hypervisor Interposer is included standard with the HX5.




Figure 3-80 Placement of VMware USB key in HX5

3.7.13 Partitioning capabilities



When you have a four-socket HX5 composed of two HX5 blade servers, you can use partitioning to manage the complex and control through software, whether you want the system to boot as one four-socket system or two two-socket systems, without making any physical changes to the hardware. You can schedule such configuration changes through IBM Systems Director and the Advanced Management Module of the chassis for even greater flexibility.

The configuration of a multinode HX5 is controlled through the Advanced Management Module (AMM) web interface. The AMM communicates with the onboard IMM to enable and disable the scaling capabilities of the system, allowing for a great deal of flexibility.

Figure 3-81 shows the AMM status page with two HX5 blades in slots six and seven with a scalability card attached to the two blades. When these blades are in scaled mode, the power buttons, video buttons, and media tray buttons all synchronize.

Blades ?

Click the icon in the Status column to view detailed information about each blade.

Bay	Status	Name	Pwr	Owner**		cKVM*	I/O Compatibility	WOL*	Local Control			BEM*
				KVM	MT*				Pwr	KVM	MT*	
1		No blade present										
2		No blade present										
3		No blade present										
4		No blade present										
5		No blade present										
6		SN#YK34C09CJ02W	Off				OK	On	✓	✓	✓	---
7		SN#YK34C09CW00N	Off		✓		OK	On	✓	✓	✓	---
8		No blade present										
9		No blade present										
10		No blade present										
11		No blade present										
12		No blade present										
13		No blade present										
14		No blade present										

* MT = Media Tray (CD/ USB) , WOL = Wake on LAN , BEM = Blade Expansion Module
 BSE1 (BSE2,BSE3) = Blade Storage Expansion 1st Generation (2nd Generation, 3rd Generation)
 PEU1 = PCI Expansion Unit 1st Generation PEU2 = PCI Expansion Unit II BPE3 = PCI Express Expansion Unit
 cKVM = Concurrent KVM Expansion BIE = Blade I/O Expansion BPR = Blade Processor Expansion

** You can change the KVM and Media Tray ownership on the Remote Control panel (under Blade Tasks).

Figure 3-81 Advanced Management Module system status

Figure 3-82 shows the two HX5 blades from the AMM web interface in an unpartitioned state. As the Available actions menu indicates, these two blades can be selected and a partition can be created, enabling the scaling between these blades.

Scalable Complex Information

A summary of scalable complex configuration, each tab indicating a unique complex and the slots that are occupied.

Complex (6 - 7)

Help

Assigned Nodes

Click the checkboxes in the first column to select one or more partitions then, click one of the actions in the action list below the table and click "Perform Action" to perform the desired action.

	Partition	Mode	Bay	Name	Status	Processors/Memory	Config Status	Primary
No partitions present								

Available actions

Power Off Partition

Perform action

Unassigned Nodes

Click the checkboxes in the first column to select one or more nodes then, click one of the actions in the action list below the table and click "Perform Action" to perform the desired action.

	Bay	Name	Status	Processors/Memory
<input checked="" type="checkbox"/>	6	SN#YK34C09CJ02W	Powered Off	2 Intel Xeon / 8 DIMMS 2GB
<input checked="" type="checkbox"/>	7	SN#YK34C09CW00N	Powered Off	8 DIMMS 2GB

Available actions

Create Partition

Create Partition

Power On Node

Power Off Node

Perform action

Figure 3-82 Two unpartitioned HX5s

Figure 3-83 on page 375 shows the result of creating a partition using the two HX5 blades. The mode is shown as "Partition" and the two blades are scaled together as a dual node, 4-socket system. The menu shows the available options that can be performed on the current partition. The following options are available:

- ▶ **Power Off Partition**
This option powers off all blades in the partition (in this example, it powers off blade 6 and blade 7).
- ▶ **Power On Partition**
This option powers on all blades in the partition.
- ▶ **Remove Partition**
This option removes the partitioning of the selected partitions. The nodes in those partitions then are listed in the Unassigned Nodes section.
- ▶ **Toggle Stand alone / Partition Mode**
This option keeps the selected partition in place, but moves the nodes into either a set of stand-alone blades or into a single partitioned system.

The mode in Figure 3-83 shows the current mode, which can be either Partition or Stand-alone state.

Scalable Complex Information

A summary of scalable complex configuration, each tab indicating a unique complex and the slots that are occupied.

Complex (6 - 7) Help

Assigned Nodes

Click the checkboxes in the first column to select one or more partitions then, click one of the actions in the action list below the table and click "Perform Action" to perform the desired action.

<input checked="" type="checkbox"/>	Partition	Mode	Bay	Name	Status	Processors/Memory	Config Status	Primary
<input checked="" type="checkbox"/>	1	Partition	6	SN#YK34C09CJ02W	Powered Off	2 Intel Xeon / 8 DIMMS 2GB	Valid	✓
			7	SN#YK34C09CW00N	Powered Off	8 DIMMS 2GB	Valid	

Available actions

- Power Off Partition
- Power On Partition
- Power Cycle Partition
- Remove Partition
- Toggle Stand-alone/Partition Mode

to select one or more nodes then, click one of the actions in the action list below the table and click "Perform Action" to perform the desired action.

<input type="checkbox"/>	Bay	Name	Status	Processors/Memory
No unassigned nodes present				

Available actions

Tues, 02 Mar 2010 14:32:39

Figure 3-83 Two HX5 blades in Partitioned mode

Figure 3-84 shows the two HX5 blades configured in a Stand-alone mode. In this mode, the blades start separately and can run two separate operating systems.

Scalable Complex Information

A summary of scalable complex configuration, each tab indicating a unique complex and the slots that are occupied.

Complex (6 - 7) Help

Assigned Nodes

Click the checkboxes in the first column to select one or more partitions then, click one of the actions in the action list below the table and click "Perform Action" to perform the desired action.

<input type="checkbox"/>	Partition	Mode	Bay	Name	Status	Processors/Memory	Config Status	Primary
<input type="checkbox"/>	1	Stand-alone	6	SN#YK34C09CJ02W	Powered Off	2 Intel Xeon / 8 DIMMS 2GB	Valid	✓
<input type="checkbox"/>			7	SN#YK34C09CW00N	Powered Off	2 Intel Xeon / 8 DIMMS 2GB	Valid	

Available actions

Power Off Partition Perform action

Unassigned Nodes

Click the checkboxes in the first column to select one or more nodes then, click one of the actions in the action list below the table and click "Perform Action" to perform the desired action.

<input type="checkbox"/>	Bay	Name	Status	Processors/Memory
No unassigned nodes present				

Available actions

Create Partition Perform action

Tues, 02 Mar 2010 14:35:45

Figure 3-84 Two HX5 blades in Stand-alone mode

The flexible partitioning of the HX5 systems through the AMM can also be scripted by using the AMM command-line interface. This approach allows scheduled partitioning or a mass configuration partition.

Several use case examples for flexible partitioning are as follows:

- ▶ A single hardware configuration is designed at the rack level. It includes scalability cards for all of the HX5 blades, but the solution usage of the rack design might or might not require scaling. A single hardware configuration can be designed and reused multiple times for many separate workloads, with the scaling turned on or off, on one or more of the HX5 blades, depending on the needs and requirements of the solution.
- ▶ The repurposing of hardware. For example, suppose that a given software application is currently unable take advantage of anything larger than a single node system. Now, a future version of this software is developed to take advantage of the dual node capabilities. The initial configuration of the HX5 blades is set up in Stand-alone mode, but when the software is upgraded, the HX5 blades are repurposed to run in a dual-node configuration.

3.8 IBM BladeCenter PS700, PS701, and PS702

The BladeCenter PS700, PS701, and PS702 blades are IBM Power Systems™ blades for 64-bit applications. They are designed to minimize complexity, improve efficiency, automate processes, reduce energy consumption, and scale easily. The POWER7 processor-based blades support AIX, IBM i, and Linux operating systems. Their ability to coexist in the same chassis with other IBM BladeCenter blade servers enhances the ability to deliver the rapid return of investment (ROI) that is demanded by clients and businesses.

Figure 3-85 shows the IBM BladeCenter PS700, PS701, and PS702 blade servers.



Figure 3-85 The IBM BladeCenter PS702, BladeCenter PS701, and BladeCenter PS700

The following topics are covered in this section:

- ▶ Overview
- ▶ Supported BladeCenter chassis
- ▶ Minimum features
- ▶ Processor
- ▶ Memory
- ▶ I/O
- ▶ Storage
- ▶ Standard onboard features

3.8.1 Overview

This section introduces the three blade servers.

The PS700 blade server

The PS700 blade server (8406-70Y) is a single socket, single wide 4-core 3.0 GHz POWER7 processor-based server. The POWER7 processor is a 64-bit, 4-core with 256 KB L2 cache per core and 4 MB L3 cache per core.

The PS700 blade server has eight DDR3 memory DIMM slots. The industry standard VLP DDR3 Memory DIMMs are either 4 GB or 8 GB running at 1066 MHz. The memory is supported in pairs, so the minimum memory required for the PS700 blade server is 8 GB (two 4 GB DIMMs). The maximum memory that can be supported is 64 GB (eight 8 GB DIMMs).

The blade server has two Host Ethernet Adapters (HEA) 1 GB integrated Ethernet ports that are connected to the BladeCenter chassis fabric (midplane). The PS700 has an integrated SAS controller that supports local (onboard) storage, integrated USB controller and Serial over LAN console access through the service processor, and the BladeCenter Advance Management Module.

It supports two onboard disk drive bays. The onboard storage can be one or two 2.5-inch SAS HDD or SSD drives. The integrated SAS controller supports RAID 0, RAID 1, and RAID10 hardware when two HDDs or solid-state drives (SSDs) are used.

The PS700 also supports one PCIe CIOv expansion card slot and one PCIe CFFh expansion card slot.

Figure 3-86 shows the BladeCenter PS700, model 8406-70Y.




Figure 3-86 Top view of PS700 blade server

The features of the server are as follows:

- ▶ Machine type and model number: 8406-70Y
- ▶ Form factor: Single-wide (30 mm) blade
- ▶ Processors:
 - Single-socket 4-core 64-bit POWER7 processor operating at a 3.0 GHz clock speed
 - Based on CMOS 12S 45 nm SOI (silicon-on-insulator) technology
 - Power consumption: 150 W/socket
 - Single-wide (SW) Blade package
- ▶ Memory:
 - Eight DIMM slots

- Minimum capacity 8 GB, maximum capacity 64 GB
- ▶ Industry-standard VLP DDR3 DIMMs
- ▶ Disk:
 - Two disk drive bays supporting one or two SAS HDD
 - Hardware mirroring RAID 0, RAID 1, and RAID 10
- ▶ Onboard integrated features:
 - Service processor (SP)
 - Two 1 GB Ethernet ports (HEA)
 - SAS Controller
 - USB Controller that routes to the USB 2.0 port on the media tray
 - One Serial over LAN (SOL) Console through SP
- ▶ Expansion card I/O options:
 - One CIOv expansion card slot (PCIe)
 - One CFFh expansion card slot (PCIe)

The PS701 blade server

The PS701 blade server (8406-71Y) is a single socket, single-wide 8-core 3.0 GHz POWER7 processor-based server. The POWER7 processor is a 64-bit, 8-core with 256 KB L2 cache per core and 4 MB L3 cache per core.

The PS701 blade server has 16 DDR3 memory DIMM slots. The industry-standard VLP DDR3 memory DIMMs are either 4 GB or 8 GB running at 1066 MHz. The memory is supported in pairs, thus the minimum memory required for PS701 blade server is 8 GB (two 4 GB DIMMs). The maximum memory that can be supported is 128 GB (16x 8 GB DIMMs).

The PS701 blade server has two Host Ethernet Adapters (HEAs) 1 GB integrated Ethernet ports that are connected to the BladeCenter chassis fabric (midplane). The PS701 also has an integrated SAS controller that supports local (onboard) storage, integrated USB controller and Serial over LAN console access through the service processor, and the BladeCenter Advance Management Module.

The PS701 has one onboard disk drive bay. The onboard storage can be one 2.5-inch SAS HDD or SSD drive. The PS701 also supports one PCIe CIOv expansion card slot and one PCIe CFFh expansion card slot.

Figure 3-87 shows the BladeCenter PS701.

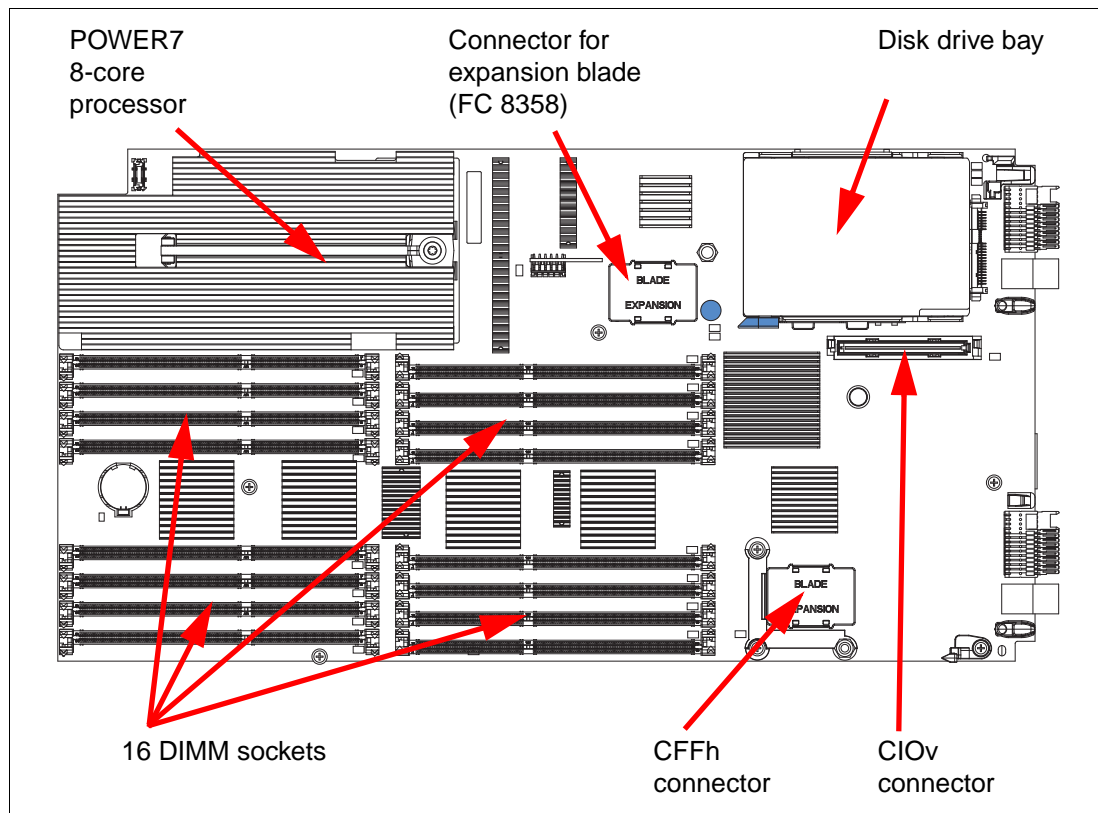


Figure 3-87 Top view of the PS701 blade server

The features of the server are as follows:

- ▶ Machine type and model number: 8406-71Y
- ▶ Form factor: Single-wide (30 mm) blade
- ▶ Processors:
 - Single-socket 8-core 64-bit POWER7 processor operating at a 3.0 GHz clock speed
 - Based on CMOS 12S 45 nm SOI (silicon-on-insulator) technology
 - Power consumption: 150 W/socket
 - Single-wide (SW) Blade package
- ▶ Memory:
 - Sixteen DIMM slots
 - Minimum capacity 8 GB, maximum capacity 128 GB
 - Industry standard VLP DDR3 DIMMs
- ▶ Disk
 - One disk drive bay supports one SAS HDD (hard disk drive). Support for SSD is planned.

- ▶ Onboard integrated features:
 - Service processor (SP)
 - Two 1 GB Ethernet ports (HEA)
 - SAS Controller
 - USB Controller that routes to the USB 2.0 port on the media tray
 - One Serial over LAN (SOL) Console through SP
- ▶ Expansion card I/O options:
 - One CIOv expansion card slot (PCIe)
 - One CFFh expansion card slot (PCIe)

The PS702 blade server

The PS702 blade server (8406-71Y +FC 8358) is a two socket, double-wide 16-core 3.0 GHz POWER7 processor-based server. The POWER7 processor is a 64-bit, 8-core with 256 KB L2 cache per core and 4 MB L3 cache per core.

The PS702 combines a single-wide base blade (PS701) and an expansion unit (feature code 8358), referred to as double-wide blade, which occupies two adjacent slots in the BladeCenter chassis.

The PS702 blade server has 32 DDR3 memory DIMM slots. The industry-standard VLP DDR3 memory DIMMs are either 4 GB or 8 GB running at 1066 MHz. The memory is supported in pairs. The minimum memory required for the PS702 blade server is 8 GB (two 4 GB DIMMs). The maximum memory that can be supported is 256 GB (32x 8 GB DIMMs).

The PS702 blade server has four Host Ethernet Adapter 1 GB integrated Ethernet ports that are connected to the BladeCenter chassis fabric (midplane). The PS702 also has an integrated SAS controller that supports local (onboard) storage, integrated USB controller and Serial over LAN console access through the service processor, and the BladeCenter Advance Management Module.

The PS702 blade server has two disk drive bays, one on the base blade and one on the expansion unit. The onboard storage can be one or two 2.5-inch SAS HDD or SSD drives. The integrated SAS controller supports RAID 0, RAID 1, and RAID 10 hardware when two HDDs or SSDs are used.

The PS702 supports two PCIe CIOv expansion card slots and two PCIe CFFh expansion card slots.

Figure 3-88 shows the two halves of the BladeCenter PS702.

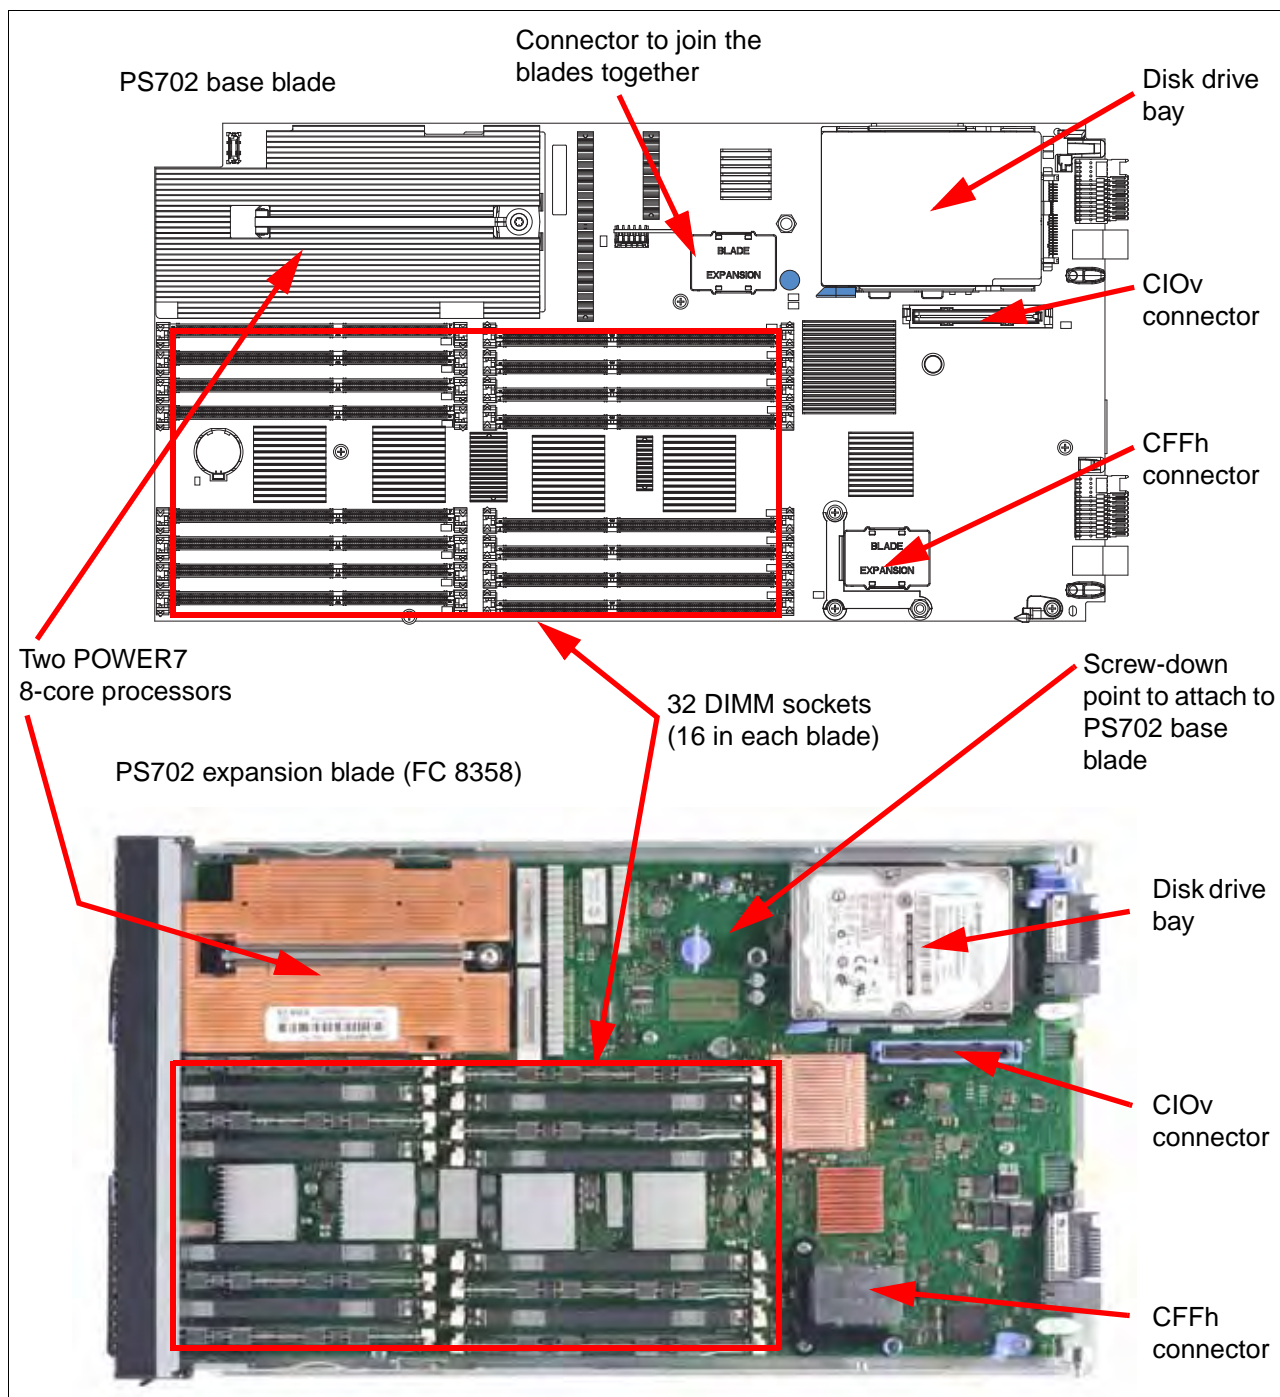


Figure 3-88 Top view of PS702 blade server

The features of the server are as follows:

- ▶ Machine type and model number: 8406-71Y + FC 8358
- ▶ Form factor: Double-wide (60 mm) blade
- ▶ Processors:
 - Two-socket 16-core 64-bit POWER7 processor operating at a 3.0 GHz clock speed
 - Based on CMOS 12S 45 nm SOI (silicon-on-insulator) technology
 - Power consumption: 150 W/socket
 - Double-wide (SW) Blade package
- ▶ Memory:
 - Sixteen DIMM Slots
 - Minimum capacity 8 GB, maximum capacity 256 GB
 - Industry Standard VLP DDR3 DIMMs
- ▶ Disk:
 - Two disk drive bays (one on each blade) supports one or two SAS HDD
 - Support for SSD planned
 - Hardware mirroring RAID 0, RAID 1, RAID 10
- ▶ Onboard integrated features:
 - Service processor (one on each blade¹)
 - Four 1 GB Ethernet ports (HEA)
 - SAS Controller
 - USB Controller which routes to the USB 2.0 port on the media tray
 - One Serial over LAN (SOL) Console through FSP
- ▶ Expansion card I/O options:
 - One CIOv expansion card slot (PCIe)
 - One CFFh expansion card slot (PCIe)

Important: The PS702 is 16-core POWER7 processor-based blade server that is a combination of a single-socket 8-core blade, model 8406-71Y (PS701) and a PS702 expansion blade.

3.8.2 Supported BladeCenter chassis

The PS700, PS701, and PS702 blades are supported in the IBM BladeCenter chassis, as listed in Table 3-129.

Table 3-129 The blade servers supported in each BladeCenter chassis

Blade	Machine type-model	Blade width	BC S 8886	BC E 8677 ^a	BC T 8720	BC T 8730	BC H 8852	BC HT 8740	BC HT 8750
PS700	8406-70Y	1 slot	Yes	Yes ^b	No	No	Yes	Yes	Yes
PS701	8406-71Y	1 slot	Yes	No	No	No	Yes	Yes	Yes
PS702	8406-71Y	2 slot	Yes	No	No	No	Yes	Yes	Yes

a. BladeCenter E requires an Advanced Management Module and a minimum of two 2000 W power supplies.

b. Only specific models of the BladeCenter E support the PS700. See Table 3-130.

¹ The service processor (or flexible service processor) on the expansion unit provides control but does not offer redundancy with the SP on the base unit.

A detailed description of each supported BladeCenter for the PS700, PS701, and PS702 blades is contained in the following sections:

- ▶ IBM BladeCenter E (PS700 only support) provides the greatest density and common fabric support. It is the lowest entry cost option. For details about the chassis, see 2.1, “IBM BladeCenter E” on page 38.

Only specific models of the BladeCenter E models listed are supported, as listed in Table 3-130.

Table 3-130 BladeCenter E models that support the PS700

BladeCenter E models ^a	Supports the PS700
8677-3Xx	No
8677-3Rx	No
8677-E2x	No
8677-3Sx	Yes ^b
8677-4Sx	Yes
8677-3Tx	Yes ^b
8677-4Tx	Yes

a. The x is the country-specific letter (for example, EMEA MTM is 8677-4SG, and the US MTM is 8677-4SU).

b. The 3Sx and 3Tx models are supported but only with upgraded (2320W) power supplies.

- ▶ IBM BladeCenter H delivers high performance, extreme reliability, and ultimate flexibility for the most demanding IT environments. See 2.2, “IBM BladeCenter H” on page 46.
- ▶ IBM BladeCenter HT models are designed for high-performance flexible telecommunications environments by supporting high-speed Internet working technologies (such as 10Gb Ethernet). They provide a robust platform for NGNs. See 2.3, “IBM BladeCenter HT” on page 56.
- ▶ 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 environments. See 2.4, “IBM BladeCenter S” on page 67.

The number of POWER7 processor-based blades that can be installed in a BladeCenter chassis depends on several factors:

- ▶ BladeCenter chassis type
- ▶ Number of power supplies installed
- ▶ Power supply voltage option (BladeCenter S only)
- ▶ BladeCenter power domain configuration

Table 3-131 shows the maximum number of PS700, PS701, and PS702 blades running in a maximum configuration (memory, disk, and expansion cards) for each supported BladeCenter chassis that can be installed with fully redundant power and without performance reduction. IBM blades that are based on processor types other than POWER7 might reduce these numbers.

Table 3-131 PS700, PS701, and PS702 blades per chassis type

	BladeCenter E ^a		BladeCenter H		BladeCenter HT		BladeCenter S			
	14 slots total		14 slots total		12 slots total		6 slots total			
							110 VAC		208 VAC	
	2 PS	4 PS	2 PS	4 PS	2 PS	4 PS	2 PS	4 PS	2 PS	4 PS
PS700	6	14	7	14	6	12	2	6	2	6
PS701	None	None	7	14	6	12	2	6	2	6
PS702	None	None	3	7	3	6	1	3	1	3

a. BladeCenter E requires 2000 or 2300 W power supplies.

When mixing blades of different processor types in the same BladeCenter, the BladeCenter Power Configurator tool determines whether the combination is valid. For more information about this update, see the following website:

<http://www.ibm.com/systems/bladecenter/powerconfig>

3.8.3 Minimum features

At a minimum, the PS700, PS701, and PS702 have the following requirements:

- ▶ A BladeCenter chassis and one processor socket per blade (a four-core single socket in PS700, an eight-core single socket in PS701, and two eight-core single sockets in the PS702 blade servers)
- ▶ Minimum memory (8 GB) and zero or one DASD
- ▶ A Language Group Specify (mandatory for the order voltage nomenclature/language)

Each system has a minimum feature set that must be met. Table 3-132 lists the minimum system configuration for the PS700, PS701, and PS702 blade servers.

Table 3-132 Minimum features for PS700, PS701, and PS702 blade server

Category	Minimum features required
BladeCenter chassis	Supported BladeCenter chassis (See 3.8.2, "Supported BladeCenter chassis" on page 383.)
Processor	<ul style="list-style-type: none"> ▶ 4-core 3.0 GHz PS700 Blade (#8406-70Y) ▶ 8-core 3.0 GHz PS701 Blade (#8406-71Y) ▶ 16-core 3.0 GHz PS702 Blade(#8406-71Y + FC 8358)
Memory	Two DDR3 Memory DIMM: <ul style="list-style-type: none"> ▶ For P700 and P701 8 GB (2x 4 GB) Memory DIMMs, 1066 MHz (#8208) ▶ For P702 8 GB (2x 4 GB) Memory DIMMs, 1066 MHz (#8208) on each base board

Category	Minimum features required
Storage For AIX and Linux: 1x disk drive For IBM i: 2x disk drives	AIX/Linux/Virtual I/O Server; either 300 <i>or</i> 600 GB: <ul style="list-style-type: none"> ▶ 300 GB SAS 2.5-inch HDD (#8274) ▶ 600 GB SAS 2.5 inch HDD (#8276) IBM i (Required VIOS partition); either 300 <i>or</i> 600 GB: <ul style="list-style-type: none"> ▶ 300 GB SAS 2.5-inch HDD (#8274) ▶ 600 GB SAS 2.5 inch HDD (#8276) If Boot from SAN 8 GB Fibre Channel HBA is selected with FC #8240, #8242, or #8271, or Fibre Channel over Ethernet Adapter, FC #8275 must be ordered.
1x Language Group	Country specific (selected by the customer)
Operating system	1x primary operating system; one of the following systems: <ul style="list-style-type: none"> ▶ AIX (#2146) ▶ Linux (#2147) ▶ IBM i (#2145) plus IBM i 6.1.1 (#0566)

3.8.4 Processor

The POWER7 3.0 GHz 64-bit POWER7 processor for blade servers is available in four-core (PS700), eight-core (PS701), or two eight-core (PS702) configurations. They are optimized to achieve maximum performance for both the system and its virtual machines. Couple that performance with PowerVM and you are now empowered for massive workload consolidation to drive maximum system use, predictable performance, and cost efficiency.

POWER7 Intelligent Threads Technology enables workload optimization by selecting the most suitable threading mode (Single thread (per core) or Simultaneous Multi-thread 2 or 4 modes, also called 2-SMT and 4-SMT). The Intelligent Threads Technology can provide improved application performance. In addition, POWER7 processors can maximize cache access to cores, improving performance, using Intelligent Cache technology.

POWER7 offers Intelligent Energy Management features that can dramatically and dynamically conserve power and further improve energy efficiency. These features enable the POWER7 processor to operate at a higher frequency if environmental conditions permit, for increased performance and performance per watt. Alternatively, if user settings permit, these features allow the processor to operate at a reduced frequency for significant energy savings.

The following list indicates the key processor feature for each P700 blade server:

- ▶ The PS700 blade server contains one four-core, 64-bit POWER7 3.0 GHz processor with 256 KB per processor core L2 cache and 4 MB per processor core L3 cache. No processor options are available.
- ▶ The PS701 blade server contains one eight-core, 64-bit POWER7 3.0 GHz processor with 256 KB per processor core L2 cache and 4 MB per processor core L3 cache. No processor options are available.
- ▶ The PS702 blade server is a double-wide that supports two eight-core, 64-bit POWER7 3.0 GHz processor with 256 KB per processor core L2 cache and 4 MB per processor core L3 cache. No processor options are available.

3.8.5 Memory

The PS700, PS701, and PS702 blade servers use industry-standard VLP DDR3 memory DIMMs. Memory DIMMs must be installed in matched pairs with the same size and speed.

The PS700, PS701, and PS702 blade servers have eight (PS700), 16 (PS701), and 32 (PS702) DIMM slots. Memory is available in 4 GB or 8 GB DIMMs, both operating at a memory speed of 1066 MHz. The memory sizes can be mixed within a system. You can use pairs of 4 GB DIMMs with pairs of 8 GB DIMMs.

The POWER7 DDR3 memory uses a new memory architecture to provide greater bandwidth and capacity. This configuration enables operating at a higher data rate for larger memory configurations. Table 3-133 shows the DIMM features.

Table 3-133 Memory DIMM options

Feature code	DIMM size	Quantity	Speed
8208	4 GB	2	1066 MHz
8209	8 GB	2	1066 MHz

DIMM support: The DDR2 DIMMs used in JS23 and JS43 blade servers are not supported in the POWER7 blade servers.

The announcement letter for the POWER7 processor-based blades incorrectly lists the memory speed of the 8 GB DIMMs to be 800 MHz.

3.8.6 I/O

The PS700 and PS701 have one CIOv PCIe expansion card slot and one CFFh PCIe high-speed expansion card slot. The PS702 blade server has two CIOv expansion card slots and two CFFh expansion card slots.

Table 3-134 shows the supported CIOv and CFFh expansion cards in POWER7 processor-based servers.

Table 3-134 Supported I/O expansion cards on POWER7 blades

Card description	Feature code
CIOv	
Broadcom 2-port Gigabit Ethernet Expansion Card (CIOv)	8243
QLogic 8 Gb Fibre Channel Expansion Card (CIOv)	8242
QLogic 4 Gb Fibre Channel Expansion Card (CIOv)	8241
Emulex 8 Gb Fibre Channel Expansion Card (CIOv)	8240
3 GB SAS Pass-thru Expansion Card (CIOv)	8246
CFFh	
QLogic 8 GB Fibre Channel Expansion Card (CFFh)	8271
QLogic Ethernet and 4 Gb Fibre Channel Expansion Card (CFFh)	8252
QLogic 2-port 10 Gb Converged Network Adapter (CFFh)	8275
4X InfiniBand DDR Expansion Card (CFFh)	8258
Broadcom 2/4-port Ethernet Expansion Card (CFFh)	8291

3.8.7 Storage

The PS700 blade server has two disk bays:

- ▶ In the first bay, it can have one 2.5-inch SAS HDD.
- ▶ In second bay, it can have one 2.5-inch SAS HDD.

The PS701 blade server has one disk bay. In this bay, it can have one 2.5-inch SAS HDD.

The PS702 blade server has two disk bays (one on each blade):

- ▶ On the base card, it can have one 2.5-inch SAS HDD.
- ▶ On the expansion unit, it can have one 2.5-inch SAS HDD.

Table 3-135 lists supported disk features on PS700, PS701, and PS701 blade servers.

Table 3-135 Supported disk drives

Feature code	Description
8274	300 GB 10 K SFF SAS HDD
8276	600 GB 10 K SFF SAS HDD

3.8.8 Standard onboard features

There are many standard onboard features.

Service processor

The service processor (or flexible service processor (FSP)) is the main part of the blade server. It monitors and manages system hardware, resources, and devices. It does the system initialization, configuration, and thermal/power management. It takes corrective action if required.

The PS700 and PS701 have only one service processor. The PS702 blade server has two FSPs (one on each blade). However, the second service processor is only in IO mode and is not redundant to the one on the base blade.

Host Ethernet Adapter

The integrated I/O Hub provides two 1 GB Ethernet ports called Host Ethernet Adapters (HEAs). HEAs are part of the Integrated Virtual Ethernet subsystem (IVE). Each HEA has its own MAC address and can have a maximum of 16 logical ports. These logical ports can be used to communicate to the multiple LPARs, which helps in virtualization/sharing of the Ethernet port without using the Ethernet bridge on the Virtual IO Server.

The PS700 and PS701 blade servers have two 1 GB HEAs. The PS702 has four 1 GB HEAs (two on each system board).

SAS controller

The integrated SAS controller is used to drive the local SAS storage. This SAS controller can also support the onboard SATA SSD by using the Vega SAS to SATA conversion module (also called the SAS to SATA interposer assembly).

The 3 GB SAS Pass-thru Expansion Card can be used to connect to the BladeCenter SAS switch, which can be connected to the external storage. This SAS Pass-thru Expansion Card can also be used to connect to BladeCenter S internal drive SAS drives.

The PS700 and PS701 blade servers have one integrated SAS controller. The PS702 has two integrated SAS controllers, one on each blade. The SAS controller host PCI-X interface to P5IOC2 I/O Hub is 64 bits wide and operates at 133 MHz.

The integrated SAS controller supports hardware mirroring RAID 0, RAID 1, or RAID 10 when two HDD or SSD disks are used in P701 or P702 blade servers.

USB controller

The USB controller connects the USB bus to the midplane, which is then routed to the media tray in the BladeCenter chassis to connect to USB devices (such as an optical drive or diskette drive).

Serial over LAN (SOL)

The integrated SOL function routes the console data stream over standard dual 1 GB Ethernet ports to the AMM. The PS700, PS701, and PS702 do not have onboard video chips and do not support KVM connections. Console access is only by SOL connection. Each blade can have a single SOL session. However, there can be multiple Telnet or SSH sessions to the BladeCenter AMM, each acting as a SOL connection to a different blade.

For more information about the BladeCenter POWER7 products, see *IBM BladeCenter PS700, PS701, and PS702 Technical Overview and Introduction*, REDP-4655.

3.9 IBM BladeCenter PS703 and PS704

Built on the proven foundation of the IBM BladeCenter family of products, the BladeCenter PS703 and PS704 are the premier blade servers for 64-bit applications. Using these blade servers, you can minimize complexity, improve efficiency, automate processes, reduce energy consumption, and scale easily. The new POWER7 processor-based IBM Power Systems blades automatically optimize performance and capacity at either a system or virtual machine level and benefit from the new POWER7 processor, which contains innovative technologies to help maximize performance and optimize energy efficiency. It represents one of the most flexible and cost-efficient solutions for UNIX, IBM i, and Linux deployments available in the market. Further enhanced by its ability to be installed in the same chassis with other IBM BladeCenter blade servers, the PS703 and PS704 blades can deliver the rapid return on investment that clients and businesses demand.

Figure 3-89 shows the IBM BladeCenter PS704 and PS703 blade servers.



Figure 3-89 The IBM BladeCenter PS704 (left) and BladeCenter PS703 (right)

The following topics are covered in this section:

- ▶ Overview
- ▶ Supported BladeCenter chassis
- ▶ Minimum features
- ▶ Processor
- ▶ Memory
- ▶ I/O
- ▶ Storage
- ▶ Standard onboard features

3.9.1 Overview

This section introduces the two blade servers.

The PS703 blade server

The IBM BladeCenter PS703 (7891-73X) blade server is a single-wide blade with two 8-core, 64-bit 2.4 GHz processors, and 16 DIMM slots that can contain up to a maximum of 128 GB of memory. One optional SAS Small Form Factor (SFF) 2.5-inch drive, or up to two optional 1.8-inch SATA solid-state drives (SSDs), can be installed in the PS703.

The PS703 has one Broadcom BCM5709S providing two 1 GB integrated Ethernet ports that are connected to the BladeCenter chassis fabric (midplane). The PS703 has an integrated SAS controller that can support a single 2.5-inch non-hot swap SAS hard disk drive or two 1.8-inch SATA solid-state devices through an optional SAS to SATA interposer card. The integrated SAS controller supports RAID 0, RAID 1, and RAID 10 when two HDDs or solid-state drives (SSDs) are used. The integrated SAS controller can also connect, through an optional CIOv 3 Gb SAS Pass-thru Expansion Card, to external storage attached to the SAS switch modules in the BladeCenter chassis.

The PS703 has one PCIe CIOv expansion card slot and one PCIe CFFh expansion card slot.

Figure 3-90 shows the BladeCenter PS703 model 7891-73X.




Figure 3-90 Top view of the PS703 blade server

The features of the PS703 blade server are as follows:

- ▶ Machine type and model number: 7891-73X
- ▶ Form factor: Single-wide (30 mm) blade
- ▶ Processors:
 - Two 8-core 64-bit POWER7 processor operating at a 2.4 GHz clock speed
 - Based on CMOS 12S 45 nm SOI (silicon-on-insulator) technology
 - Power consumption: 110 W per socket
 - Single-wide (SW) Blade package
- ▶ Memory:
 - Sixteen DIMM slots.
 - The minimum memory is 8 GB. The maximum capacity 128 GB.
 - Industry-standard VLP DDR3 DIMMs.
 - Optional Active Memory Expansion.

- ▶ Disk:
 - One disk drive bay supports one 2.5-inch SAS HDD or two 1.8-inch SAS SSDs.
 - A 3 Gb SAS disk storage controller.
 - Hardware mirroring RAID 0, RAID 1, and RAID 10.
- ▶ Onboard integrated features:
 - Service processor (SP)
 - Two 1 GB Ethernet ports
 - SAS Controller
 - USB Controller that routes to the USB 2.0 port on the media tray
 - One Serial over LAN (SOL) Console through SP
- ▶ Expansion card I/O options:
 - One CIOv expansion card slot (PCIe)
 - One CFFh expansion card slot (PCIe)

The PS704 blade server

The IBM BladeCenter PS704 is a double-wide blade with four 8-core, 64-bit 2.4 GHz processors, and 32 DIMM slots that can contain up to a maximum of 256 GB of memory. One or two optional SAS Small Form Factor (SFF) 2.5-inch drives or one, two, three, or four optional 1.8-inch SATA SSDs, can be installed in the PS704.

The PS704 combines a single-wide base blade and an SMP unit, to operate as a double-wide blade, which occupies two adjacent slots in the IBM BladeCenter chassis.

The PS704 blade server has 32 DDR3 memory DIMM slots. The industry standard VLP DDR3 memory DIMMs are either 4 GB or 8 GB running at 1066 MHz. The minimum memory required for PS704 blade server is 32 GB. The maximum memory that can be supported is 256 GB (32x 8 GB DIMMs).

The PS704 blade server has two disk drive bays, one on the base blade and one on the SMP unit. The onboard storage can be one or two 2.5-inch SAS HDDs or up to four SSD drives. The integrated SAS controller supports RAID 0, RAID 1, and RAID 10 hardware when two HDDs or SSDs are used.

The PS704 blade server has four onboard 1 Gb integrated Ethernet ports that are connected to the BladeCenter chassis fabric (midplane). The PS704 also has an integrated USB controller and Serial over LAN console access through the service processor, and the BladeCenter Advance Management Module.

The PS704 supports two PCIe CIOv expansion card slots and two PCIe CFFh expansion card slots.

Figure 3-91 shows the two halves of the BladeCenter PS704, model 7891-74X.




Figure 3-91 Top view of PS704 blade server (base and SMP unit)

The features of the server are as follows:

- ▶ Machine type and model number: 7891-74X
- ▶ Form factor: Double-wide (60 mm) blade
- ▶ Processors:
 - Four eight-core 64-bit POWER7 processors, operating at a 2.4 GHz clock speed for a total of 32 cores in the blade server
 - Based on CMOS 12S 45 nm silicon-on-insulator (SOI) technology
 - Power consumption is 110 W per socket
- ▶ Memory:
 - 32 DIMM slots.
 - Minimum memory is 32 GB. Maximum capacity is 256 GB (using 8 GB DIMMs).
 - Industry standard VLP DDR3 DIMMs.
 - Optional Active Memory Expansion.
- ▶ Disk:
 - One 3 Gb SAS disk storage controller (on the SMP unit).
 - Two disk drive bays support up to two 2.5-inch SAS HDDs or up to four 1.8-inch SAS SSDs with optional SAS to SATA interposers.
- ▶ Hardware mirroring:
 - One HDD: RAID 0.
 - One SSD: RAID 0.
 - Two HDDs: RAID 0 or RAID 10.
 - One HDD and one SSD: RAID 0 on each disk. Combining HDD and SSD in the one RAID configuration is not allowed.
 - Two SSDs: RAID 0 or RAID 10.
 - Three SSDs: RAID 0, RAID 5, or RAID 10 (RAID 10 with only two disks).
 - Four SSDs: RAID 0, RAID 5, RAID 6, or RAID 10.
- ▶ Onboard integrated features:
 - Service processor (one on each blade)
 - Four 1 GB Ethernet ports
 - SAS controller
 - USB Controller that routes to the USB 2.0 port on the media tray
 - One Serial over LAN (SOL) Console through FSP
- ▶ Expansion card I/O options:
 - Two CIOv expansion card slots (PCIe)
 - Two CFFh expansion card slots (PCIe)

Service processor: The service processor (or flexible service processor) on the expansion unit provides control but does not offer redundancy with the SP on the base unit.

3.9.2 Supported BladeCenter chassis

The PS703 and PS704 are supported in the IBM BladeCenter chassis, as listed in Table 3-136.

Table 3-136 The blade servers supported in each BladeCenter chassis

Blade	Machine type-model	Blade width	BC S 8886	BC E 8677	BC T 8720	BC T 8730	BC H 8852	BC HT 8740	BC HT 8750
PS703	7891-73X	1 slot	Yes	No	No	No	Yes	Yes	Yes
PS704	7891-74X	2 slot	Yes	No	No	No	Yes	Yes	Yes

A detailed description of each supported BladeCenter for the PS703 and PS704 blade servers is contained in the following sections:

- ▶ IBM BladeCenter H delivers high performance, extreme reliability, and ultimate flexibility for the most demanding IT environments. See 2.2, “IBM BladeCenter H” on page 46.
- ▶ IBM BladeCenter HT models are designed for high-performance flexible telecommunications environments by supporting high-speed Internet working technologies (such as 10Gb Ethernet). They provide a robust platform for NGNs. See 2.3, “IBM BladeCenter HT” on page 56.
- ▶ 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 environments. See 2.4, “IBM BladeCenter S” on page 67.

The number of POWER7 processor-based blades that can be installed in a BladeCenter chassis depends on several factors:

- ▶ BladeCenter chassis type
- ▶ Number of power supplies installed
- ▶ Power supply voltage option (BladeCenter S only)
- ▶ BladeCenter power domain configuration

Table 3-137 shows the maximum number of PS703 and PS704 blades running in a maximum configuration (memory, disk, and expansion cards) for each supported BladeCenter chassis that can be installed with fully redundant power and without performance reduction. IBM blades that are based on processor types other than POWER7 might reduce these numbers.

Table 3-137 PS703 and PS704 per chassis type

Server	BladeCenter H		BladeCenter HT		BladeCenter S			
	14 slots total		12 slots total		6 slots total			
					110 VAC		208 VAC	
	2 PS	4 PS	2 PS	4 PS	2 PS	4 PS	2 PS	4 PS
PS703	7	14	6	12	2	6	2	6
PS704	3	7	3	6	1	3	1	3

When mixing blades of different processor types in the same BladeCenter, the BladeCenter Power Configurator tool determines whether the combination is valid. For more information about this update, go to the following address:

<http://www.ibm.com/systems/bladecenter/powerconfig>

3.9.3 Minimum features

At minimum, a PS703 requires a BladeCenter chassis, two eight-core 2.4 GHz processors, a minimum memory of 16 GB, zero or one disks, and a Language Group Specify (mandatory for the order voltage nomenclature/language).

At minimum, a PS704 requires a BladeCenter chassis, four eight-core 2.4 GHz processors, a minimum memory of 32 GB, zero or one disks, and a Language Group Specify (mandatory for the order voltage nomenclature/language).

Each system has a minimum feature set to be valid. The minimum system configurations for a PS703 and PS704 are shown in Table 3-138.

Table 3-138 Minimum features for PS703 and PS704 blade server

Category	Minimum features required
BladeCenter chassis	Supported BladeCenter chassis (See 3.9.2, "Supported BladeCenter chassis" on page 395.)
Processor	<ul style="list-style-type: none"> Two 8-core 2.4 GHz processors in a PS703 blade (7891-73X) Four 8-core 2.4 GHz processors in a PS704 blade (7891-74X)
Memory	DDR3 Memory DIMM: For PS703, 16 GB (in either of the following ways): <ul style="list-style-type: none"> Two 8 GB (2x 4 GB DIMMs) DDR3 1066 MHz (#8196) One 16 GB (2x 8 GB DIMMs) DDR3 1066 MHz (#8199) For PS704, 32 GB (in either of the following ways): <ul style="list-style-type: none"> Four 8 GB (2x 4 GB DIMMs) DDR3 1066 MHz (#8196) Two 16 GB (2x 8 GB DIMMs) DDR3 1066 MHz (#8199)
Storage	AIX, Linux, Virtual I/O Server, and IBM i (requires a VIOS partition); either 300, 600, or 177 GB: <ul style="list-style-type: none"> 300 GB SAS 2.5-inch HDD (#8274) 600 GB SAS 2.5-inch HDD (#8276) 177 GB SATA SSD (#8207) If Boot from SAN 8 GB Fibre Channel HBA is selected with FC #8240, #8242, or #8271 or Fibre Channel over Ethernet Adapter FC, #8275 must be ordered. FC #8207 requires FC #4539 - Interposer for 1.8-inch solid-state drives.
1x Language Group	Country specific (selected by the customer)
Operating system	1x primary operating system (one of these) <ul style="list-style-type: none"> AIX (#2146) Linux (#2147) IBM i (#2145) plus IBM i 6.1.1 (#0566) IBM i (#2145) plus IBM i 7.1 (#0567)

3.9.4 Processor

The processors used in the PS703 and PS704 are 64-bit POWER7 processors operating at 2.4 GHz. They are optimized to achieve maximum performance for both the system and its virtual machines. Couple that performance with PowerVM and you are now empowered for massive workload consolidation to drive maximum system use, predictable performance, and cost efficiency.

POWER7 Intelligent Threads Technology enables workload optimization by selecting the most suitable threading mode (Single thread (per core) or Simultaneous Multi-thread 2 or 4 modes (also called 2-SMT and 4-SMT)). The Intelligent Threads Technology can provide improved application performance. In addition, POWER7 processors can maximize cache access to cores, improving performance, using Intelligent Cache technology.

EnergyScale Technology offers Intelligent Energy management features, which can dramatically and dynamically conserve power and further improve energy efficiency. These Intelligent Energy features enable the POWER7 processor to operate at a higher frequency if environmental conditions permit, for increased performance per watt. Alternatively, if user settings permit, these features allow the processor to operate at a reduced frequency for significant energy savings.

The PS703 and PS704 come with a standard processor configuration. There are no optional processor configurations for the PS703 and PS704. The PS703 and PS704 processor configurations are as follows:

- ▶ The PS703 blade server is a single-wide blade that contains two eight-core, 64-bit POWER7 2.4 GHz processors with 256 KB per processor core L2 cache and 4 MB per processor core L3 cache. No processor options are available.
- ▶ The PS704 blade server is a double-wide blade that supports four eight-core, 64-bit POWER7 2.4 GHz processors with 256 KB per processor core L2 cache and 4 MB per processor core L3 cache. No processor options are available.

3.9.5 Memory

The PS703 and PS704 blade servers use industry standard VLP DDR3 memory DIMMs. Memory DIMMs must be installed in matched pairs and have the same size and speed.

The PS703 and PS704 blade servers have 16 (PS703) and 32 (PS704) DIMM slots. Memory is available in 4 GB or 8 GB DIMMs, both operating at a memory speed of 1066 MHz. The memory sizes can be mixed within a system. You can use pairs of 4 GB DIMMs with pairs of 8 GB DIMMs.

The POWER7 DDR3 memory uses a new memory architecture to provide greater bandwidth and capacity, which enables operating at a higher data rate for larger memory configurations.

Table 3-139 shows the DIMM features.

Table 3-139 Memory DIMM options

Feature code	DIMM size	Package includes	Speed
8196	8 GB	Two 4 GB DIMMs	1066 MHz
8199	16 GB	Two 8 GB DIMMs	1066 MHz

Unsupported:

- ▶ The DDR2 DIMMs used in JS23 and JS43 blade servers are not supported in the POWER7 blade servers.
- ▶ The DDR3 DIMMs used in PS700, PS701, and PS702 blade servers are not supported in the PS703 and PS704 servers.

The optional Active Memory Expansion is a POWER7 technology that allows the effective maximum memory capacity to be much larger than the true physical memory. Compression and decompression of memory content using processor cycles can allow memory expansion up to 100%. This configuration enables an AIX V6.1 (or later) partition to do more work with the same physical amount of memory or a server to run more partitions and do more work with the same physical amount of memory.

3.9.6 I/O

The PS703 has one CIOv PCIe expansion card slot and one CFFh PCIe expansion card slot. The PS704 blade server has two CIOv expansion card slots and two CFFh expansion card slots.

Table 3-140 shows the supported CIOv and CFFh expansion cards in the Ps703 and PS704 blade servers.

Table 3-140 Supported I/O expansion cards on POWER7 blades

Card description	Feature code
CIOv	
QLogic 8 Gb Fibre Channel Expansion Card (CIOv)	8242
QLogic 4 Gb Fibre Channel Expansion Card (CIOv)	8241
Emulex 8 Gb Fibre Channel Expansion Card (CIOv)	8240
3 GB SAS Pass-thru Expansion Card (CIOv)	8246
Broadcom 2-Port Gigabit Ethernet Expansion Card (CIOv)	8243
CFFh	
QLogic 8 Gb Fibre Channel Expansion Card (CFFh)	8271
QLogic Ethernet and 4 Gb Fibre Channel Expansion Card (CFFh)	8252
QLogic 2-port 10 Gb Converged Network Adapter (CFFh)	8275
2-port QDR 40 GBps InfiniBand Expansion Card (CFFh)	8272
Broadcom 2/4-port Ethernet Expansion Card (CFFh)	8291

3.9.7 Storage

The PS703 blade server has one disk bay. The bay supports either of the following disks:

- ▶ One 2.5-inch SAS HDD
- ▶ One or two 1.8-inch SATA SSDs

If you choose to use SSDs, then the Interposer for 1.8-inch SSDs, feature code 4539, must also be installed.

The PS704 blade servers have two disk bays (one on the base blade and one in the SMP unit blade):

- ▶ On the base unit, it can have one 2.5-inch SAS HDD.
- ▶ On the base unit, it can have up to two 1.8-inch SATA SSDs.
- ▶ On the expansion unit, it can have one 2.5-inch SAS HDD.
- ▶ On the expansion unit, it can have up to two 1.8-inch SATA SSDs

Table 3-141 lists the supported disk features on the PS703 and PS704 blade servers.

Table 3-141 Supported disk drives

Feature code	Description
2.5-inch SAS drives	
8274	300 GB 10 K SFF SAS HDD
8276	600 GB 10 K SFF SAS HDD
1.8-inch solid-state drive (SSD) and interposer	
8207	177 GB SATA SSD (requires feature code #4539)
4539	Interposer for 1.8-inch SSDs

3.9.8 Standard onboard features

There are many standard onboard features.

Service processor

The service processor (or flexible service processor (FSP)) is the main part of the blade server. It monitors and manages system hardware, resources, and devices. It does system initialization, configuration, and thermal/power management. It takes corrective action as required.

The PS703 has only one service processor. The PS704 blade server has two FSPs (one on each blade). However, the second service processor is only in IO mode and is not redundant to the one on the base blade.

Ethernet ports

The PS703 has a 2-port onboard integrated Ethernet adapter for a total of two Ethernet ports. The PS704 has two 2-port onboard integrated Ethernet adapters with one in the base blade and the second one in the SMP blade for a total of four Ethernet ports.

Ports: The PS703 and PS704 do not have the Host Ethernet Adapters (HEAs) and Integrated Virtual Ethernet (IVE) ports that previous POWER processor-based blade servers have included. A virtual Ethernet can be provided from the Virtual I/O Server virtual network environment.

SAS controller

The PS703 blade server has one integrated SAS controller. The PS704 has one integrated SAS controller on the SMP blade.

The integrated SAS controller is used to drive the local SAS storage. This SAS controller can also support SATA SSD with the addition of the SSD Interposer for 1.8-inch SSDs.

The integrated SAS controller supports hardware mirroring for RAID 0, RAID 1, or RAID 10, depending on the number of drives installed

The 3 Gb SAS Pass-thru Expansion Card can be used to connect to the BladeCenter SAS Connectivity Module, which can be connected to external storage. This SAS Pass-thru Expansion Card can also be used to connect to BladeCenter S chassis internal SAS drives.

USB controller

The USB controller connects the USB bus to the midplane, which is then routed to the media tray in the BladeCenter chassis to connect to USB devices (such as an optical drive or diskette drive).

Serial over LAN (SOL)

The integrated SOL function routes the console data stream over standard dual 1 Gb Ethernet ports to the AMM. The PS703 and PS704 servers do not have onboard video chips and do not support KVM connections. Console access is only by SOL connection. Each blade can have a single SOL session, but there can be multiple Telnet or SSH sessions to the BladeCenter AMM, each acting as a SOL connection to a different blade.

For more information about the BladeCenter PS703 and PS704 products, see *IBM BladeCenter PS703 and PS704 Technical Overview and Introduction*, REDP-4744.

3.10 Expansion blades

There are 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 3-142 lists the available expansion blades and the servers that support them.

Table 3-142 Currently available expansion blades and supported server blades

Expansion unit	Part number	Feature code	HS12 - 8028	HS22 - 7870	HS22V - 7871	HS23 - 7875	HS23E - 8038	HX5 - 7872	HX5 - 7873	PS700/701/702	PS703/704	Page
PCI Express I/O Expansion Unit	43W4391	4339	Y	Y	Y	N	N	N	N	N	N	401
PCI Express Gen 2 Expansion Blade	46M6730	9295	N	Y	N	N	Y	Y	Y	N	N	403
PCI Express Gen 2 Expansion Blade II	68Y7484	A247	N	Y	N	Y	Y	N	Y	N	N	403
GPU Expansion Blade with NVIDIA Tesla M2070	81Y4103 ^a	5090	N	Y	N	N	N	N	N	N	N	408
GPU Expansion Blade with NVIDIA Tesla M2075	46M6771	A25F	N	Y	N	N	Y	N	N	N	N	408
GPU Expansion Blade with NVIDIA Tesla M2070Q	46M6772	A10R	N	Y	N	N	Y	N	N	N	N	408
GPU Expansion Blade II with NVIDIA Tesla M2070Q	68Y7479	A246	N	Y	N	Y	Y	N	N	N	N	408
GPU Expansion Blade II with NVIDIA Tesla M2075	68Y7478	A245	N	Y	N	Y	Y	N	N	N	N	408
GPU Expansion Blade II with NVIDIA Tesla M2090	00D6881	A2VW	N	Y	N	Y	Y	N	N	N	N	408

a. The IBM BladeCenter GPU Expansion Blade is only available through special bid or configure-to-order (CTO) in some locations.

As you can see in Table 3-142, some blades, such as the HS22 type 7870, support more than one expansion blade. It is not possible to install more than one type of expansion blade on a blade server at the same time. Depending on the expansion unit type, it is possible to install several expansion units of that type on a single blade, for example, PCI Express Gen 2 Expansion Blade.

3.10.1 PCI Express I/O Expansion Unit

The PCI Express I/O Expansion Unit (PEU3e) attaches to supported blade servers through a connector on the server system board. The PEU3e has two standard full-length PCI Express slots that support the attachment of two full sized, 25 W PCI Express adapters. These two slots are PCI Express x16 slots mechanically, but they can operate only at speeds up to x8. The PEU3e provides two power connectors that can be used to provide 12 V, through a cable, to PCI Express adapters that have external power connectors.

PCI Express support: Using PCI Express adapters that require more than 25 W requires special testing and support to be completed. The cable or cables used to connect to the 12 V connectors on the PEU3e are only available if this testing is completed successfully. Contact your IBM representative to engage the SPORE process, which provides this special testing.

Table 3-143 shows the ordering part number and feature code for the PCI Express I/O Expansion Unit.

Table 3-143 Ordering part number and feature code

Description	Part number	Feature code
PCI Express I/O Expansion Unit	43W4391	4339

Only one PEU3e can be attached to the blade server. For a discussion of all expansion units and the supported combinations for each blade server type, see 3.10, “Expansion blades” on page 400. Table 3-142 on page 400 lists the blade servers that support the attachment of the PEU3e.

Figure 3-92 shows the PEU3e.



Figure 3-92 PCI Express I/O Expansion Unit (PEU3e)

No PCI-X or PCI Express sockets on the PEU3e 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 3-93 shows the internal circuitry of the PEU3e.




Figure 3-93 PCI Express I/O Expansion Unit internals

Supported PCIe adapter

Table 3-144 lists PCIe adapters that are supported in the PCI Express I/O Expansion Unit.

Table 3-144 PCIe adapters that are supported in the PCI Express I/O Expansion Unit

Description	Part number	Feature code
640 GB High IOPS MLC Duo Adapter	81Y4519	5895
IBM 160 GB High IOPS SS Class SSD PCIe Adapter	46M0877	0096
IBM 320 GB High IOPS MS Class SSD PCIe Adapter	46M0898	1649
IBM Ultra 320 SCSI Controller PCIe	43W4324	3569
ServeRAID-MR10M SAS/SATA Controller	43W4339	3559
NetXtreme II 1000 Express Ethernet Adapter	39Y6066	1485
PRO/1000 PT Quad Port Server Adapter by Intel	39Y6136	2974

Consult the IBM ServerProven website to see the latest list of expansion cards that are supported in the PCIe Expansion Unit:

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

For more information, see the *IBM BladeCenter PCI Express I/O Expansion Unit (PEU3e) Installation Guide*:

<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5076806>

3.10.2 PCI Express Gen 2 Expansion Blade/Expansion Blade II

The IBM BladeCenter PCI Express Gen 2 Expansion Blade and Expansion Blade II provide the capability to attach selected PCI Express cards on selected blade servers. This capability is ideal for many applications that require special telecommunications network interfaces or hardware acceleration using a PCI Express card. It provides one full-height and full-length PCI Express slot and one full-height and half-length PCI Express slot with a 75 W power supply for each slot. It integrates PCI Express card support capability into the BladeCenter architecture.

The unique stacking capability of the IBM BladeCenter PCI Express Gen 2 Expansion Blades allows you to connect up to four of them to a single blade server, providing up to eight PCIe 2.0 slots. In addition, you can still use a CFFh I/O expansion card adapter by installing it in the topmost expansion blade.

Table 3-145 shows the ordering part number and feature code for the PCI Express Gen 2 Expansion Blades.

Table 3-145 Ordering part number and feature code

Description	Part number	Feature code
PCI Express Gen 2 Expansion Blade	46M6730	9295
PCI Express Gen 2 Expansion Blade II	68Y7484	A247

Figure 3-94 shows the IBM BladeCenter PCI Express Gen 2 Expansion Blade attached to an HS22 blade server.




Figure 3-94 IBM BladeCenter PCI Express Gen 2 Expansion Blade (attached to an HS22 blade server)

The only difference between Expansion Blade and Expansion Blade II is cut-outs on the sides of Expansion Blade II, as shown in Figure 3-95.




Figure 3-95 PCI Express Gen 2 Expansion Blade and Expansion Blade II

The IBM BladeCenter PCI Express Gen 2 Expansion Blade and Expansion Blade II have the following features:

- Support for PCIe 2.0 adapters in an expansion blade

Use the expansion blade to install one or two standard form factor PCIe 2.0 adapters in a BladeCenter environment. These adapters enable the usage of specialized adapters or adapters that would not otherwise be available to BladeCenter customers. Each of the two adapters can consume up to 75 W.

- Ability to stack up to four expansion blades on a single base blade

You can attach up to four expansion blades and maintain the BladeCenter density advantage while still giving you the option to install PCIe cards as needed without the need for each server blade and the added complexity and cost that brings. The first expansion blade connects to the server blade using the CFFh expansion slot of the server blade. The second expansion blade attaches to the CFFh connector on the first expansion blade, and so on.

- CFFh slot still available

The CFFh expansion connector is accessible on the topmost expansion blade, even with four expansion blades attached. Use this connector to maintain the integrated networking capabilities of the blade server when it is installed in a BladeCenter H or BladeCenter HT chassis.

Figure 3-96 shows the layout of the IBM BladeCenter PCI Express Gen 2 Expansion Blade.




Figure 3-96 Layout of the IBM BladeCenter PCI Express Gen 2 Expansion Blade

The IBM BladeCenter PCI Express Gen 2 Expansion Blade has a removable PCI Express riser tray that is removed from the frame (slides out of the front of the Expansion Blade) to facilitate the installation of the PCI Express adapters. The PCI Express riser tray is connected to the rest of the PCI Express Expansion Blade circuitry by a flexible cable connector.

Figure 3-97 shows the PCI Express riser extended out from the frame of the expansion blade.




Figure 3-97 Showing the PCI Express riser extended out from the frame

Specifications

The expansion blade has the following specifications:

- ▶ A 30 mm expansion blade that attaches to a supported server blade.
- ▶ Two PCI Express 2.0 slots:
 - Each slot is PCIe x8 electrically, and PCIe x16 mechanically.
 - Each slot supports adapters up to 75 W power consumption.
 - Slot 1 supports full-length/full-height cards.
 - Slot 2 supports half-length/full-height cards.
 - Table 3-146 on page 407 lists the supported cards.
- ▶ Up to four expansion blades can be attached to the base server, depending on the server to which they are attached:
 - HS22: Up to four expansion blades.
 - HS23: Up to four expansion blades.
 - Single-node HX5: Up to three expansion blades.
 - Two-node HX5: Up to two expansion blades.
- ▶ The CFFh slot remains usable on the topmost expansion blade for use in BladeCenter H and HT configurations. Table 3-147 on page 407 lists the supported CFFh cards.

HX5 support: HX5 with MAX5 is not supported with PCI Express Gen2 Expansion Blades.

Supported PCIe adapter

Table 3-146 lists the PCIe adapter that is supported in the PCI Express Gen 2 Expansion Blade and Expansion Blade II.

Table 3-146 PCIe adapter that is supported in the PCI Express Gen 2 Expansion Blades

Description	Part number	Feature code	PEU Gen 2	PEU II Gen 2
IBM 320 GB High IOPS SD Class SSD PCIe Adapter	46M0878	0097	Yes	Yes
IBM 365GB High IOPS MLC Mono Adapter	46C9078	A3J3	Yes	Yes
IBM 785GB High IOPS MLC Mono Adapter	46C9081	A3J4	Yes	Yes
IBM 1.2TB High IOPS MLC Mono Adapter	90Y4377	A3DY	Yes	Yes
IBM 2.4TB High IOPS MLC Duo Adapter	90Y4397	A3DZ	Yes	Yes
IBM 300GB High IOPS MLC Modular Adapter	90Y4361	A3MZ	No	Yes
IBM 600GB High IOPS MLC Modular Adapter	90Y4365	A3N0	No	Yes
IBM 800GB High IOPS MLC Modular Adapter	90Y4369	A3N1	No	Yes
IBM 300GB High IOPS SLC Modular Adapter	90Y4373	A3N2	No	Yes

Consult the IBM ServerProven website to see the most recent list of supported adapter:

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

Adapter availability: If the PCI Express adapter that you require is not on the ServerProven website, use the IBM ServerProven Opportunity Request for Evaluation (SPORE) process to confirm compatibility in the configuration.

Supported CFFh expansion cards

Table 3-147 lists the CFFh expansion cards that are supported in the Expansion Blades.

Table 3-147 Supported CFFh expansion cards

Description	Part number	Feature code	Expansion Blade, 46M6730	Expansion Blade II, 68Y7484
2/4 Port Ethernet Expansion Card (CFFh)	44W4479	5476	Yes	Yes
2-port 40 Gb InfiniBand Expansion Card (CFFh)	46M6001	0056	Yes	Yes
Broadcom 10 Gb Gen 2 4-port Ethernet Expansion Card (CFFh)	46M6164	0098	Yes	Yes
Brocade 2-port 10 GbE Converged Network Adapter	81Y1650	5437	No	Yes
Emulex 10 GbE Virtual Fabric Adapter II	00Y3266	A3NV	No	Yes
Emulex 10 GbE Virtual Fabric Adapter Advanced II	00Y3264	A3NW	No	Yes
Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh)	90Y3570	A1NW	No	Yes
QLogic 2-port 10 Gb Converged Network Adapter (CFFh)	00Y3280 ^a	A3JB	Yes	Yes
QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)	00Y3270 ^b	A3JC	Yes	Yes

a. Replaces 42C1830.

b. Replaces 44X1940.

Consult the IBM ServerProven website to see the latest list of supported CFFh cards:

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

Related publications

For more information, see these documents:

- *IBM BladeCenter PCI Express Gen 2 Expansion Blade Installation and User's Guide:*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5085445>
- *IBM BladeCenter PCI Express Gen 2 Expansion Blade:*
<http://www.redbooks.ibm.com/abstracts/tips0783.html>

3.10.3 IBM BladeCenter GPU Expansion Blade

The IBM BladeCenter GPU Expansion Blade and GPU Expansion Blade II provide the capability to attach next-generation graphics processing unit (GPU) technology on supported server blades. This capability enables many applications to take advantage of the acceleration and visualization performance advantages that are offered in general-purpose computing on GPUs. These products offered are the NVIDIA Tesla M2070, M2070Q, M2075, and M2090 GPUs.

The unique stacking capability of the IBM BladeCenter GPU Expansion Blade allows you to connect up to three of them to a single blade server. In addition, you can still use a CFFh I/O expansion card adapter by installing it in the topmost expansion blade.

Table 3-148 lists the ordering information for the IBM BladeCenter GPU Expansion Blade variants.

Table 3-148 Ordering part numbers and feature codes

Description	Part number	Feature code
GPU Expansion Blade with NVIDIA Tesla M2070	81Y4103 ^a	5090
GPU Expansion Blade with NVIDIA Tesla M2075	46M6771	A25F
GPU Expansion Blade with NVIDIA Tesla M2070Q	46M6772	A10R
GPU Expansion Blade II with NVIDIA Tesla M2070Q	68Y7479	A246
GPU Expansion Blade II with NVIDIA Tesla M2075	68Y7478	A245
GPU Expansion Blade II with NVIDIA Tesla M2090	00D6881	A2VW

a. The IBM BladeCenter GPU Expansion Blade is only available through special bid or configure-to-order (CTO) in some locations.

Figure 3-98 shows the IBM BladeCenter GPU Expansion Blade attached to an HS22 blade server.




Figure 3-98 IBM BladeCenter GPU Expansion Blade (attached to an HS22 blade server)

The IBM BladeCenter GPU Expansion Blade has the following features:

- Support for a high-performance NVIDIA adapter in each expansion blade

The expansion blade includes either an NVIDIA Tesla M2070, NVIDIA Tesla M2075, NVIDIA Tesla M2070Q, or NVIDIA Tesla M2090 adapter. The usage of these expansion blades enables the use of the specialized adapters that would not otherwise be available to BladeCenter customers.

- Ability to stack up to four expansion blades on a single base blade

You can attach up to four expansion blades (it is possible to attach four when they are ordered through the special bid process) and maintain the BladeCenter density advantage. You still can access the power of these high-performance GPUs as needed without needing additional server blades and the added complexity and cost those blades bring. The first expansion blade connects to the server blade using the CFFh expansion slot of the server blade. The second expansion blade attaches to the CFFh connector on the first expansion blade, and so on.

Mixing expansion blades: No mixing of other expansion blades with the IBM BladeCenter GPU Expansion Blade is possible. You cannot attach an IBM BladeCenter PCI Express Gen 2 Expansion Blades to the same stack as the IBM BladeCenter GPU Expansion Blades.

► CFFh slot still available

The CFFh expansion connector is accessible on the topmost expansion blade, even with four expansion blades attached. You can maintain the integrated networking capabilities of the blade server when installed in a BladeCenter H or BladeCenter HT chassis.

The IBM BladeCenter GPU Expansion Blade has a removable PCI Express riser tray that is removed from the frame (slides out of the front of the Expansion Blade) to facilitate access to the NVIDIA GPU adapter. The PCI Express riser tray is connected to the rest of the IBM BladeCenter GPU Expansion Blade circuitry by a flexible cable connector.

Figure 3-99 shows the layout of the IBM BladeCenter GPU Expansion Blade.




Figure 3-99 Layout of the IBM BladeCenter GPU Expansion Blade

Specifications

The expansion blade has the following specifications:

- A 30 mm expansion blade that attaches to a supported server blade.
- A standard NVIDIA GPU computing module.
- Up to four expansion blades can be attached to the base server. With four expansion blades attached, the total width of the server is 150 mm.
- The CFFh slot remains usable on the topmost expansion blade for usage in BladeCenter H and HT configurations. Table 3-149 on page 411 lists the supported CFFh cards.

The NVIDIA M2070, M2075, and M2070Q cards have the following specifications:

- ▶ An NVIDIA Fermi GPU engine (M2070 and M2070Q - Tesla T20 GPU, M2075 - Tesla T20A GPU)
- ▶ 448 processor cores operating at 1.15 GHz
- ▶ 6 GB GDDR5 (graphics DDR) memory operating at 1.566 GHz
- ▶ PCIe x16 host interface
- ▶ 225 W (TDP) power consumption: M2070 and M2070Q
- ▶ 220 W (TDP) power consumption: M2075
- ▶ Double Precision floating point performance (peak): 515 Gflops
- ▶ Single Precision floating point performance (peak): 1.03 Tflops

The NVIDIA M2070Q combines the Fermi GPU with an NVIDIA Quadro advanced visualization engine in the same GPU.

New model: The NVIDIA Tesla M2075 is a next revision of the NVIDIA Tesla M2070. It uses Tesla T20A GPU instead of Tesla T20 GPU on M2070 and has lower power consumption.

The NVIDIA Tesla M2090 card has the following specifications:

- ▶ NVIDIA Fermi GPU engine (Tesla T20A GPU)
- ▶ 512 processor cores operating at 1.3 GHz
- ▶ 6 GB GDDR5 (graphics DDR) memory operating at 1.85 GHz
- ▶ PCIe x16 host interface
- ▶ 225W (TDP) power consumption
- ▶ Double Precision floating point performance (peak): 665 Gflops
- ▶ Single Precision floating point performance (peak): 1.33 Tflops

Supported CFFh expansion cards

Table 3-149 lists the CFFh expansion cards that are supported in the IBM BladeCenter GPU Expansion Blade and GPU Expansion Blade II.

Table 3-149 Supported CFFh expansion cards

Description	Part number	Feature code	GPU Expansion Blade	GPU Expansion Blade II
2/4 Port Ethernet Expansion Card (CFFh)	44W4479	5476	Yes	Yes
2-port 40 Gb InfiniBand Expansion Card (CFFh)	46M6001	0056	Yes	Yes
Broadcom 10 Gb Gen 2 4-port Ethernet Expansion Card (CFFh)	46M6164	0098	Yes	Yes
Brocade 2-port 10 GbE Converged Network Adapter	81Y1650	5437	No	Yes
Emulex 10GbE Virtual Fabric Adapter II	00Y3266	A3NV	Yes	Yes
Emulex 10 GbE Virtual Fabric Adapter Advanced II	00Y3264	A3NW	Yes	Yes
Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh)	90Y3570	A1NW	Yes	Yes
QLogic 2-port 10 Gb Converged Network Adapter (CFFh)	00Y3280 ^a	A3JB	Yes	Yes
QLogic Ethernet and 8 Gb Fibre Channel Expansion Card (CFFh)	00Y3270 ^b	A3JC	Yes	Yes

a. Replaces 42C1830.

b. Replaces 44X1940.

Consult the IBM ServerProven website to see the latest list of CFFh expansion cards that are supported in the IBM BladeCenter GPU Expansion Blade. Go to the following addresses:

- ▶ <http://ibm.com/servers/eserver/serverproven/compat/us/peu/46M6740CT0.html>
- ▶ <http://ibm.com/servers/eserver/serverproven/compat/us/peu/46M6771.html>
- ▶ <http://ibm.com/servers/eserver/serverproven/compat/us/peu/46M6772.html>
- ▶ <http://ibm.com/servers/eserver/serverproven/compat/us/peu/68Y7479.html>
- ▶ <http://ibm.com/servers/eserver/serverproven/compat/us/peu/68Y7478.html>
- ▶ <http://ibm.com/servers/eserver/serverproven/compat/us/peu/46M6772.html>

Related publications

For more information, see these documents:

- ▶ *IBM BladeCenter GPU Expansion Blade Installation Guide:*
<http://publib.boulder.ibm.com/infocenter/bladectr/documentation/topic/com.ibm.bladecenter.25k8373.doc/81Y1131.pdf>
- ▶ *IBM BladeCenter GPU Expansion Blade and GPU Expansion Blade II:*
<http://www.redbooks.ibm.com/abstracts/tips0798.html>

4



Network integration

This chapter describes different aspects of planning and implementing an IBM BladeCenter based network infrastructure. You must 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 other considerations.

This chapter includes the following topics:

- ▶ 4.1, “Ethernet switch module selection” on page 414
- ▶ 4.2, “VLAN” on page 415
- ▶ 4.3, “High availability and redundancy” on page 416
- ▶ 4.4, “Performance” on page 420
- ▶ 4.5, “Systems management” on page 422
- ▶ 4.6, “IBM switch stacking” on page 430
- ▶ 4.7, “Virtual Blade Switch technology” on page 432
- ▶ 4.8, “IBM Fabric Manager” on page 433
- ▶ 4.9, “BladeCenter Virtual Fabric vNIC solution” on page 439

4.1 Ethernet switch module selection

Several I/O modules can be used to provide network connectivity. They include Ethernet switch modules that provide integrated switching capabilities for BladeCenter, and pass-through 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 functions and simplified cabling. However, under some circumstances (for example, specific security policy or certain network requirements), it is not possible to use integrated switching capability, so pass-through modules are needed.

Ensure that the external interface ports of the switches selected are compatible with the physical cabling used or planned to be used in the client's data center (for information about supported cables, see 2.7, "Ethernet Switch Modules" on page 84).

Also, ensure that features and functions required in the network are supported by the proposed switches modules (see Table 2-34 on page 85 and switch module-related information in 2.7, "Ethernet Switch Modules" on page 84).

Currently, you can use a wide range of network technologies in the BladeCenter chassis. You can start 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). In addition, converged 10 Gb Ethernet-based networking solutions are available to further expand BladeCenter networking capabilities.

Table 4-1 lists common selection considerations that might be useful to you when you select an appropriate switch module. Table 2-34 on page 85 has an expanded feature comparison.

Table 4-1 Switch module selection criteria

Suitable switch module →	IBM Server Connectivity Module	IBM Layer 2/3 Copper GbE Switch Module	IBM Layer 2/3 Fiber GbE Switch Module	IBM 1/10 Gb Uplink Ethernet Switch Module	IBM Layer 2-7 Gigabit Ethernet Switch Module	Cisco Catalyst Switch Module 3012	Cisco Catalyst Switch Module 3110G	Cisco Catalyst Switch Module 3110X	Cisco Nexus 4001I Switch Module	IBM Virtual Fabric 10 Gb Switch Module	Brocade Converged 10 GbE Switch Module
Requirement											
Gigabit Ethernet to Blades/ Gigabit Ethernet Uplinks	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	No
Gigabit Ethernet to Blades/ 10 Gb Ethernet Uplinks	No	No	No	Yes	No	No	No	Yes	Yes	Yes	No
10 Gb Ethernet to Blades/ 10 Gb Ethernet Uplinks	No	No	No	No	No	No	No	No	Yes	Yes	Yes
Basic Layer 2 switching (VLAN, port aggregation)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Advanced Layer 2 switching: IEEE based features (STP, QoS)	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Advanced Layer 2 switching: proprietary features (CDP, VTP)	No	No	No	No	No	Yes	Yes	Yes	Yes	No	No

Suitable switch module →	IBM Server Connectivity Module	IBM Layer 2/3 Copper GbE Switch Module	IBM Layer 2/3 Fiber GbE Switch Module	IBM 1/10 Gb Uplink Ethernet Switch Module	IBM Layer 2-7 Gigabit Ethernet Switch Module	Cisco Catalyst Switch Module 3012	Cisco Catalyst Switch Module 3110G	Cisco Catalyst Switch Module 3110X	Cisco Nexus 4001 Switch Module	IBM Virtual Fabric 10 Gb Switch Module	Brocade Converged 10 GbE Switch Module
Requirement											
Layer 3 IPv4 switching (forwarding, routing, ACL filtering)	No	Yes	Yes	Yes	Yes	Yes ^a	Yes ^b	Yes ^b	No	Yes	Yes ^c
Layer 3 IPv6 switching (forwarding, routing, ACL filtering)	No	No	No	Yes ^d	No	No	Yes ^b	Yes ^b	No	Yes ^d	No
Layer 4-7 switching (content switching, traffic inspection, SLB)	No	No	No	No	Yes	No	No	No	No	No	No
10 Gb Ethernet CEE/FCoE	No	No	No	No	No	No	No	No	Yes ^e	Yes	Yes
Switch stacking	No	No	No	Yes	No	No	Yes	Yes	No	Yes	No
vNIC support (see 4.9, “BladeCenter Virtual Fabric vNIC solution” on page 439)	No	No	No	No	No	No	No	No	No	Yes	No

a. Supports only IP forwarding based on static routes and routes learned through RIP.

b. IP Services software is required to support both IPv4 and IPv6 routing and ACL filtering. (IP forwarding is supported in base software with static routes and routes learned through RIP.)

c. Supports only static routes.

d. Supports IPv6 forwarding based on static routes and OSPF v3 (no IPv6 ACL filtering support).

e. A Software Upgrade License is required to support this feature.

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 the 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 described in more detail in 4.5, “Systems management” on page 422.

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 configuration can be set by using the Broadcom Advanced Server Program (BASP) software (which is also used for NIC teaming, as described in 4.4.3, “NIC teaming” on page 421). Each logical interface appears 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 is 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 example, not enough interfaces or low traffic). It might also help to implement strict security policies for separating network traffic by using VLANs while having access to server resources from different VLANs, without needing to implement Layer 3 routing in the network.

To ensure 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.

For information about the switch module's supported features and protocols, see Table 2-34 on page 85 and the switch module-related information in 2.7, "Ethernet Switch Modules" on page 84.

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 requirement 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 describes 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 the cables that connect 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; 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 that illustrate this redundancy, see Figure 2-3 on page 42 and Figure 2-9 on page 51.

- ▶ Implement technologies that provide automatic failover in case of any failure. This task can be accomplished 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 achieve a higher level of availability in a 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 (Figure 4-1).

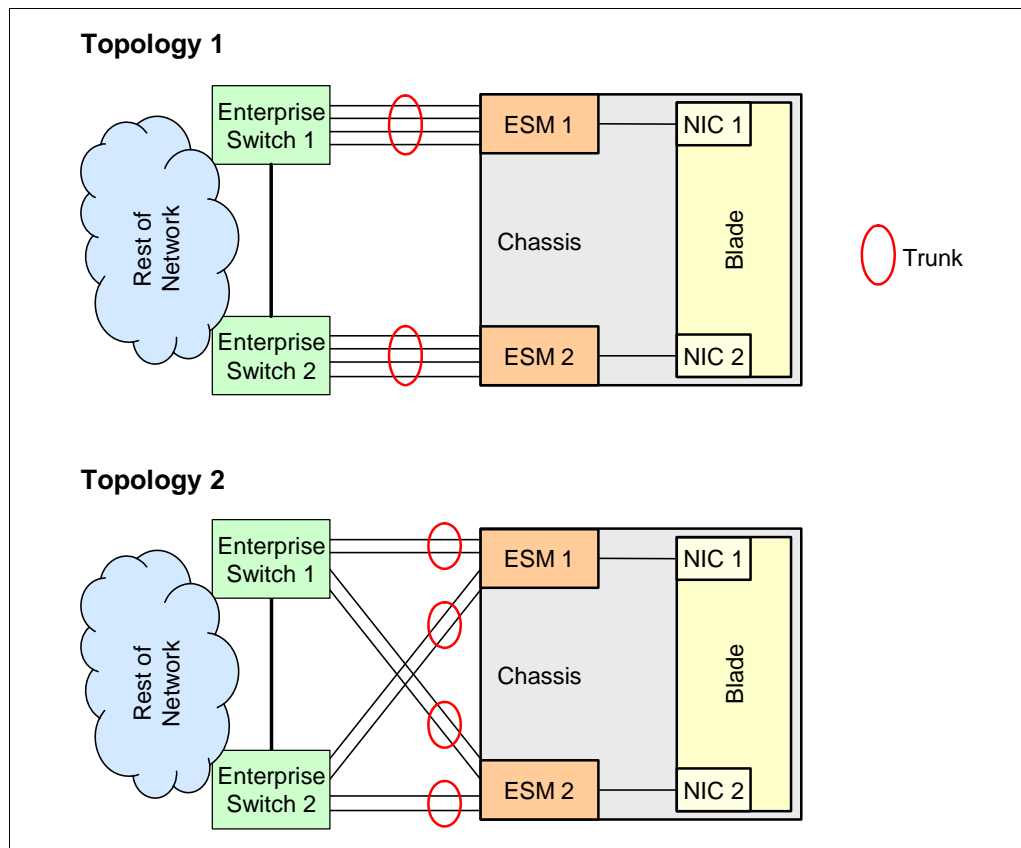


Figure 4-1 IBM BladeCenter redundant LAN integration topologies

Topology 1 in Figure 4-1 has each ESM in BladeCenter directly connected to the one of enterprise switches through aggregation links by 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, must be implemented. This configuration 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 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) 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 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. If there is a link failure, STP reconfigures the network and activates the previously disabled link. The process of reconfiguration can take tenths of a second, and service is unavailable during this period.

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 417). In other cases, STP is not required.

Whenever possible, plan to use trunks (aggregated links) with VLAN tagging for interswitch connections. This configuration can help achieve higher performance by increasing interswitch bandwidth, and 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 420.

STP modifications, such as Port Fast Forwarding or Uplink Fast, might help improve STP convergence time and 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 that 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 is disabled and the other carries traffic from all VLANs. However, if two STP instances are running, then one link is disabled for one set of VLANs while carrying traffic from another set of VLANs, and vice versa. In other words, both links are active, thus enabling more efficient use of the available bandwidth.

For a list of switches that support STP, see Table 2-34 on page 85. For more information, see 2.7, “Ethernet Switch Modules” on page 84.

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 configuration is known as *NIC teaming*. 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 on standby (waiting for the failure of the active port). One port (for example, the active port) should be connected to the ESM in I/O bay 1, and another port (for example, the standby port) should be connected to the 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).

If there is a 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 action is done quickly, within a few seconds. After the restoration of the failed link, the teaming driver can perform a failback or can do nothing, depending on the configuration.

Look at topology 1 in Figure 4-1 on page 417. Assume that NIC teaming is on and that the blade NIC port connected to ESM 1 is active and the other is on standby. If something goes wrong with the internal link to ESM 1, then the teaming driver performs a failover. But what happens 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. Trunk Failover is a special feature supported only on IBM BladeCenter ESMs. Thus, if Trunk Failover is enabled and you lose connectivity to Enterprise Switch 1, then the NIC teaming driver performs a failover and service is 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 example, 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 and 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. For more information, consult the corresponding documentation.

For a list of switches that support trunk failover, see Table 2-34 on page 85.

Automatic fallback: Do not use automatic fallback for NIC teaming to avoid possible issues when you replace a failed ESM. A newly installed ESM has no configuration data, and it can cause service disruption.

For more information, see 2.7, “Ethernet Switch Modules” on page 84.

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 is up and running on another Layer 3 device.

Some ESMs for IBM BladeCenter can be used for VRRP. For a list of switches that support VRRP, see Table 2-34 on page 85. For more information, see 2.7, “Ethernet Switch Modules” on page 84.

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 that exist 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).

For a list of switches that support routing protocols, see Table 2-34 on page 85. For more information, see 2.7, “Ethernet Switch Modules” on page 84.

4.4 Performance

Another major topic to be considered during BladeCenter network planning is network performance. Planning network performance is a complicated task, so the following sections provide guidance about 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. If one of the physical links in the trunk fails, then traffic is 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
- ▶ Brocade Converged 10 GbE Switch Module
- ▶ IBM L2/3 Copper GbE Switch Module for BladeCenter
- ▶ IBM L2/3 Fiber GbE Switch Module for BladeCenter
- ▶ IBM Layer 2-7 GbE Switch Module
- ▶ IBM 1/10 Gb Uplink Ethernet Switch Module
- ▶ IBM Virtual Fabric 10 Gb Ethernet Switch Module
- ▶ Cisco Catalyst Switch Module 3110G
- ▶ Cisco Catalyst Switch Module 3110X
- ▶ Cisco Catalyst Switch Module 3012
- ▶ Cisco Nexus 4001I Switch Module

Additionally, Cisco Systems uses its own protocol called Port Aggregation Protocol (PAgP). It is supported by the following Cisco ESMs:

- ▶ Cisco Catalyst Switch Module 3110G
- ▶ Cisco Catalyst Switch Module 3110X
- ▶ Cisco Catalyst Switch Module 3012

For more information, see 2.7, “Ethernet Switch Modules” on page 84.

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 increase in size can speed up server network processing, and can provide better usage of the network. For example, jumbo frames can be helpful in iSCSI implementations.

Jumbo frames must be supported by all network devices in the communication path. For example, 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. For a list of switches that support jumbo frames, see Table 2-34 on page 85. For more information, see 2.7, “Ethernet Switch Modules” on page 84.

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 418. However, in some cases, getting more network bandwidth for certain servers is necessary. NIC teaming can help you get that bandwidth.

Static trunking or 802.1ad modes of NIC teaming can be used for interfaces connected to the same ESM (for example, a blade’s integrated Ethernet port and a port on an expansion card installed in 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.

For more information, see the relevant BASP documentation.

4.4.4 Server Load Balancing

In a scale-out environment, the performance of network applications can be increased by implementing load balancing clusters. Various methods exist to complete this task:

- ▶ IP load balancing (for example, Microsoft Network Load Balancing)
- ▶ Application load balancing that uses specific software features (for example, IBM WebSphere® Load Balancer)
- ▶ Application load balancing that uses hardware features integrated into network devices (for example, Server Load Balancing with IBM L2/7 GbE Switch for IBM BladeCenter)

Consider using load balancing features of IBM 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 a virtual server concept similar to the virtual router described in 4.3.4, "Virtual Router Redundancy Protocol" on page 419, 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 IBM L2/7 GbE Switch Module for IBM BladeCenter.

For more information, see 2.7, "Ethernet Switch Modules" on page 84.

4.5 Systems management

An important part of the 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):

- ▶ Advanced Management Module (AMM), external 100 Mbps Ethernet
- ▶ AMM, internal 100 MBps Ethernet interfaces for I/O module and blade server management
- ▶ AMM, internal RS-485 interface for blade server management
- ▶ AMM, internal I2C bus interface for I/O module and other modules management

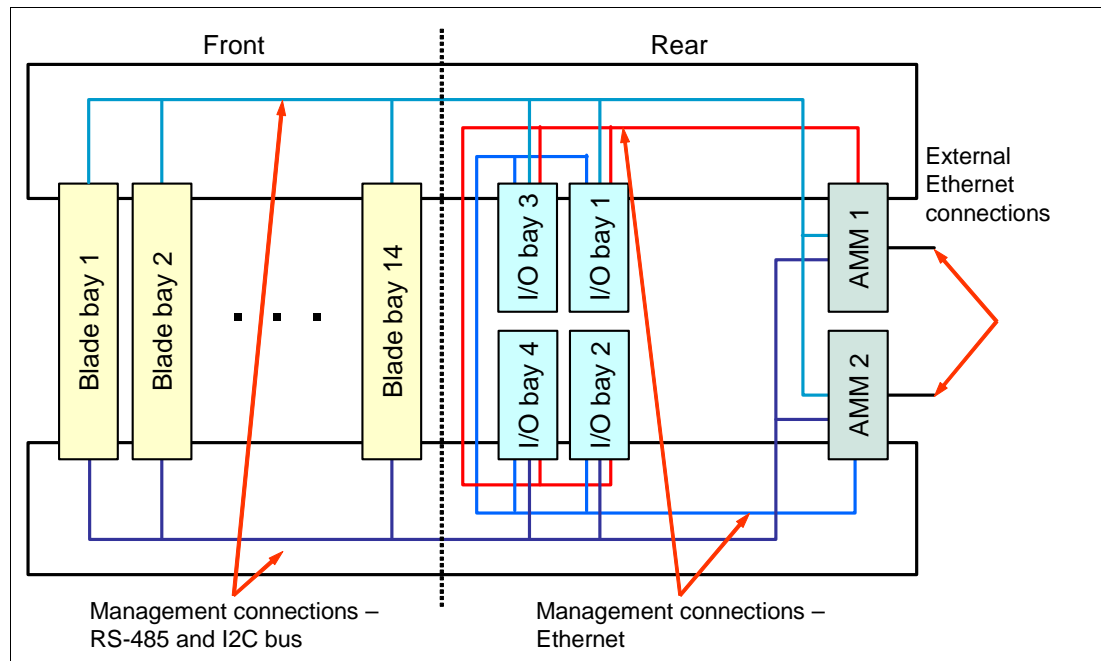


Figure 4-2 Management paths in IBM BladeCenter E chassis

Although Figure 4-2 on page 422 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. BladeCenter S has six blade bays and only one AMM installed.

For redundancy purposes, each component has dual management paths. For example, 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.

If you plan to use management module redundancy, use another set of KVM cables for cabling redundant management modules. If a management module fails in this case, you can select another console port without manually reconfiguring the KVM infrastructure.

The same rule applies to Ethernet cabling. Connect Ethernet ports on both management modules to the network to be able to manage remotely in case of failover without any reconfiguration.

Important:

- ▶ Only one KVM cable is required both for single and redundant management modules for BladeCenter T, because only one KVM module is in the chassis.
- ▶ LAN module on BladeCenter T has two Ethernet ports, so connect both interfaces to the network to enable redundancy.

Devices requiring management in the IBM BladeCenter chassis include I/O modules and blade servers:

- ▶ Pass-through 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 AMM management path. If you plan to use external interfaces of ESMs for management purposes, then certain considerations must be taken into account, as described in 4.5.2, “Management network topology” on page 424.

Each blade server has two RS-485 management paths, and each I/O module and other module, such as FAN or a power supply, has two I2C bus management paths (one for each management module). Only one path is active at the same time. This path is used for functions such as blade power control, VPD, error logs, analog remote control, and others.

Each blade server also has two internal management paths through an I/O module (one for each management module). Only one path is active at the same time, and the path is routed to the AMM through I/O module. This path is used for functions such as Serial-over-LAN (SOL), Concurrent Keyboard Video and Mouse (cKVM), BSMP service data and image flashing, Chassis Internal Network (CIN), and others.

By default, the internal management path uses the integrated Ethernet port and is routed to the AMM through the I/O module in bay 1 (or bay 2). With Management Channel Auto Discovery (MCAD), which is a feature of the AMM, the AMM uses a path based on the expansion cards installed on a blade and I/O modules installed in the chassis. The path is given a higher preference on a high speed I/O module over low speed I/O modules and is done automatically. The user cannot select the internal management path. MCAD is only supported on IMM based systems such as HS22, HS22V, and HX5, with corresponding adapters and I/O modules.

SOL traffic uses an internal management path to route the management module. You can implement SOL only through the management module, and a SOL-capable Ethernet switch or the Intelligent Copper Pass-thru Module is required (other pass-through modules do not support SOL). For a list of SOL-capable switches, see Table 2-34 on page 85.

The cKVM uses the same technology as SOL to transmit KVM signals, including video from multiple blades simultaneously. This setup applies to HS22 with integrated cKVM functionality and to other blades that support cKVM with an optional cKVM feature card. Standard remote KVM control (non-concurrent or analog KVM) uses the RS-485 management path and does not require a SOL-capable I/O module.

Whenever possible, plan to perform all management tasks through management module management paths. Furthermore, management tasks for some I/O modules, such as FCSM or IBM SCM, can be done only through the management module interface.

4.5.2 Management network topology

Figure 4-3 shows the preferred management network topology.

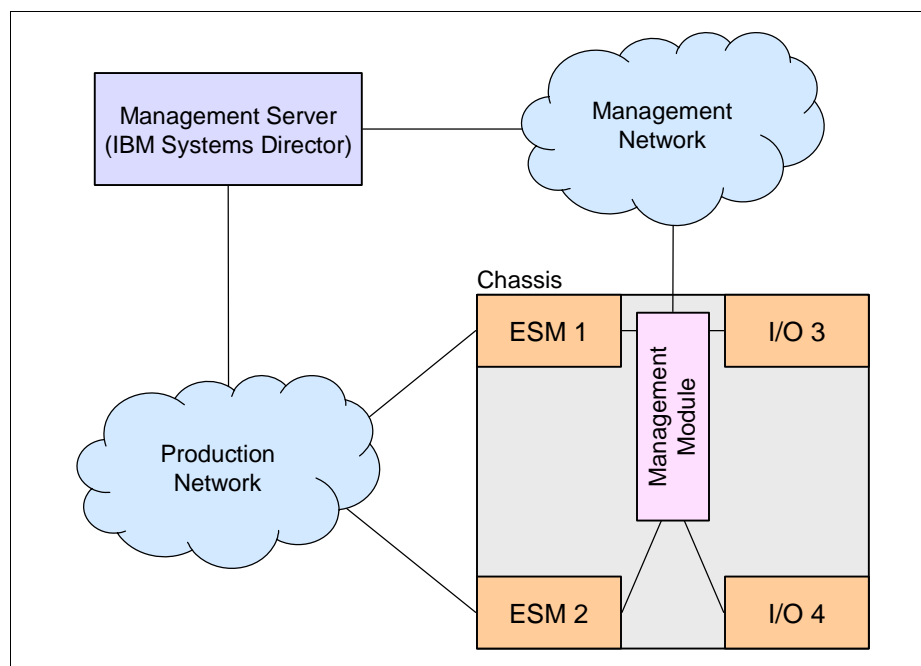


Figure 4-3 Preferred management network topology

Although AMM has two Ethernet ports (one internal and one external), these two ports share a single IP address used for both external and internal communications. You should therefore plan your management network so. For example, use the same subnet for management module and I/O module management interfaces, as we describe later in this chapter.

Static IP addresses: Use static IP addresses (or at least reservations for DHCP) for the AMM to avoid management issues.

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) that differs from the production network.

However, if you plan to use a management server such as IBM Systems Director (and it is preferable 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. This network cannot be combined with the management network (that is, these networks *must* have different IP subnets or IP networks). Otherwise, this setup can lead to unpredictable network behavior and connectivity losses.

Sometimes a management VLAN is implemented across the enterprise network for network management. In this case, the correct way to integrate BladeCenter into the VLAN is to connect the AMM external interface to a port of an external switch assigned to the management VLAN. Then, assign IP addresses from this management subnet to the AMM interface and all internal management interfaces of I/O modules.

Management interfaces: 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 scenario.

All internal management ports on all IBM switches are always in VLAN 4095, and this configuration cannot be changed. IP addresses assigned by the management module are also always in VLAN 4095. So, if you want to manage a switch through 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.

Under some circumstances, such as when the server administration team and the network administration team are two separate teams, a requirement is that the server team not be able to control network switch settings. IBM implemented the Protected Mode feature in the AMM to deal with these situations. Protected Mode, when enabled, prevents the management module from performing switch administration tasks, such as changing management IP addresses, resetting a switch configuration, or answering ARP requests on behalf of the switch IP address. The switch's firmware must support Protected Mode also.

For list of switches that support protected mode, see Table 2-34 on page 85.

Managing Ethernet switches: Even if you are not planning to manage Ethernet switches through the management module and 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.

Use a dedicated management server (such as IBM Systems 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 IBM Tivoli Provisioning Manager for OS Deployment, then some requirements must be met, especially in the case of a routed network. For example, 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. See 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.

IBM ToolsCenter

IBM ToolsCenter is a collection of server management tools to help manage your IBM BladeCenter environment. ToolsCenter makes managing your server environment less complicated, more productive, and cost-effective. IBM ToolsCenter provides tools that are available for the following four areas:

- ▶ Deployment
- ▶ Configuration
- ▶ Updates
- ▶ Diagnostics

Table 4-2 lists the tools that ToolsCenter provides for each area.

Table 4-2 Tools that ToolsCenter provides for each area

Task	Tool	Description
Deployment	ServerGuide	You can use the ServerGuide to implement an easy Windows installation. (Unattended mode)
	Start Now Advisor	You can use the Start Now Advisor to implement a quick configuration for BladeCenter S components.
	Bootable Media Creator (BoMC)	You can use the Bootable Media Creator to perform a firmware update by using a DVD, CD, or USB memory key. An OS is not required.
	ServerGuide Scripting Toolkit	You can use the ServerGuide Scripting Toolkit to implement a hardware deployment solution for Windows and Linux installations.
Configuration	Advanced Settings Utility (ASU)	You can use the Advanced Settings Utility to modify the firmware settings by using the command line on multiple operation system platforms.
	Storage Configuration Manager (SCM)	You can use the Storage Configuration Manager to manage SAS-based BladeCenter storage.
	MegaRAID Storage Manager (MRSM)	You can use the MegaRAID Storage Manager to configure and monitor LSI AID controllers.

Task	Tool	Description
Updates	UpdateXpress System Pack Installer (UXSPI)	The UpdateXpress System Packs contains an integration tested bundle of online firmware and device driver updates for IBM BladeCenter.
Diagnostics	Dynamic System Analysis (DSA)	You can use Dynamic System Analysis to collect system information and analyze the system. Preboot DSA is also available.
	Service Advisor (SA)	Service Advisor provides automate hardware problem reporting by electronically sending BladeCenter chassis hardware events to IBM for analysis.

For additional information about IBM ToolsCenter, go to the following address:

<http://ibm.com/support/entry/portal/docdisplay?brand=5000008&Indocid=TOOL-CENTER>

IBM Systems Director

IBM Systems Director is the platform management backbone that is used to achieve Smart Computing. An integral component of the IBM Smarter Systems portfolio, IBM Systems Director enables integration with IBM Tivoli products and third-party management platforms, providing the building block for integrated services management.

With Systems Director, you can achieve the following goals:

- ▶ Automate data center operations by implementing cloud-ready virtual infrastructures
- ▶ Unify the management of physical and virtual resources for IBM servers, storage, and networks
- ▶ Simplify the management of optimized systems
- ▶ Achieve a single view of the actual energy usage throughout your data center

IBM Systems Director provides, through a single user interface, a consistent view of managed systems, determines how these systems relate to one another, and identifies their statuses, thus correlating technical resources with business needs. A set of common tasks included with IBM Systems Director provides many of the core capabilities required for basic management, which means instant business value. These common tasks include discovery, inventory, configuration, system health, monitoring, updates, event notification, and automation across managed systems.

You can add plug-ins to IBM Systems Director that provide advanced functions or functions tailored to a particular environment. The following IBM Systems Director plug-ins are available for IBM BladeCenter:

- ▶ Active Energy Manager
- ▶ AIX profile Manager for IBM Systems Director
- ▶ IBM Systems Director Network Control
- ▶ Service and Support Manager
- ▶ IBM Systems Director Storage Control
- ▶ IBM Systems Director VMControl™
- ▶ IBM PowerVM Workload Partitions Manager™ for AIX
- ▶ IBM PowerHA® SystemMirror® for IBM Systems Director
- ▶ Various IBM BladeCenter I/O Module Plug-in Extensions (see Table 4-3 on page 428)

IBM BladeCenter I/O Module Plug-in Extensions, together with the IBM Systems Director Network Control add-on, provide centralized and unified management of both servers and I/O modules from a single console.

Table 4-3 lists the BladeCenter I/O modules supported by I/O Modules Plug-in Extensions.

Table 4-3 BladeCenter I/O modules supported by IBM Systems Director I/O Module Plug-in Extensions

Description	Part number	Feature code
Ethernet switch modules		
Cisco Nexus 4001I Switch Module	46M6071	0072
IBM Server Connectivity Module	39Y9324	1484
IBM L2/3 Copper GbE Switch Module	32R1860	1495
IBM L2/3 Fiber GbE Switch Module	32R1861	1496
IBM L2-7 Gb Ethernet Switch Module	32R1859	1494
IBM 1/10 Gb Uplink Ethernet Switch Module	44W4404	1590
IBM Virtual Fabric 10 Gb Switch Module	46C7191	1639
Fibre Channel Switch Modules		
QLogic Virtual Fabric Extension Module	46M6172	4799

IBM BladeCenter management can also be integrated into enterprise management systems, such as Tivoli products and third-party management platforms. With an extensive suite of available platform management tasks and automated tools, IBM Systems Director assists systems management personnel in increasing productivity, resulting in improved responsiveness and service.

For more information about IBM Systems Director and plug-ins, see the following websites:

- IBM Systems Director Partner Integration:
<http://www.ibm.com/systems/software/director/downloads/integration.html>
- IBM Systems Director 6.3 Information Center:
<http://publib.boulder.ibm.com/infocenter/director/pubs/index.jsp>

IBM Systems Director Editions

IBM Systems Director Editions provide options to select the level of leading management solutions from IBM based on your needs, from basic hardware management to more advanced cross enterprise management. IBM Systems Director Editions are packaged to make it simpler, easier to deploy, and more cost-effective to purchase the features needed to manage servers, storage, and networks in your data center.

IBM offers integrated management solutions, from platform management through service management, to drive greater operational management efficiencies and break down the barriers that get in the way to achieving Smart Computing. IBM Systems Director Editions provide three types of Editions for IBM x86 and Power Systems:

- IBM Systems Director Express Edition
- IBM Systems Director Standard Edition
- IBM Systems Director Enterprise Edition

These IBM Systems Director Editions provide an integrated set of tools to empower your IT organization to address major concerns associated with managing both physical and virtual server infrastructures. IBM Systems Director Editions reduce operational complexity, improve the efficiency of IT staff and systems, and control costs while meeting the demanding business requirements for service delivery.

Table 4-4 lists the features that IBM Systems Director Editions provide. IBM Systems Director Express Edition and IBM Systems Director Standard Edition are available for IBM x86 systems. All IBM Systems Director Editions are available for Power Systems.

Table 4-4 The features of IBM Systems Director Editions

Features	Express	Standard	Enterprise
Visualize physical/virtual system relationship	Y	Y	Y
Monitor system health	Y	Y	Y
Provide threshold and error alerts	Y	Y	Y
Update operation systems and firmware	Y	Y	Y
Simplify deployment with virtual images		Y ^a	Y
Control energy use within existing capacity		Y	Y
Monitor network system health with servers and storage		Y	Y
Automate configuration and placement for new workloads			Y
Manage workload availability end-to-end			Y
Understand capacity			Y
Analyze and report historical performance			Y

a. Power Systems only

For more information about IBM Systems Director Editions, see the IBM Systems Director editions website:

<http://www.ibm.com/systems/software/director/editions/>

IBM Fabric Manager

IBM Fabric Manager (IFM) is a solution that enables you to quickly replace and recover IBM BladeCenter blade servers and IBM Flex System compute nodes in your environment. IFM assigns Ethernet Media Access Control (MAC), Fibre Channel worldwide name (WWN), and SAS WWN addresses so that any servers plugged into those bays take on the assigned addresses. This design enables the Ethernet and Fibre Channel infrastructure to be configured once and before any servers are connected to the chassis. For more information, see 4.8, “IBM Fabric Manager” on page 433.

Other management tools

For the management of individual servers and modules, the simplest (from the implementation 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 the following tools:

- ▶ IBM System Networking Switch Center (SNSC) for centralized management of IBM ESMs using SNMP. This tool can be integrated with IBM Systems Director and VMware Virtual Center and vSphere clients.
- ▶ Cisco Network Assistant to graphically manage groups of Cisco ESMs.
- ▶ SAN Surfer or QuickTools for managing QLogic Fibre Channel switch modules.
- ▶ Brocade DataCenter Fabric Manager to manage Brocade-based Fibre Channel fabrics.

4.6 IBM switch stacking

Starting with IBM Networking OS (N/OS) 6.3, the IBM 1/10 Gb Ethernet Uplink Switch and IBM 10 Gb Virtual Fabric Switch support switch stacking. A stack is a group of up to eight Virtual Fabric 10 Gb Switch Modules or 1/10 Gb Uplink ESM switches (but not both) with N/OS that work together as a unified system and are managed as a single entity.

A stack has the following properties, regardless of the number of switches included:

- ▶ The network views the stack as a single entity.
- ▶ The stack can be accessed and managed as a whole using standard switch IP interfaces.
- ▶ After the stacking links are established, the number of ports available in a stack equals the total number of remaining ports of all the switches that are part of the stack.
- ▶ The number of available IP interfaces, VLANs, Trunks, Trunk Links, and other switch attributes are not aggregated among the switches in a stack. The totals for the stack as a whole are the same as for any single switch configured in stand-alone mode.

Before IBM switches can form a stack, they must meet the following requirements:

- ▶ All switches must be the same model (Virtual Fabric 10 Gb Switch Module or /10 Gb Uplink ESM, but not both).
- ▶ Each switch must be installed with N/OS Version 6.3 or later. The same release version is not required, as the Master switch pushes a firmware image to each differing switch that is part of the stack.
- ▶ The preferred stacking topology is a bidirectional ring (see Figure 4-4 on page 431). To achieve this topology, two external 10 Gb Ethernet ports on each switch must be reserved for stacking. By default, the first two 10 Gb Ethernet ports are used.
- ▶ For enhanced redundancy when creating port trunks, include ports from different stack members in the trunks.




Figure 4-4 IBM Switch stacking bidirectional topology

When switches operate in stacking mode, the following features are not supported:

- ▶ Active Multipath Protocol (AMP)
- ▶ Port flood blocking
- ▶ Link Layer Detection Protocol (LLDP)
- ▶ Protocol-based VLANs
- ▶ Routing protocols (RIP, OSPF, OSPFv3, and BGP)
- ▶ IPv6
- ▶ Virtual Router Redundancy Protocol (VRRP)
- ▶ Route maps
- ▶ MAC address notification
- ▶ Adding static MAC addresses
- ▶ Static multicast
- ▶ IGMP Relay and IGMPv3
- ▶ CEE and FCoE

4.7 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. Using this technology, you can 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.

Figure 4-5 shows a sample topology.




Figure 4-5 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 a VBS stack, 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.
 - There is a single configuration file per stack.
 - There is a single software upgrade for the entire stack.
- ▶ A single Spanning Tree Protocol (STP) instance for Layer 2 networks.
- ▶ A 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 consider building one VBS stack (or more VBS stacks, if the network paths from different server NICs must be physically separated) from the switches in the same rack whenever possible.
- ▶ In a master switch failure, any other member switch is eligible to become the master to provide availability and redundancy.
- ▶ In 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 by using supported link aggregation protocols such as 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 that go out of the rack; 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 that go out of rack, you may choose to use just two 10 Gb links going out of rack.
- 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 is 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, and MSTP), link aggregation, link state tracking, and routing, are supported on the VBS level as well.

Switch stack formation: When the switch stack is formed, the AMM cannot manage any member of the stack (including the master switch) over IP. Management over external ports or a serial console cable is required.

4.8 IBM Fabric Manager

IBM Fabric Manager for IBM BladeCenter allows address virtualization, server pre-provisioning, and automatic failover of compute nodes to enable speed of deployment and high availability through the dynamic provisioning of IT resources in a production environment. This application replaces both BladeCenter Open Fabric Manager Basic and Advanced applications and combines them in one package.

IBM Fabric Manager has the following features:

- High availability with redundant Fabric Manager servers
- Option to run Fabric Manager server as a service
- Protection with unique addresses for a newly defined profile
- Remote console
- Profile-based user interface
- Combination of both BladeCenter Open Fabric Manager Basic and Advanced applications in one easy-to-use package

IBM Fabric Manager overview

Fabric Manager software for BladeCenter is designed to easily manage I/O and network interconnects by virtualizing network parameters, such as the worldwide name (WWN) and Media Access Control (MAC) addresses. When a compute node is replaced or failover occurs from one compute node to another, the LAN and SAN configurations are not affected. This application replaces both BladeCenter Open Fabric Manager Basic and BladeCenter Open Fabric Manager Advanced, which were previously used on BladeCenter servers.

After the Fabric Manager for BladeCenter is installed, you can preconfigure LAN and SAN connections. The I/O connections are made automatically when compute nodes are plugged into the chassis slots. No special tools and training are required; you manage the program with an easy-to-use web-based user interface.

IBM Fabric Manager offers these benefits:

- ▶ Time savings: Preconfigure more than 1,400 LAN and SAN connections once for each blade server.
- ▶ Simplified management: Manage growth and complexity from a single Fabric Manager user interface.
- ▶ Flexibility: Fabric Manager works with Ethernet, Fibre Channel, and SAS switch modules.
- ▶ Ease of use: Profile-based user interface enables easy setup, deployment, and management.
- ▶ Reduced risk: Manage risk to keep business operations running with available I/O failover to standby compute nodes and redundant Fabric Manager servers.
- ▶ Convenient packaging: Purchase both BladeCenter Open Fabric Manager Basic and Advanced functions in one easy-to-use package.
- ▶ New Enterprise Licensing offering: Provides time-saving benefits in deploying multiple chassis.
- ▶ Flexibility: Ability to upgrade from Basic or Advanced through a single part number.

With IBM Fabric Manager, you can do the following actions:

- ▶ Quickly replace and recover compute nodes in the data center. IBM Fabric Manager assigns Ethernet MAC, Fibre Channel WWN, and SAS WWN addresses so that any compute nodes plugged into those slots take on the assigned addresses. The Ethernet and Fibre Channel infrastructure is configured once before any compute nodes are connected to the chassis.
- ▶ Monitor the health of compute nodes. Automatically, without user intervention, replace a failed compute node from a designated pool of spare compute nodes. After receiving a failure alert, Fabric Manager attempts the following actions:
 - Power off the failing compute node.
 - Read the Fabric Manager virtualized addresses and boot target parameters.
 - Apply these parameters to the next compute node in the standby pool.
 - Power on the standby compute node.
 - Preassign MAC and WWN addresses, as well as storage boot targets, for up to 256 chassis or 3,584 compute nodes.
- ▶ Create addresses for compute nodes, save the address profiles, and deploy the addresses to the slots in the same chassis or in up to 256 different chassis. This task can be performed without any compute nodes installed in the chassis.
- ▶ Create profiles for chassis that have not been installed in the environment by simply associating an IP address to the future chassis.

Ordering information

With the exception of the IBM Fabric Manager Enterprise Key offering, licenses for offerings in this part include Software Subscription and Support (SW S&S). Software Subscription and Support contracts are managed in Passport Advantage.

Table 4-5 lists the part numbers for single chassis entitlement for the type of chassis with the Subscription and Support (maintenance) term listed.

Table 4-5 Ordering part numbers for IBM Fabric Manager

Description	Part number for EMEA and AP	Part number for US, Canada, AP, and Japan
IBM Fabric Manager for BladeCenter Per Managed Chassis with 1 Year SW S&S	00Y6198	00Y6192
IBM Fabric Manager for BladeCenter Per Managed Chassis with 3 Year SW S&S	00Y6199	00Y6193
IBM Fabric Manager for BCS Per Managed Chassis with 1 Year SW S&S	00Y6200	00Y6194
IBM Fabric Manager for BCS Per Managed Chassis with 3 Year SW S&S	00Y6201	00Y6195
IBM Fabric Manager for BladeCenter Upgrade from BOFM Per Managed Chassis with 1 Year SW S&S	00Y6202	00Y6196
IBM Fabric Manager for BCS Upgrade from BOFM Per Managed Chassis with 1 Year SW S&S	00Y6203	00Y6197

Managed Chassis: Managed Chassis is a unit of measure by which the program can be licensed. A *chassis* is a physical enclosure structure that provides services, such as power, cooling, networking, various interconnects, and management, to IT elements that dock inside the enclosure. The licensee must obtain Managed Chassis entitlements for each chassis that is managed by the program.

Table 4-6 lists the part number for an enterprise key through an IBM Fabric Manager server.

Table 4-6 Part number for IBM Fabric Manager Enterprise Key

Description	Part number for EMEA and AP	Part number for US, Canada, AP, and Japan
Enterprise Key		
IBM Fabric Manager Enterprise Key Per Managed Chassis	00Y6213	00Y6212
Software and Subscription for Enterprise Key		
IBM Fabric Manager for Enterprise Key for BladeCenter, Per Managed Chassis with 1 Year SW S&S	00Y6220	00Y6215
IBM Fabric Manager for Enterprise Key Upgrade from BOFM for BladeCenter Per Managed Chassis with 1 Year SW S&S	00Y6222	00Y6217
IBM Fabric Manager for Enterprise Key for BCS Per Managed Chassis with 1 Year SW S&S	00Y6221	00Y6216
IBM Fabric Manager for Enterprise Key Upgrade from BOFM for BCS Per Managed Chassis with 1 Year SW S&S	00Y6223	00Y6218

The IBM Fabric Manager Enterprise Key offering allows purchase of an Enterprise key that can be distributed across all chassis in the datacenter through an IFM server. The offerings associated with Software and Subscription provide the software maintenance, and must be ordered by type and quantity to match the hardware being managed by the Enterprise key.

To allow for easier order submission, the offerings in this chapter can be configured in conjunction with eligible System x hardware, software, and license-only configurations. The configurations can be created for this product based on existing supported System x hardware system units in a specific business environment.

With the exception of the IBM Fabric Manager Enterprise Key offering, licenses for offerings in this chapter include one year or three years of Software Subscription and Support. Software Subscription and Support renewals for all features are aligned through Passport Advantage Express or Passport Advantage, reducing the administrative burden of keeping your Subscription and Support up-to-date.

If a new release of a specific offering in this chapter is made generally available in Passport Advantage after the announcement of the offering in the System x ordering systems, it is possible to elect to take delivery of any release generally available for that offering during the product registration process.

Eligible software products, offered and selected during the ordering process with each hardware system, are enabled for electronic download for the system unit for which the product is ordered. Optionally, software and license-only orders can be created for these products based on existing supported hardware system units in the business environment. Offerings included in this chapter are available exclusively by electronic download.

With the exception of the IBM Fabric Manager Enterprise Key offering, orders that are placed for the eligible offerings in this announcement are routed and fulfilled through Passport Advantage Express. Invoicing for the initial license orders occurs through the System x order management systems. Subscription and Support renewals are managed through Passport Advantage.

Passport Advantage allows you to have a common anniversary date for Software Subscription and Support renewals. The anniversary date, established at the order ship date for the initial product acquisition, remains unchanged while the Passport Advantage Agreement remains in effect. New software purchases initially include 12 or 36 full months of Subscription and Support coverage.

Subscription and Support in the second year (the first year of renewal) can be prorated to be coterminous with the common anniversary date but managed and invoiced through Passport Advantage. Thereafter, all Software Subscription and Support renewals are on the common anniversary date and include 12 full months of Subscription and Support.

Packaging

Important: For software products in this chapter, physical media and hardcopy documentation are not included in the deliverable package. The software code is delivered only through electronic download.

After the receipt of an order, the IBM Business Partner or client is required to register and accept Passport Advantage or Passport Express terms and conditions. Upon registration, the client is provided with a website for the electronic download. The product includes the license and one year or three years of Software Subscription and Support, unless otherwise noted.

The Proof of Entitlement for these programs is a copy of a paid sales receipt, purchase order, invoice, or other sales record from IBM or its authorized reseller from whom the program was acquired, provided that it states the license charge unit (the characteristics of intended use of the program; for example, number of processors and number of users) and quantity acquired.

The program, when downloaded from a website, contains the applicable IBM license agreement and license information, if appropriate, and will be presented for acceptance at the time of installation of the program. For future reference, the license and license information are stored in a directory, such as LICENSE.TXT.

System requirements

IBM Fabric Manager can be installed on the IBM System x and BladeCenter servers that are listed on the IBM ServerProven® website:

<http://www-03.ibm.com/servers/eserver/serverproven/compat/us/>

IBM System x and BladeCenter servers must be Intel 32-bit or 64-bit systems with the supported operating systems:

- ▶ Microsoft Windows Server 2003 (SP1, SP2) (32-bit or 64-bit)
- ▶ Windows Server 2008 (SP1, SP2) (32-bit or 64-bit)
- ▶ Windows Server 2008 R2 (64-bit)
- ▶ Red Hat Enterprise Linux 4 (up to SP8) (32-bit or 64-bit)
- ▶ Red Hat Enterprise Linux 5 (up to SP5)
- ▶ SUSE Linux Enterprise Server 9 (up to SP4)
- ▶ SUSE Linux Enterprise Server 10 (up to SP2)
- ▶ SUSE Linux Enterprise Server 11

Important: A minimum of 1 GB RAM and 300 MB free disk space during installation and Sun Java™ SE 1.6, or later, are required.

The program specifications and specified operating environment information are in documentation accompanying the program, if available, such as a README file, or other information published by IBM, such as an announcement letter. Documentation and other program content might be supplied only in the English language.

Hardware support

Table 4-7 lists the chassis, servers, and options that support the IBM Fabric Manager offerings.

Table 4-7 IBM Fabric Manager support matrix: chassis, blade servers, and expansion cards

Description (Yes = supported, No = not supported)	Machine type or part number	Feature code	IBM Fabric Manager V4.0
Chassis			
IBM BladeCenter E	8677	N/A	Yes
IBM BladeCenter H	8852	N/A	Yes
IBM BladeCenter S	8886	N/A	Yes
IBM BladeCenter T	8720, 8730	N/A	Yes
IBM BladeCenter HT	8740, 8750	N/A	Yes
Blade servers			
IBM BladeCenter HS12	8028	N/A	Yes
IBM BladeCenter HS22	7870	N/A	Yes
IBM BladeCenter HS22V	7871	N/A	Yes
IBM BladeCenter HS23	7875, E5-2600	N/A	Yes

Description (Yes = supported, No = not supported)	Machine type or part number	Feture code	IBM Fabric Manager V4.0
IBM BladeCenter HS23	7875, E5-2600 v2	N/A	Yes
IBM BladeCenter HS23E	8038	N/A	Yes
IBM BladeCenter HX5	7872	N/A	Yes
IBM BladeCenter HX5	7873	N/A	Yes
IBM BladeCenter PS700/701/702	8406	N/A	Yes
IBM BladeCenter PS703/704	7891	N/A	Yes
Ethernet expansion cards			
Integrated Gigabit Ethernet	None	None	Yes
10Gb Interposer Card for HS23	94Y8550	A244	Yes
Ethernet Expansion Card (CFFv)	39Y9310	2969	Yes
Ethernet Expansion Card (CIOv)	44W4475	1039	Yes
2/4 Port Ethernet Exp. Card (CFFh)	44W4479	5476	Yes
QLogic 2-port 10Gb CNA (CFFh)	42C1830	3592	Yes
QLogic 2-port 10Gb CNA (CFFh)	00Y3280	A3JB	Yes
QLogic Virtual Fabric Adapter (CFFh)	00Y3332	A4AC	Yes
Qlogic Virtual Fabric CNA (CFFh)	00Y5618	A4AD	Yes
Broadcom 10Gb Gen 2 4-pt. Eth. (CFFh)	46M6164	0098	Yes
Broadcom 10Gb Gen 2 2-pt. Eth. (CFFh)	46M6168	0099	Yes
Broadcom 2-port 10Gb VFA (CFFh)	81Y3133	A1QR	Yes
Emulex Virtual Fabric Adapter (CFFh)	49Y4235	5755	Yes
Emulex Virtual Fabric Adpt. Adv. (CFFh)	49Y4275	2435	Yes
Emulex Virtual Fabric Adapter II (CFFh)	90Y3550	A1XG	Yes
Emulex Virt. Fabric Adpt. Adv. II (CFFh)	90Y3566	AIXH	Yes
Emulex Virtual Fabric Adapter II (CFFh)	00Y3266	A3NV	Yes
Emulex Virt. Fabric Adpt. Adv. II (CFFh)	00Y3264	A3NW	Yes
Emulex 10GbE VFA II for HS23	81Y3120	A287	Yes
Emulex 10GbE VFA Adv. II for HS23	90Y9332	A2ZN	Yes
Intel 10Gb 2-port Ethernet (CFFh)	42C1810	3593	No
Brocade 2-port 10Gb CNA (CFFh)	81Y1650	5437	No
Fibre Channel expansion cards			
Emulex 8Gb FC (CIOv)	46M6140	3598	Yes
QLogic 4Gb FC (CFFv)	41Y8527	2970	Yes
QLogic 4Gb FC (CIOv)	46M6065	3594	Yes

Description (Yes = supported, No = not supported)	Machine type or part number	Feature code	IBM Fabric Manager V4.0
QLogic 8Gb FC (CIOv)	44X1945	1462	Yes
QLogic Ethernet and 8Gb FC (CFFh)	44X1940	5485	Yes
QLogic Ethernet and 8Gb FC (CFFh)	00Y3270	A3JC	Yes
SAS expansion cards			
SAS Expansion Card (CFFv)	39Y9190	2979	Yes
SAS Expansion Card (CFFv)	44E5688	A3J9	Yes
SAS Connectivity Card (CIOv)	43W4068	1041	Yes
ServeRAID MR10ie (CIOv)	46C7167	5752	No
ServeRAID H1135 (CIOv)	90Y4750	A1XJ	No

4.9 BladeCenter Virtual Fabric vNIC solution

Currently, deployment of server virtualization technologies in data centers requires significant effort to provide sufficient network I/O bandwidth to satisfy the demands of virtualized applications and services. For example, every virtualized system can host several dozen network applications and services, and each of these services requires certain bandwidth (or speed) to function properly. Furthermore, because of different network traffic patterns relevant to different service types, these traffic flows might interfere with each other, leading to serious network problems, including the inability of the service to perform its functions. This inability becomes important when I/O disk storage data traffic uses the same physical infrastructure (for example, iSCSI or FCoE).

Providing sufficient bandwidth and isolation to virtualized applications might be particularly challenging for blade-based deployments where the number of I/O ports per blade is limited. For example, for the maximum of eight ports per single-wide blade, at least two ports are used to connect to the external storage. Therefore, only six ports are available with a total network bandwidth of 6 Gb per blade for Gigabit Ethernet infrastructures, leaving no room for future growth. In addition, the traffic flows are isolated on a physical port basis with a minimum bandwidth of 1 Gb per flow, thus limiting the flexibility of bandwidth usage.

The IBM BladeCenter Virtual Fabric virtual network interface controller (vNIC) solution addresses the issues by increasing the number of available Ethernet ports and providing more flexibility in allocating the available bandwidth to meet specific application requirements.

The BladeCenter Virtual Fabric vNIC solution is based on the IBM BladeCenter H or HT chassis with a 10 Gb Converged Enhanced Ethernet (CEE) infrastructure and 10 Gb Virtual Fabric Adapters (VFAs) installed in each blade server.

The concept of vNIC is based on the ability to split one physical 10 Gb Ethernet interface into up to four logical subinterfaces called vNICs, where every vNIC is characterized by a certain amount of bandwidth allocated to it. Each vNIC is represented as a separate network adapter to the host operating system. The bandwidth of a single vNIC can be from 100 Mb to 10 Gb with 100 Mb increments, and the combined bandwidth of all vNICs on a same port cannot exceed 10 Gb. With dual-port 10 Gb VFAs, the maximum number of vNICs is eight, and the total bandwidth for all vNICs cannot exceed 20 Gb.

Two operational modes are available, depending on the type of 10 Gb expansion card installed in the server and the type of 10 Gb switch installed in the BladeCenter chassis:

- **IBM Virtual Fabric Mode**

In this mode, the 10 Gb switch is vNIC-aware, that is, the configuration of vNICs is done on a switch, then it propagates vNIC parameters to VFA using the DataCenter Bridging eXchange (DCBX) protocol. vNIC bandwidth allocation and metering is performed by both the switch and the VFA. In such a case, a bidirectional virtual channel of an assigned bandwidth is established between them for every defined vNIC. This mode is supported only by IBM Virtual Fabric 10 Gb Switch Module together with Emulex Virtual Fabric Adapters or Emulex Virtual Fabric Adapters II.

- **Switch Independent Mode**

In this mode, the 10 Gb switch is not aware of vNICs, that is, the configuration of vNICs is done on a VFA, and vNIC bandwidth allocation and metering is only performed by VFA itself. In such a case, a unidirectional virtual channel is established where the bandwidth management is only performed for the outgoing traffic on a VFA side (server-to-switch). The incoming traffic (switch-to-server) uses the all available physical port bandwidth, as there is no metering performed on either the VFA or a switch side. This mode is supported by different switches and pass-through modules together with Broadcom Virtual Fabric Adapters or Emulex Virtual Fabric Adapters II.

With either mode, the integrated Gigabit Ethernet ports on the blade servers are still available, connecting to 1 Gb Ethernet switches installed in bays 1 and 2 of the chassis. Adding another CIOv expansion card in the blade servers and I/O modules in bays 3 and 4 of the BladeCenter H or HT chassis makes additional network or storage ports available to applications to provide additional expansion capabilities.

The combination of the integrated Gigabit Ethernet ports, the CIOv expansion card, and the eight vNICs of the Virtual Fabric Adapter means that up to 12 I/O ports are available to each single-wide blade in the chassis.

In addition, Emulex Virtual Fabric Adapter Advanced II supports FCoE and iSCSI hardware initiator functionality on both physical ports that provide external storage connectivity. When the ports operate in pNIC mode, the Advanced adapter or the Advanced upgrade operates in traditional Converged Network Adapter (CNA) mode with two Ethernet ports and two storage ports (iSCSI or FCoE) available to the operating system. When the ports are configured in vNIC mode and storage protocols are enabled on the Emulex 10 GbE Virtual Fabric Adapter Advanced II, six ports are Ethernet, and two ports are either iSCSI or FCoE. Unlike Emulex Virtual Fabric Adapter II (part number 90Y3566), the Emulex Virtual Fabric Adapter Advanced (part number 49Y4275) supports only iSCSI in either pNIC or vNIC mode.

4.9.1 IBM Virtual Fabric Mode

Figure 4-8 provides a general view of IBM Virtual Fabric Mode solution architecture and offers an example of configuring the vNICs to have different bandwidth values. The different colors represent different traffic flows for different services / applications, and the thickness of links represent the amount of assigned bandwidth for particular traffic flow. The Virtual Fabric architecture provides the capability to subdivide a single 10 Gb link into up to four virtual channels.

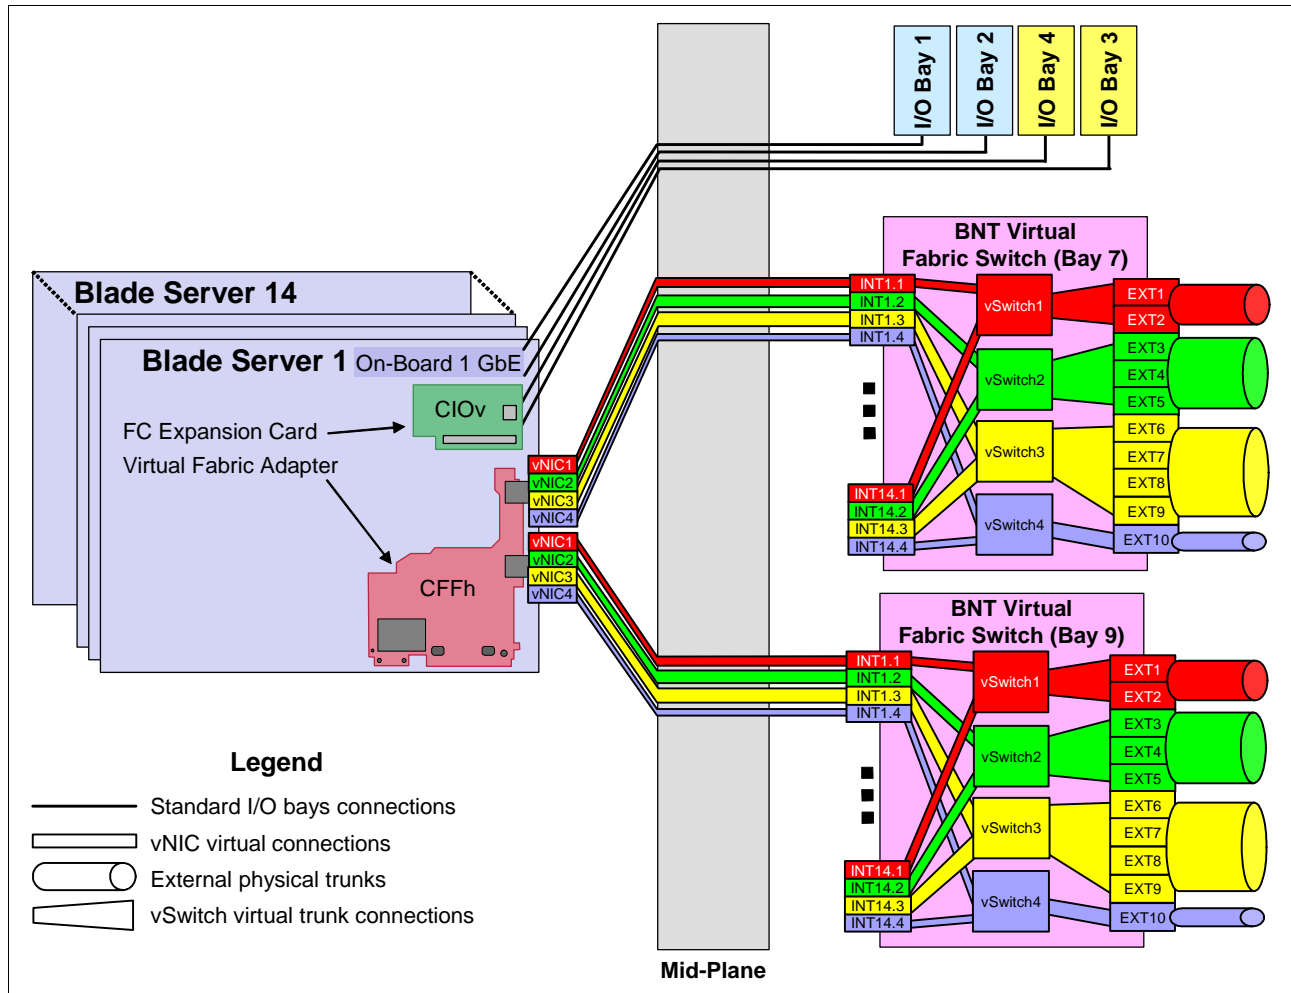


Figure 4-6 BladeCenter Virtual Fabric mode solution architecture

Configuration is provided through the switch user interfaces. Parameters are passed to the adapter by the switch. Channel bandwidth is dynamic in that it may be changed (from the switch) without resetting the VFA.

The VFA enforces bandwidth limits for each virtual channel. The VFS enforces the same virtual channel bandwidth limits on traffic that flows to the server.

Bandwidth allocation and metering is controlled by the VFS, and in case incoming or outgoing traffic flow for a particular vNIC exceeds established bandwidth limits, the switch begins to drop packets until the traffic falls below configured thresholds.

To isolate vNICs traffic flows from each other, vNICs can be assigned to vNIC logical groups (or vSwitches) that behave as independent switches inside VFS. vNICs in the same group can communicate with each other while vNICs in different groups cannot. This situation is enforced by assigning a unique VLAN ID to every vSwitch, and that VLAN ID becomes the VLAN number for all vNICs belonging to the vSwitch. In addition, external or internal switch ports and external trunks can also be assigned to a particular vNIC group to isolate upstream connections.

The IBM Virtual Fabric 10 Gb Switch Module disables several Layer 2 and Layer 3 switch functions when supporting the IBM Virtual Fabric Mode. For example, ACL filters, VMready, Layer 2 spanning tree, and Layer 3 routing are some of the functions that are not supported in

conjunction with Virtual Fabric Mode operation. These functions are disabled at the port level, not the entire switch. For example, you cannot configure vNIC and VMready on the same switch port (blade server), but you can configure them on different ports on the same switch module.

Trunk failover is supported on a per-vNIC basis, so you can set up NIC teaming in a failover mode for vNICs. If there is a external trunk failure, the switch shuts down only those vNICs that belong to the same vNIC group as the failed external trunk belongs to. Other vNICs on the same pNIC port are untouched and continue to perform their functions.

There are two uplink modes of operation for vNICs:

► **Dedicated uplink**

In this mode, an external port or external trunk uplink can belong only to a single vNIC group.

► **Shared uplink**

In this mode, an external port or external trunk uplink can belong to multiple vNIC groups.

The vNICs are configured on IBM Virtual Fabric Switch by dividing internal VFS ports into subinterfaces (or logical ports). Each internal interface of VFS can be divided into up to four vNICs, and they have strict one-to-one mapping to vNICs on VFA.

Table 4-8 lists vNIC mappings for VFA and VFS and vNIC identification data to differentiate vNICs in the host operating system (vNIC PCIe Function ID).

Table 4-8 vNIC mappings and identification

vNIC PCIe Function ID	VFA port number	VFS bay number	vNIC number	vNIC ID ^a
0	0	7	1	INTx.1
2	0	7	2	INTx.2
4	0	7	3	INTx.3
6	0	7	4	INTx.4
1	1	9	1	INTx.1
3	1	9	2	INTx.2
5	1	9	3	INTx.3
7	1	9	4	INTx.4

a. The letter x in a vNIC ID represents the VFS internal port number to which the VFA physical port is connected. (The x is in the range of 1 - 14.)

Consider several factors when you plan a Virtual Fabric vNIC solution:

- The Virtual Fabric Adapter must operate in vNIC mode.
- Up to four vNICs can be configured on each pNIC on the VFA (a total of eight vNICs on a 2-port VFA).
- Up to two vNICs can be used as hardware iSCSI initiators or FCoE ports when the advanced feature license is activated on Virtual Fabric Adapter II (either by purchasing Advanced Upgrade, 49Y4265, for 2-port VFA II, 90Y3550, or by purchasing VFA Advanced II, 90Y3566).
- The bandwidth for vNIC can be assigned in the range from 100 Mb - 10 Gb with 100 Mb increments (the default is 2.5 Gb.)

- ▶ The sum of bandwidth allocated for all four vNICs on the same pNIC cannot exceed 10 Gb.
- ▶ Up to 32 vNIC groups are supported on VFS.
- ▶ One vNIC can belong to only one vNIC group.
- ▶ All vNICs on the same VFA physical port must belong to different vNIC groups.
- ▶ Only one external trunk or one external port is allowed per one vNIC group.
- ▶ One internal port can belong only to one vNIC group.
- ▶ In dedicated mode, one external trunk or external port can belong only to a single vNIC group.
- ▶ In shared mode, one external trunk or external port can belong to multiple vNIC groups.
- ▶ Trunk failover is supported on a per-vNIC basis.
- ▶ Virtual machines (VMs) and other virtual entities (VEs) associated with vNICs are automatically detected by VFS when VMready is enabled. However, these VEs cannot be assigned to any of the VMready VM groups (because vNIC groups used by vNICs are fully isolated from other switch entities).
- ▶ vNICs are supported with IBM Fabric Manager.

Table 4-9 lists the components supported by BladeCenter Virtual Fabric Mode solutions.

Table 4-9 BladeCenter Virtual Fabric Mode solutions

Description	Machine type/part number	Feature	Quantity
Chassis			
IBM BladeCenter H	8852	None	Varies
IBM BladeCenter HT	8740/8750	None	Varies
Blade servers			
HS22	7870	None	Varies
HS22V	7871	None	Varies
HS23 (E5-2600)	7875	None	Varies
HS23 (E5-2600 v2)	7875	None	Varies
HX5	7872	None	Varies
HX5	7873	None	Varies
Virtual Fabric Adapters			
Emulex Virtual Fabric Adapter II (CFFh)	90Y3550	A1XG	1 per blade server
Emulex Virtual Fabric Adapter Advanced II (CFFh)	90Y3556	A1XH	1 per blade server
Emulex 10 GbE Virtual Fabric Adapter II (CFFh)	00Y3266	A3NV	1 per blade server
Emulex 10 GbE Virtual Fabric Adapter Advanced II (CFFh)	00Y3264	A3NW	1 per blade server
Emulex 10GbE VFA II (CFFh) for HS23	81Y3120	A287	1 per HS23 blade
Emulex 10GbE VFA Advanced II (CFFh) for HS23	90Y9332	A2ZN	1 per HS23 blade

Description	Machine type/part number	Feature	Quantity
Virtual Fabric Switches			
IBM Virtual Fabric 10 Gb Switch Module	46C7191	1639	1 or 2 per chassis
External Virtual Fabric Switch connections			
Copper-based			
1 m IBM Passive DAC SFP+ cable	90Y9427	A1PH	1 per external port
3 m IBM Passive DAC SFP+ cable	90Y9430	A1PJ	1 per external port
5 m IBM Passive DAC SFP+ cable	90Y9433	A1PK	1 per external port
8.5 m IBM Passive DAC SFP+ cable	90Y9436	A1PL	1 per external port
IBM SFP RJ45 Transceiver	81Y1618	3268	1 per external port
Fiber Optics-based			
IBM 10GBase-SR 10 GbE 850 nm Fiber SFP+ Transceiver	44W4408	4942	1 per external port
IBM SFP+ SR Transceiver	46C3447	5053	1 per external port
IBM SFP+ LR Transceiver	90Y9412	A1PM	1 per external port
IBM SFP+ ER Transceiver	90Y9415	A1PP	1 per external port

4.9.2 Switch Independent Mode

Switch Independent Mode architecture has the following key differences from IBM Virtual Fabric mode:

- ▶ A physical port on a switch is not divided into subinterfaces. The switch treats the traffic from the specific vNIC in the same way as a traffic from any other regular host port, that is, identifying it by VLAN ID and MAC address.
- ▶ The traditional VLAN approach is used to isolate the traffic from different vNICs rather than the vSwitch concept.
- ▶ Bandwidth metering is only performed by VFA itself for server-to-switch traffic. Switch-to-server traffic is not metered.

The vNICs are configured in the operating system by using the Emulex or Broadcom software in the pre-boot environment. The VFA configuration and setup is performed inside the network controller rather than the switch.

You must take several considerations into account when planning to use the components in Switch Independent Mode:

- ▶ The Virtual Fabric adapter must operate in Switch Independent mode (Emulex) or in NIC Partition mode (Broadcom).
- ▶ Up to four vNICs can be configured on each physical NIC (pNIC) on the VFA (a total of eight vNICs on a two-port VFA).
- ▶ The bandwidth for each vNIC can be assigned in a range of 100 Mb - 10 Gb in 100 Mb increments (the default is 2.5 Gb)
- ▶ The total bandwidth of the four vNICs per port cannot exceed 10 Gbps. The Broadcom 2-port 10Gb Virtual Fabric Adapter supports oversubscription, where the aggregate bandwidth can be as high as 4x 10 Gb. This oversubscription allows any of the vNICs to

consume as much as 10 Gb where there is no contention or less when there is contention (based on the weight assigned to that vNIC).

- ▶ The sum of the bandwidth allocated for all four vNICs on the same pNIC cannot exceed 10 Gb.
- ▶ Up to two vNICs can be used as hardware iSCSI initiators or FCoE ports when the advanced feature license is activated on Virtual Fabric Adapter II (either by purchasing Advanced Upgrade, 49Y4265, for 2-port VFA II, 90Y3550, or by purchasing VFA Advanced II, 90Y3566).
- ▶ vNICs are supported IBM Fabric Manager.

Table 4-9 on page 443 lists the components supported by BladeCenter Switch Independent Mode solution.

Table 4-10 BladeCenter Switch Independent Mode solutions

Description	Machine type or part number	Feature	Quantity
Chassis			
IBM BladeCenter H	8852	None	Varies
IBM BladeCenter HT	8740/8750	None	Varies
Blade servers			
HS22	7870	None	Varies
HS22V	7871	None	Varies
HS23	7875, E5-2600	None	Varies
HS23	7875, E5-2600 v2	None	Varies
HX5	7872	None	Varies
HX5	7873	None	Varies
Virtual Fabric Adapters			
Emulex Virtual Fabric Adapter II (CFFh)	90Y3550	A1XG	1 per blade server
Emulex Virtual Fabric Adapter Advanced II (CFFh)	90Y3566	ALXH	1 per blade server
Emulex 10 GbE Virtual Fabric Adapter II (CFFh)	00Y3266	A3NV	1 per blade server
Emulex 10 GbE Virtual Fabric Adapter Advanced II (CFFh)	00Y3264	A3NW	1 per blade server
Emulex 10 GbE VFA II (CFFh) for HS23	81Y3120	A287	1 per HS23 blade
Emulex 10 GbE VFA Advanced II (CFFh) for HS23	90Y9332	A2ZN	1 per HS23 blade
QLogic 10Gb Virtual Fabric Adapter (CFFh)	00Y3332	A4AC	1 per blade server
QLogic 10 Gb Virtual Fabric CNA (CFFh)	00Y5618	A4AD	1 per blade server
Broadcom 2-port 10 Gb Virtual Fabric Adapter (CFFh)	81Y3133	2436	1 per blade server
Virtual Fabric I/O Modules			
IBM Virtual Fabric 10 Gb Switch Module	46C7191	1639	1 or 2 per chassis
Cisco Nexus 4001I Switch Module	46M6071	0072	1 or 2 per chassis
10 Gb Ethernet Pass-thru Module	46M6181	1641	1 or 2 per chassis

Description	Machine type or part number	Feature	Quantity
External Virtual Fabric I/O Module connections			
IBM Virtual Fabric 10G Switch Module			
1 m IBM Passive DAC SFP+ cable	90Y9427	A1PH	1 per external port
3 m IBM Passive DAC SFP+ cable	90Y9430	A1PJ	1 per external port
5 m IBM Passive DAC SFP+ cable	90Y9433	A1PK	1 per external port
8.5 m IBM Passive DAC SFP+ cable	90Y9436	A1PL	1 per external port
IBM SFP RJ45 Transceiver	81Y1618	3268	1 per external port
IBM 10GBase-SR 10 GbE 850 nm Fiber SFP+ Transceiver	44W4408	4942	1 per external port
IBM SFP+ SR Transceiver	46C3447	5053	1 per external port
IBM SFP+ LR Transceiver	90Y9412	A1PM	1 per external port
IBM SFP+ ER Transceiver	90Y9415	A1PP	1 per external port
Cisco Nexus 4001I Switch Module			
Cisco 10BASE-SR SFP+ Transceiver	88Y6054	A1A6	1 per external port
Cisco 1000BASE-T SFP Transceiver	88Y6058	A1A7	1 per external port
Cisco 1000BASE-SX SFP Transceiver	88Y6062	A1A8	1 per external port
Cisco 10BASE-LR SFP+ Transceiver	SFP-10G-LR ^a		1 per external port
Cisco 10BASE-CU SFP+ Cable 1 Meter	SFP-H10GB-CU1M(=) ^a		1 per external port
Cisco 10BASE-CU SFP+ Cable 3 Meter	SFP-H10GB-CU3M(=) ^a		1 per external port
Cisco 10BASE-CU SFP+ Cable 5 Meter	SFP-H10GB-CU5M(=) ^a		1 per external port
Cisco GE SFP, LC connector LX/LH transceiver	GLC-LH-SM(=) ^a		1 per external port
10 Gb Ethernet Pass-thru Module			
IBM 10GBase-SR 10 GbE 850 nm Fiber SFP+ Transceiver	44W4408	4942	1 per external port
Twinax Active 1 m	45W2398	2711	1 per external port
Twinax Active 3 m	45W2408	2731	1 per external port
Cisco 10BASE-CU SFP+ Cable 1 Meter	SFP-H10GB-CU1M(=) ^a		1 per external port
Cisco 10BASE-CU SFP+ Cable 3 Meter	SFP-H10GB-CU3M(=) ^a		1 per external port
Cisco 10BASE-CU SFP+ Cable 5 Meter	SFP-H10GB-CU5M(=) ^a		1 per external port

a. The option is available from Cisco.

For more information about the vNIC solution, see the following documentation:

- *Virtual Fabric Adapter Installation and User's Guide:*
<http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5082945>
- *IBM BladeCenter vNIC Tutorial:*
<http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5080917>

- ▶ *IBM Virtual Fabric 10 Gb Switch Module Application Guide:*
<http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5080917>
- ▶ *IBM BladeCenter Interoperability Guide:*
<http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5073016>
- ▶ *IBM BladeCenter Virtual Fabric Solutions, SG24-7966:*
<http://www.redbooks.ibm.com/abstracts/sg247966.html>

5



Storage integration

IBM BladeCenter offers several possibilities for integrating into a storage infrastructure, such as SAS, Fibre Channel, and 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 so on.

This chapter includes the following topics:

- ▶ 5.1, “External storage for IBM BladeCenter” on page 450
- ▶ 5.2, “Fibre Channel storage system interoperability” on page 451
- ▶ 5.3, “iSCSI system storage interoperability” on page 457
- ▶ 5.4, “SAS storage system interoperability” on page 461
- ▶ 5.5, “High availability and redundancy” on page 463
- ▶ 5.6, “Performance” on page 464
- ▶ 5.7, “Backup solutions” on page 464
- ▶ 5.8, “SAS tape storage for IBM BladeCenter” on page 468
- ▶ 5.9, “External storage expansion (JBOD) attachments” on page 471
- ▶ 5.10, “IBM BladeCenter S integrated storage” on page 475
- ▶ 5.11, “Converged networks” on page 484

5.1 External storage for IBM BladeCenter

Several options are available for attaching external storage systems to IBM BladeCenter:

- ▶ Direct-attached SAS storage
- ▶ Storage area networks (SANs) based on iSCSI
- ▶ SANs based on Fibre Channel technologies
- ▶ Converged Networks based on 10 Gb Converged Enhanced Ethernet

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 is 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 provide all the benefits of centralized shared storage in terms of storage consolidation and adequate levels of performance, but use traditional IP address-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 on, and storage servers with an iSCSI interface, such as IBM System Storage DS3500 or IBM N series.

SAS is a technology that is used primarily for server direct-attached storage. However, it can also be used for external shared storage attachments, such as the IBM System Storage DS3500 series.

Converged Networks can carry both SAN and LAN types of traffic over the same physical infrastructure. Such consolidation decreases costs and increases efficiency in building, maintaining, operating, and managing the networking infrastructure.

iSCSI, SAS, FC-based SANs, and Converged Networks can be used for diskless blade solutions to provide greater levels of utilization, availability, and cost effectiveness.

For this book, we limit the description to the entry-level DS3000 series storage systems and entry-level tape drives, autoloaders, and libraries.

Compatibility: Compatibility tables for storage systems and tape drives are based on the latest firmware levels available for that system at the time of the writing of this book.

5.2 Fibre Channel storage system interoperability

In general, if IBM BladeCenter is integrated into FC storage fabric, ensure that the following requirements are met. Check the compatibility guides from your storage system vendor for confirmation.

- ▶ The 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).
- ▶ Clustering software is supported by the storage system (in case you plan to implement clustering technologies).

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.

If IBM BladeCenter is attached to a non IBM storage system, then support for that configuration is provided by the storage system's vendor. If non IBM storage is listed on IBM ServerProven, its configuration has been tested, however IBM does not provide support for it; support is provided by a third-party vendor. See the vendor's compatibility information for supported configurations.

For IBM System Storage compatibility information, see the IBM System Storage Interoperability Center (SSIC) website:

<http://www.ibm.com/systems/support/storage/config/ssic>

5.2.1 Fibre Channel storage support for x86 blade servers

Table 5-1 lists Fibre Channel storage support for x86 blade servers.

Table 5-1 Fibre Channel support: x86 blade servers

FC adapter	BladeCenter I/O module	External SAN fabric ^a	Operating systems	FC storage targets
<ul style="list-style-type: none"> ▶ Emulex 8Gb FC (CIOv), 46M6140 ▶ QLogic 4Gb FC (CFFv), 41Y8527 ▶ QLogic 4Gb FC (CIOv), 46M6065 ▶ QLogic 8Gb FC (CIOv), 44X1945 ▶ QLogic Ethernet and 8Gb FC (CFFh), 44X1940 ▶ QLogic Ethernet and 8Gb FC (CFFh), 00Y3270 	<ul style="list-style-type: none"> ▶ Brocade Enterprise 20-port 8Gb SAN Switch Module, 42C1828 ▶ Brocade 20-port 8Gb SAN Switch Module, 44X1920 ▶ Brocade 10-port 8Gb SAN Switch Module, 44X1921 ▶ QLogic 20-Port 8Gb SAN Switch Module, 44X1905 ▶ QLogic 20-Port 4/8Gb SAN Switch Module, 88Y6406 ▶ QLogic 8Gb Intelligent Pass-thru Module, 44X1907 ▶ QLogic 4/8Gb Intelligent Pass-thru Module, 88Y6410 	<ul style="list-style-type: none"> ▶ IBM b-type ▶ Brocade ▶ Cisco MDS ▶ McDATA ▶ QLogic 	<ul style="list-style-type: none"> ▶ RHEL 5 ▶ RHEL 6 ▶ SLES 10 ▶ SLES 11 ▶ vSphere 4.1 ▶ vSphere 5.0 ▶ vSphere 5.1 ▶ Microsoft Windows Server 2003 / 2003 R2 ▶ Microsoft Windows Server 2008 / 2008 R2 ▶ Microsoft Windows 2012 	<ul style="list-style-type: none"> ▶ DCS3700 ▶ DS3000 ▶ DS5000 ▶ DS8000® ▶ SVC ▶ V3500 ▶ V3700 ▶ V7000 ▶ XIV® ▶ Tape
	<ul style="list-style-type: none"> ▶ Cisco 4Gb 20 port FC Switch Module, 39Y9280 ▶ Cisco 4Gb 20 port FC Switch Module, 44E5696 ▶ Cisco 4Gb 10 port FC Switch Module, 39Y9284 ▶ Cisco 4Gb 10 port FC Switch Module, 44E5692 	<ul style="list-style-type: none"> ▶ IBM b-type ▶ Cisco MDS 		

a. External SAN is a mandatory requirement when BladeCenter I/O module operates as a NPIV gateway (if NPIV mode is supported by this module). If an I/O module operates in Full Fabric mode, it supports direct connections to the Fibre Channel storage targets.

5.2.2 Fibre Channel storage support for Power Systems blade servers

Table 5-2 lists Fibre Channel storage support for Power Systems blade servers.

Table 5-2 Fibre Channel support: Power Systems blade servers

FC expansion card	BladeCenter I/O module	External SAN fabric	Operating systems	FC storage targets
<ul style="list-style-type: none"> Emulex 8Gb FC (CIOv), FC 8240 QLogic 4Gb FC (CIOv), FC 8241 QLogic 8Gb FC (CIOv), FC 8242 QLogic Ethernet and 8Gb FC (CFFh), FC 8271 	<ul style="list-style-type: none"> Brocade Enterprise 20-port 8Gb SAN Switch Module, 42C1828 Brocade 20-port 8Gb SAN Switch Module, FC 5869 Brocade 10-port 8Gb SAN Switch Module, FC 5045 QLogic 20-Port 8Gb SAN Switch Module, FC 3284 QLogic 20-Port 4/8Gb SAN Switch Module, 88Y6406 QLogic 8Gb Intelligent Pass-thru Module, FC 5449 QLogic 4/8 b Intelligent Pass-thru Module, 88Y6410 	<ul style="list-style-type: none"> IBM b-type Brocade Cisco MDS 	<ul style="list-style-type: none"> AIX 5.3 AIX 6.1 AIX 7.1 RHEL 4 RHEL 5 RHEL 6 SLES 10 SLES 11 VIOS 2.1 VIOS 2.2 	<ul style="list-style-type: none"> DCS3700 DS3000 DS5000 DS8000 SVC V3500 V3700 V7000 XIV Tape
	<ul style="list-style-type: none"> Cisco 4Gb 20 port FC Switch Module, FC 3242 Cisco 4Gb 10 port FC Switch Module, FC 3241 	<ul style="list-style-type: none"> IBM b-type Cisco MDS 		

For additional IBM System Storage compatibility information, see the IBM System Storage Interoperability Center:

<http://www.ibm.com/systems/support/storage/config/ssic>

5.2.3 Fibre Channel switch selection, fabric interoperability rules, NPIV

IBM BladeCenter provides integrated FC switching functionality by using several switch options (see 2.9, "Fibre Channel switch modules" on page 153):

- Brocade Enterprise 20-port, 20-port, and 10-port 8 Gb SAN Switch Modules (42C1828, 44X1920, and 44X1921)
- Cisco 10 and 20 port 4 Gb Fibre Channel Switch Modules (39Y9280 and 39Y9284)
- QLogic 20-Port 8 Gb SAN Switch Module (44X1905) and QLogic 20-Port 4/8 Gb SAN Switch Module (88Y6406)

Brocade Access Gateway functionality is integrated into the Brocade 8 Gb SAN switch modules and it enables certain NPIV (N_Port ID Virtualization) processing on it to provide intelligent pass-through capabilities for the BladeCenter chassis.

Access gateway mode: Access gateway mode is only available when 20 ports are active. If you have a 10-port switch, to enable access gateway mode, you must order and apply the 10-port Ports on Demand feature.

QLogic 8 Gb Intelligent Pass-thru Module (44X1907) and QLogic 4/8 Gb Intelligent Pass-thru Module (88Y6410) are embedded solutions that provide seamless integration of BladeCenter into an existing Fibre Channel fabric, avoiding any multivendor interoperability issues by using N_Port ID Virtualization (NPIV) technology. This module adds ports to the fabrics (not domain IDs, as switches do), and it is not apparent to the switches in the fabric.

Here is a brief overview that describes how the NPIV concept works for the Intelligent Pass-thru Module and the Brocade Access Gateway. In general, several basic types of ports are used in Fibre Channel fabrics, including these examples:

- ▶ N_Ports (node ports) represent an endpoint 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 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 configuration 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 concern with E_Ports is an interoperability issue between switches from different vendors. In many cases, only the “interoperability mode” can be used in these fabrics, thus disabling most of the vendor’s advanced features.

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. You can separate traffic between VMs and 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. 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 also; check the documentation that includes the FC switch.

The QLogic Intelligent Pass-thru Modules and the Brocade Access Gateway use the NPIV technique to present the blade’s port IDs as N_Ports to the external fabric switches, 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.

External switches connected to the IPM do not see IPM at all; they see only N_ports connected to the F_ports. So you can use IPM to achieve a higher port count for better scalability without using Domain IDs, and avoid multivendor interoperability issues. However, these intelligent pass-through modules cannot be directly attached to the storage system. They must be attached to an external NPV-capable FC switch. See the switch's documentation about NPV 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 (Table 5-3). There are no strict rules to follow during integration planning. However, several considerations must be taken into account.

Table 5-3 SAN module feature comparison and interoperability

	Brocade SAN SM	QLogic FCSM	Cisco FCSM	QLogic IPM	Brocade Access Gateway
Basic FC connectivity					
FC-SW-2 interoperability	Yes	Yes	Yes	Not applicable	Not applicable
Zoning	Yes	Yes	Yes	Not applicable	Not applicable
Maximum number of Domain IDs	239	239	239	Not applicable	Not applicable
Advanced FC connectivity					
VSAN	No	No	Yes	Not applicable	Not applicable
Port Aggregation	Yes	No	Yes	Not applicable	Not applicable
Advanced fabric security	Yes	Yes	Yes	Not applicable	Not applicable
Interoperability (existing fabric)					
Brocade interoperability	Yes	No	No	Yes	Yes
QLogic interoperability	Yes	Yes	Yes	Yes	Yes
Cisco interoperability	No	No	Yes	Yes	Yes

Almost all switches support interoperability standards, which means that almost any switch can be integrated into an existing fabric using interoperability mode. *Interoperability mode* is a special mode used to integrate FC fabrics of various vendors into one fabric, 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 5-3 provides a high-level overview of the standard and advanced functions that are available for particular BladeCenter SAN switches. It lists how these switches might be used for designing new storage networks or integrating with existing storage networks.

Important: 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.

Note the following information about the Interoperability section of Table 5-3 on page 455:

- ▶ The word “Yes” indicates that a feature is supported without any restrictions for existing fabric, but with restrictions for added fabric, and vice versa.
- ▶ The word “No” does not necessarily mean that a feature is not supported. Instead, it means that severe restrictions apply to the existing fabric; some functions of the existing fabric potentially must be disabled (if used).

For example, if we integrate Brocade FCSM into a QLogic fabric, then we are not able to use Brocade proprietary features such as ISL trunking; however, a QLogic fabric does not lose functionality. Conversely, if we integrate a QLogic fabric into an 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 configuration is possible with interoperability mode only, so extended functionality is not supported). A better way would be to use the QLogic Intelligent Pass-thru or Brocade Access Gateway.

The rules for switch selection and interoperability can be summarized as follows:

- ▶ QLogic FCSM is used when IBM BladeCenter is integrated into an existing QLogic fabric or when basic FC functionality is required; that is, with one IBM BladeCenter with a direct-connected storage server.
- ▶ Brocade SAN SM is used when IBM BladeCenter is integrated into an 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.

FC storage fabric: Using FC storage fabric from the same vendor avoids possible operational, management, and troubleshooting issues.

5.2.4 NPIV support on Power Systems blade servers

Table 5-4 lists N-Port ID Virtualization (NPIV) support on Power Systems blade servers.

Table 5-4 Fibre Channel NPIV support on Power Systems blade servers

Expansion card		AIX Client			IBM i Client			Linux Client		
Switch module	Part number or feature code	QLogic 8Gb FC (CIOv), FC 8242	QLogic Eth. & 8Gb FC (CFFh), FC 8271	Emulex 8Gb FC (CIOv), FC 8240	QLogic 8Gb FC (CIOv), FC 8242	QLogic Eth. & 8Gb FC (CFFh), FC 8271	Emulex 8Gb FC (CIOv), FC 8240	QLogic 8Gb FC (CIOv), FC 8242	QLogic Eth. & 8Gb FC (CFFh), FC 8271	Emulex 8Gb FC (CIOv), FC 8240
Brocade Enterprise 20-port 8Gb SAN SM	42C1828	Y	Y	Y	Y	Y	Y ^a	Y	Y	Y
Brocade 20-port 8Gb SAN Switch Module	FC 5869	Y	Y	Y	Y	Y	Y ^a	Y	Y	Y
Brocade 10-port 8Gb SAN Switch Module	FC 5045	Y	Y	Y	Y	Y	Y ^a	Y	Y	Y
Cisco 4Gb 20 port FC Switch Module	FC 3242	Y	Y	Y	Y	Y	Y ^a	N	N	N
Cisco 4Gb 10 port FC Switch Module	FC 3241	Y	Y	Y	Y	Y	Y ^a	N	N	N
QLogic 20-Port 8Gb SAN Switch Module	FC 3284	Y	Y	N	Y	Y	N	Y	Y	N
QLogic 20-Port 4/8Gb SAN Switch Module	88Y6406	Y	Y	N	Y	Y	N	Y	Y	N
QLogic 8Gb Intelligent Pass-Thru Module	FC 5449	Y	Y	N	N	N	N	N	N	N
QLogic 4/8Gb Intelligent Pass-Thru Module	88Y6410	Y	Y	N	N	N	N	N	N	N

a. Virtual Tape only with VIOS 2.1.2 or later.

5.3 iSCSI system storage interoperability

iSCSI uses a traditional Ethernet network for block I/O between storage systems and servers. 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.

iSCSI technology is used with BladeCenter to build a low-cost SAN based on a Gigabit Ethernet or 10 Gb Ethernet infrastructure, or to integrate IBM blade servers into an IBM System i® solution. With an IBM System i solution, blades are diskless and use the SAN Boot technique with iSCSI to load and run the operating system. Diskless blades can also be deployed in traditional SAN infrastructures with external iSCSI disk systems by using the same SAN Boot technique.

An IBM BladeCenter iSCSI solution's components include the iSCSI initiator on the blade itself, traditional Ethernet switches, an iSCSI target such as IBM System Storage DS3500 iSCSI models, an optional DHCP server, and a management station with iSCSI Configuration Manager.

Currently, iSCSI on blade servers can be implemented by using either software iSCSI initiators or hardware iSCSI initiators.

A software iSCSI initiator is specialized software that uses a server's processor for iSCSI protocol processing.

A software iSCSI initiator can exist on two separate 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 or to the Ethernet expansion card

Both implementations provide iSCSI capabilities for Ethernet NICs. However, an operating system driver can be used only 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 a software iSCSI initiator for download), or they can be obtained as a part of an NIC firmware upgrade (if supported by NIC), as in blade servers.

Software iSCSI initiators are supported on any 1 Gb and 10 Gb Ethernet expansion cards that are compatible with the blade servers.

Important: SAN Boot is not supported on software iSCSI initiators.

A hardware iSCSI initiator is an expansion card that has a dedicated ASIC for iSCSI processing.

5.3.1 Hardware-base iSCSI initiator support

Table 5-5 lists iSCSI support using a hardware-based iSCSI initiator.

Table 5-5 Hardware-based iSCSI support

Ethernet adapter	BladeCenter I/O module	External Ethernet switch	Operating systems	iSCSI storage targets
<ul style="list-style-type: none"> ▶ Emulex Virtual Fabric Adapter II, 90Y3550^a ▶ Emulex Virtual Fabric Adapter II Advanced, 90Y3566 ▶ Emulex 10GbE VFA II for HS23, 81Y3120^b ▶ Emulex 10GbE VFA Advanced II for HS23, 90Y9332 ▶ Integrated 10GbE LOM (HS23)^c 	<ul style="list-style-type: none"> ▶ IBM Virtual Fabric 10Gb Switch Module, 46C7191 (pNIC, vNIC1, vNIC2) 	<ul style="list-style-type: none"> ▶ Ethernet Switch ▶ IBM RackSwitch™ G8124E 	<ul style="list-style-type: none"> ▶ RHEL 5 ▶ RHEL 6 ▶ SLES 10 ▶ SLES 11 ▶ vSphere 4.1 ▶ vSphere 5.0 ▶ vSphere 5.1 ▶ Windows Server 2008 R2 ▶ Windows Server 2012 	<ul style="list-style-type: none"> ▶ DCS3700 ▶ DS3500 ▶ DS5000 ▶ SVC ▶ V3500 ▶ V3700 ▶ V7000 ▶ XIV
<ul style="list-style-type: none"> ▶ QLogic 10Gb Virtual Fabric Adapter, 00Y3332^d ▶ QLogic 10Gb Virtual Fabric CNA, 00Y5618 	<ul style="list-style-type: none"> ▶ IBM Virtual Fabric 10Gb Switch Module, 46C7191 (pNIC, vNIC2) 	<ul style="list-style-type: none"> ▶ Ethernet Switch 		

a. Requires an Advanced Upgrade, part number 49Y4265.

b. Requires an Advanced Upgrade, part number 90Y9350.

c. Requires an Advanced Upgrade, part number 90Y9310.

d. Requires QLogic 10Gb Virtual Fabric Advanced FoD Upgrade, part number 00Y5622.

5.3.2 iSCSI SAN boot

iSCSI SAN boot for blades enables diskless blades to be booted from an external iSCSI storage system or an IBM System i host by using a standard Ethernet port on the blade server itself or on an Ethernet expansion card that is installed in the blade server.

Two configuration techniques are available to do a SAN boot for blades using iSCSI:

- ▶ Use the basic input/output system (BIOS) and integrated network interface cards (NICs) on the blade itself to configure iSCSI initiators with the iSCSI Configuration Manager (also known as *IBM BladeBoot technology*).
- ▶ Use Unified Extensible Firmware Interface (UEFI) to configure iSCSI initiators for both integrated NICs and certain expansion cards.

Table 5-6 on page 460 lists currently available blades that support a built-in software iSCSI initiator that is used for diskless blade booting.

Upgrades: BIOS/UEFI, BMC/IMM, and NIC firmware upgrades might be required to support iSCSI boot with onboard Ethernet controllers or Ethernet expansion cards.

Table 5-6 IBM SAN boot support with built-in software iSCSI initiator

Blade server →										
Expansion card	Part number	HS12 (8028)	HS22 (7870)	HS22V (7871)	HS23 (7875)	HS23E (8038)	HX5 (7872)	HX5 (7873)	PS700/1/2 (8406)	PS703/4 (7891)
iSCSI Configuration Manager (IBM BladeBoot)										
Integrated Gigabit Ethernet	None	Y	N	N	N	N	N	N	N	N
2/4 Port Ethernet Card (CFFh)	44W4479	Y	N	N	N	N	N	N	N	N
UEFI										
Integrated Gigabit Ethernet	None	N	Y	Y	Y	Y	Y	Y	N	N
Ethernet Card (CIOv)	44W4475	N	Y	Y	Y	Y	Y	Y	N	N
2/4 Port Ethernet Card (CFFh)	44W4479	N	Y	Y	Y	Y	Y	Y	N	N
Broadcom 10Gb Gen 2 4-port. Ethernet (CFFh)	46M6164	N	Y	Y	Y	Y	Y	Y	N	N
Broadcom 10Gb Gen 2 2-port Ethernet (CFFh)	46M6168	N	Y	Y	Y	Y	Y	Y	N	N

The following operating systems are supported:

- ▶ RHEL 5
- ▶ RHEL 6
- ▶ SLES 10
- ▶ SLES 11
- ▶ Windows Server 2003
- ▶ Windows Server 2008

For additional IBM System Storage compatibility information, see the IBM System Storage Interoperability Center (SSCIC):

<http://www.ibm.com/systems/support/storage/config/ssic>

Separate network segment: Consider using a separate network segment for iSCSI traffic. NICs, switches (or VLANs), and storage system ports that participate 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 highly available (HA) clustering solutions by using iSCSI, but certain restrictions might apply (check the storage system vendor's compatibility guides). For example, Microsoft Cluster Services 2003 is not currently supported with multipathing for IBM iSCSI solutions (DS3300 and N3700).

When planning your iSCSI solution, consider the following items:

- ▶ IBM blade hardware, the initiators, and the operating system are 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 sources:

- ▶ *IBM BladeCenter iSCSI SAN Solution*, REDP-4153
- ▶ IBM System Storage N series Interoperability Matrix:
<http://www.ibm.com/support/docview.wss?uid=ssg1S7003897>

5.4 SAS storage system interoperability

In general, if IBM BladeCenter is implemented with SAS storage connectivity, ensure that the following requirements are met. Check the compatibility guides from your storage system vendor for confirmation.

- ▶ IBM BladeCenter server hardware and SAS HBA are supported by the storage system.
- ▶ The SAS Connectivity Modules are 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).
- ▶ Clustering software is supported by the storage system (in case you plan to implement clustering technologies).

If any of these requirements are not met, consider proposing another solution that is supported.

Almost every vendor of storage systems or storage fabrics has extensive compatibility matrixes, which include supported SAS HBAs, connectivity modules, and operating systems.

Table 5-7 lists external SAS storage support for blade servers.

Table 5-7 SAS support

SAS expansion card	BladeCenter I/O module	Operating systems	SAS storage targets
<ul style="list-style-type: none"> ▶ SAS Connectivity Card (CFFv), 43W3974^a ▶ SAS Expansion Card (CFFv), 39Y9190 ▶ SAS Expansion Card (CFFv), 44E5688 ▶ SAS Connectivity Card (CIOv), 43W4068 	<ul style="list-style-type: none"> ▶ SAS Connectivity Module, 39Y9195 	<ul style="list-style-type: none"> ▶ RHEL 5 ▶ RHEL 6 ▶ SLES 10 ▶ SLES 11 ▶ vSphere 4.0 ▶ vSphere 4.1 ▶ vSphere 5.0 ▶ Microsoft Windows Server 2003 / 2003 R2 ▶ Microsoft Windows Server 2008 / 2008 R2 	<ul style="list-style-type: none"> ▶ DS3200 ▶ DS3500 ▶ System Storage Tape^b ▶ System x HH LTO tape^c
	<ul style="list-style-type: none"> ▶ SAS RAID Controller Module, 43W3584^d 		<ul style="list-style-type: none"> ▶ EXP2500 ▶ EXP3000 ▶ System Storage Tape^b ▶ System x HH LTO tape^c
<ul style="list-style-type: none"> ▶ ServeRAID MR10ie (CIOv), 46C7167 	<ul style="list-style-type: none"> ▶ SAS Connectivity Module, 39Y9195 	<ul style="list-style-type: none"> ▶ RHEL 6 ▶ SLES 10 ▶ SLES 11 ▶ vSphere 5.0 	<ul style="list-style-type: none"> ▶ DS3200 ▶ DS3500
<ul style="list-style-type: none"> ▶ ServeRAID H1135 Controller (CIOv), 90Y4750 	<ul style="list-style-type: none"> ▶ SAS Connectivity Module, 39Y9195 	<ul style="list-style-type: none"> ▶ RHEL 6 ▶ SLES 10 ▶ SLES 11 ▶ vSphere 5.0 	<ul style="list-style-type: none"> ▶ DS3200 ▶ DS3500
<ul style="list-style-type: none"> ▶ SAS Connectivity Card (CIOv), FC 8246 	<ul style="list-style-type: none"> ▶ SAS Connectivity Module, FC 3267 	<ul style="list-style-type: none"> ▶ AIX 5.3 ▶ AIX 6.1 ▶ AIX 7.1 ▶ RHEL 5 ▶ SLES 10 ▶ SLES 11 ▶ VIOS 2.2 	<ul style="list-style-type: none"> ▶ DS3200 ▶ DS3500
	<ul style="list-style-type: none"> ▶ SAS RAID Controller Module, 43W3584 	<ul style="list-style-type: none"> ▶ VIOS 2.1 ▶ VIOS 2.2 	<ul style="list-style-type: none"> ▶ System Storage Tape

a. Does not support tape attachments.

b. Not supported with RHEL 6, vSphere 4.0, vSphere 4.1, and vSphere 5.0.

c. Not supported with vSphere 4.0, vSphere 4.1, and vSphere 5.0.

d. Supports System Storage tape attachments only. Does not support System x HH LTO tape.

For additional IBM System Storage compatibility information, see the IBM System Storage Interoperability Center:

<http://www.ibm.com/systems/support/storage/config/ssic>

5.5 High availability and redundancy

The IBM BladeCenter chassis has built-in redundancy. All Fibre Channel HBAs, SAS daughter cards, CNAs, or iSCSI cards for blade servers are dual port. All storage switches (Fibre Channel, SAS, or Ethernet) can be installed in pairs 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 Fibre Channel infrastructure is shown in Figure 5-1.

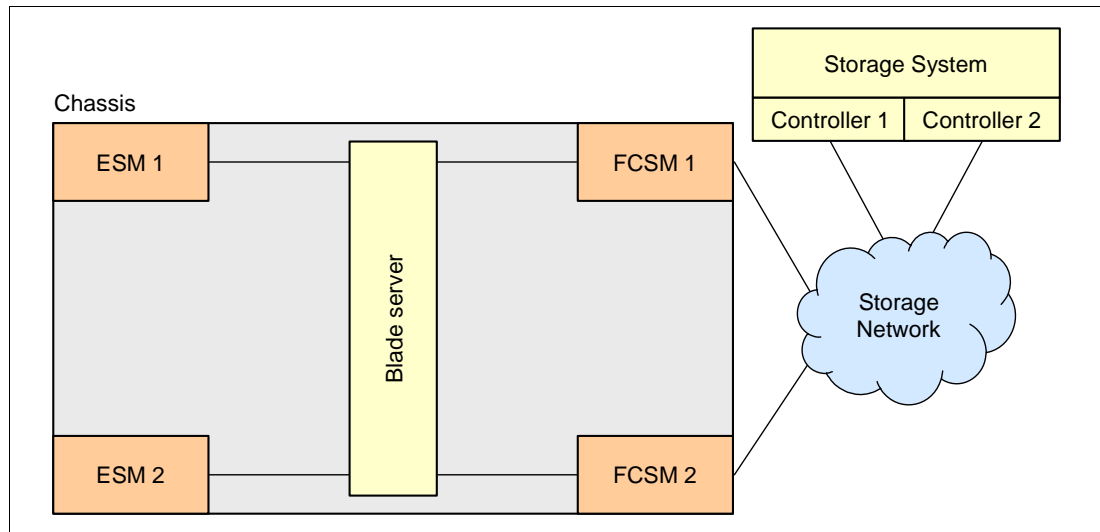


Figure 5-1 IBM BladeCenter SAN infrastructure topology

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.

If there is a failure, the specific operating system driver provided by the storage system manufacturer is responsible for the automatic failover process (also known as multipathing capabilities).

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 additional information, see the storage system vendor's documentation.

5.6 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 IBM System Storage DS8000, provides this functionality. If you plan to use such drivers, 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 are served by controller 1 and others are 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 grouping 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 feature is an optional feature that requires the purchase of an additional license.

For more information, see the storage system vendor's documentation and the switch vendor's documentation.

5.7 Backup solutions

Backup is one of the most important parts of the IBM BladeCenter solution. First, you need to decide which tape backup solution is implemented. The valid IBM BladeCenter possibilities are listed here (for more information, see 5.7.1, "A dedicated server for centralized LAN backup" on page 465):

- ▶ A centralized LAN backup with dedicated backup server (blade in chassis) with SCSI-attached or SAS-attached tape autoloader or tape library
- ▶ A centralized LAN backup with dedicated backup server (blade in chassis) with FC-attached tape autoloader or tape library
- ▶ A centralized LAN backup with dedicated backup server (non-blade) with SCSI-attached or SAS-attached tape drive, autoloader, or tape library
- ▶ A centralized LAN backup with dedicated backup server (non-blade) with FC-attached tape autoloader or tape library
- ▶ A LAN-free backup with FC-attached tape autoloader or library (see 5.7.2, "LAN-free backup for blades" on page 466.)

If you plan to use a 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 a 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. Verify that the tape device and type 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.

5.7.1 A 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 that require backup have backup agents, and backup traffic from these agents to the backup server use standard LAN paths.

If you use an FC-attached tape drive, then it should be connected to the FC fabric (or at least to the HBA) that is dedicated for backup; it should *not* be connected to the FC fabric that carries disk traffic. If you cannot use dedicated switches, then use zoning techniques on FC switches to separate these two fabrics.

Mixing storage: 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 having the tape and disk on the same HBA, you should separate these devices.

If you plan to use the blade server as a dedicated backup server with FC-attached tape, your choices are as follows:

- ▶ Use one port of the FC expansion card for tape and another for disk (there is no redundancy in this case).
- ▶ Choose the BladeCenter H with MSIM or BladeCenter HT with MSIM-HT chassis in conjunction with a suitable CFFh expansion card to increase the number of FC ports per blades to gain redundancy.

Figure 5-2 shows the possible topologies and traffic flows for LAN backups and FC-attached storage devices.

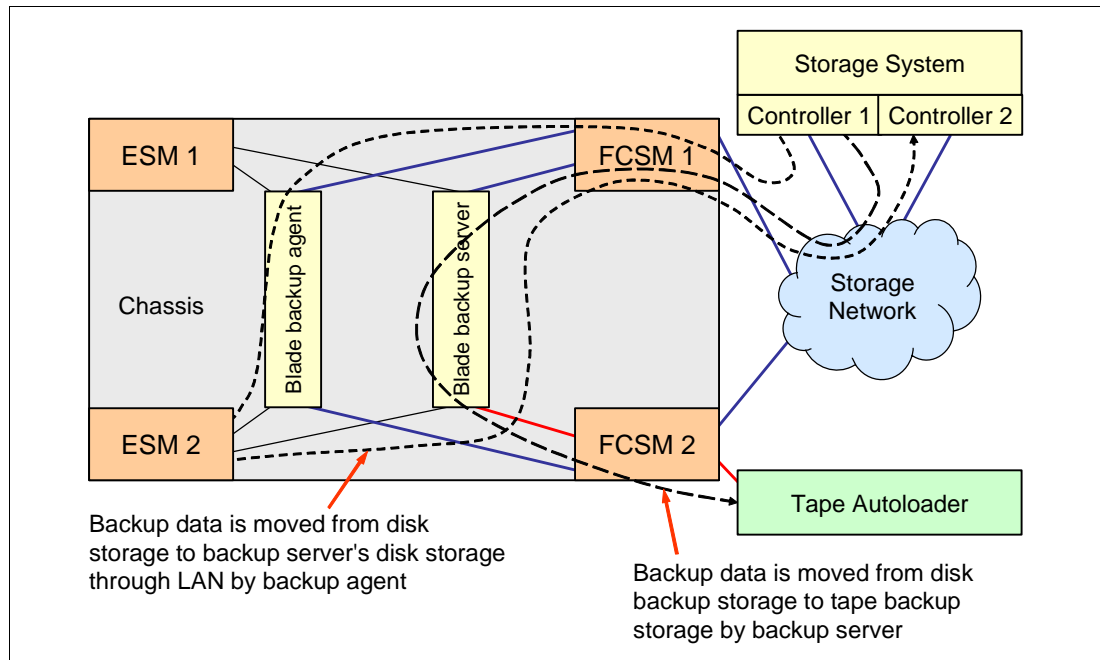


Figure 5-2 LAN backup topology and traffic flow

The topology in Figure 5-2 has the following characteristics:

- ▶ Each blade server that participates in the 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 that requires backup.

The backup traffic flow is as follows:

1. The backup agent transfers backup data from the disk storage to the backup server through LAN.
2. The backup server stores this data on its disk storage, for instance, on the same storage system.
3. The backup server transfers data from its storage directly to the tape device.
4. Zoning is implemented on an FCSM 2 to separate disk and tape data flows. Zoning is almost like VLANs in networks.

5.7.2 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.

Figure 5-3 on page 467 and Figure 5-4 on page 467 illustrate this process.

Figure 5-3 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.

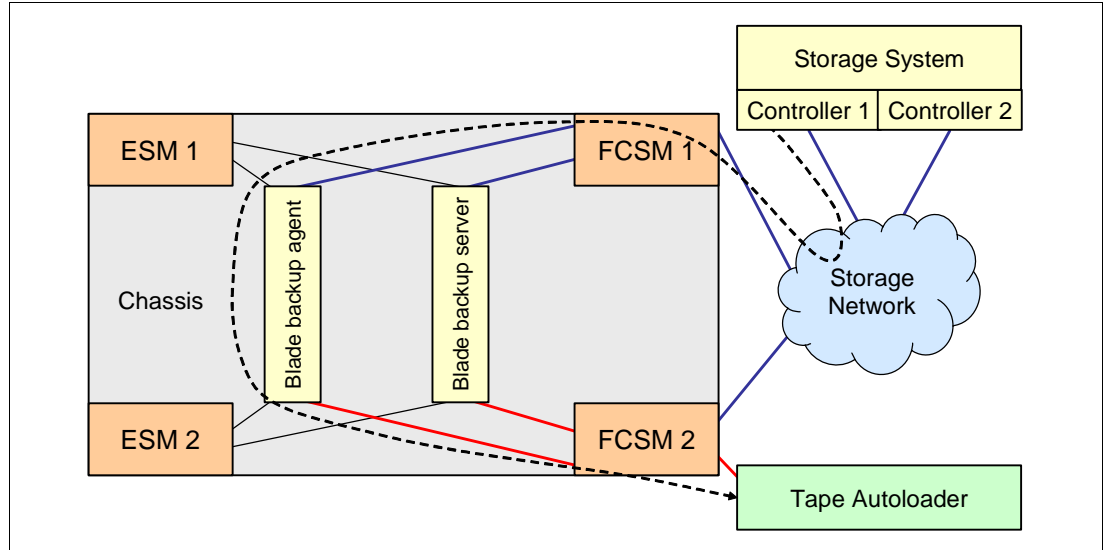


Figure 5-3 LAN-free backup without disk storage redundancy

If you need redundancy, you can implement a more complex topology, as shown in Figure 5-4.

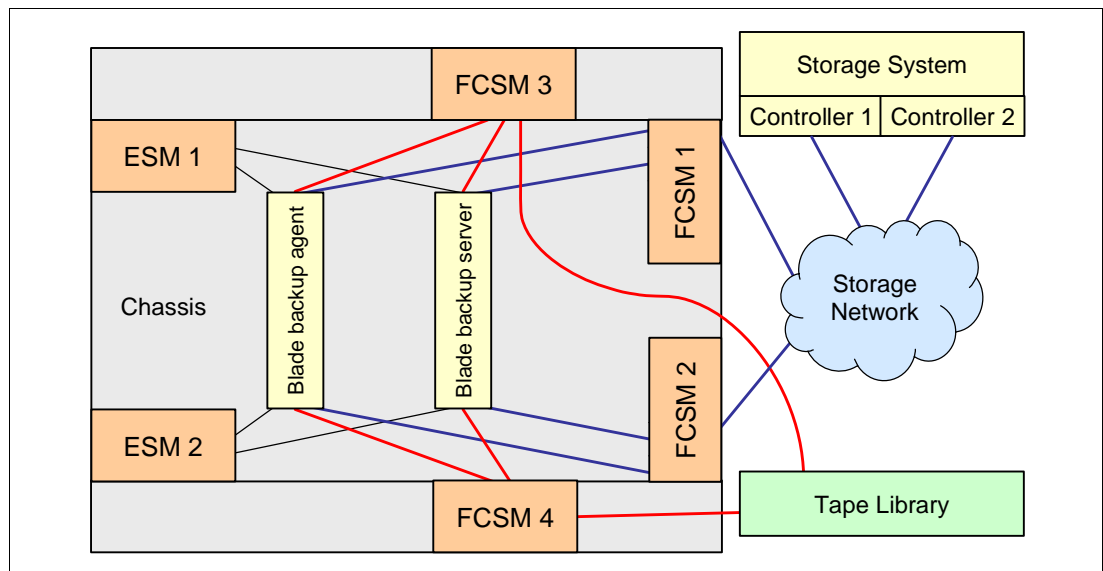


Figure 5-4 LAN-free backup with disk storage redundancy

This topology provides full redundancy both for disk storage and tape storage subsystems. In this case, you use BladeCenter H chassis with MSIM, and an FC-capable CFFh expansion card for each blade that requires backup and for the backup server. The backup traffic flow is the same as in the previous topology. FC switch zoning is not required because the *red* fabric is exclusively used for tape data flow.

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, and additional information.

5.8 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.

The BladeCenter SAS backup solution components include the following items:

- ▶ A supported BladeCenter chassis
- ▶ Supported blade servers
- ▶ A supported SAS expansion card installed in blade server
- ▶ Supported SAS I/O modules installed in the chassis
- ▶ A supported tape backup unit
- ▶ SAS cables
- ▶ Backup software

Supported SAS tape backup units include devices from both IBM System Storage (TS2230, TS2240, TS2250, TS2350, TS2900, TS3100, and TS3200) and IBM System x (HH LTO3, HH LTO4, and HH LTO5):

- ▶ IBM Half High LTO Generation 3 SAS Tape Drive, 43W8478
- ▶ IBM Half High LTO Generation 4 SAS Tape Drive, 44E8895
- ▶ IBM Half High LTO Gen 5 Internal SAS Tape Drive, 49Y9898
- ▶ IBM Internal Half High LTO Generation 5 SAS Tape Drive, 49Y9898
- ▶ IBM System Storage TS2230 Half-High LTO 3 SAS Tape Drive, 3580S3V
- ▶ IBM System Storage TS2240 Half-High LTO 4 SAS Tape Drive, 3580S4V
- ▶ IBM System Storage TS2250 Half-High LTO 5 SAS Tape Drive, 3580S5E
- ▶ IBM System Storage TS2350 LTO 5 SAS Tape Drive, 3580S5X
- ▶ IBM System Storage TS2900 Tape Autoloader with these items:
 - Half-High LTO 4 SAS Tape Drive, 3572S4R
 - Half-High LTO 5 SAS Tape Drive, 3572S5R
- ▶ IBM System Storage TS3100 Tape Library 35732UL (driveless) with these items:
 - Ultrium 3 Half-High SAS Drive V2
 - Ultrium 4 Half-High SAS Drive V2
 - Ultrium 4 SAS Drive
 - Ultrium 5 Half-High SAS Drive
 - Ultrium 5 SAS Drive
- ▶ IBM System Storage TS3200 Tape Library 35734UL (driveless) with these items:
 - Ultrium 3 Half-High SAS Drive V2
 - Ultrium 4 Half-High SAS Drive V2
 - Ultrium 4 SAS Drive
 - Ultrium 5 Half-High SAS Drive
 - Ultrium 5 SAS Drive

The HH LTO3, HH LTO 4, and HH LTO 5 are supported in external tape enclosures only.

Table 5-8 lists the compatibility information for entry-level series SAS-based tape drives, autoloaders, and libraries from IBM System Storage with Ultrium 5 HH V2 drives.

Consider the following information regarding Table 5-8:

- ▶ SAS connectivity is provided by SAS Connectivity Modules (part number 39Y9195) or SAS RAID Controller Modules (part number 43W3584).
- ▶ The following external SAS cables are supported with the TS family of tape devices:
 - 2.0 m Mini-SAS/Mini-SAS 1x cable, part number 95P4713
 - 5.5 m Mini-SAS/Mini-SAS 1x cable, part number 95P4714

Table 5-8 Entry IBM Tape Storage - SAS compatibility matrix

Tape Storage System →			TS2230, TS2240, TS2250, TS2350, TS2900, TS3100, and TS3200											
Operating system →			Microsoft Windows Server			Novell		Red Hat			IBM			
Expansion card			2008 R2	2008	2003 R2	SLES 11	SLES 10	RHEL 6	RHEL 5	RHEL 4	AIX 7.1	AIX 6.1	AIX 5.3	
Description	Part number or feature code													
IBM BladeCenter HS12 8028														
SAS Expansion Card (CFFv)	39Y9190	Y	Y	Y	Y	Y	N	Y	Y	N	N	N		
ServeRAID MR10ie (CIOv) ^a	46C7167	Y	Y	Y	Y	Y	N	Y	Y	N	N	N		
IBM BladeCenter HS22 7870														
SAS Connectivity Card (CIOv)	43W4068	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N		
ServeRAID MR10ie (CIOv) ^a	46C7167	Y	Y	Y	Y	Y	N	Y	Y	N	N	N		
IBM BladeCenter HS22V 7871														
SAS Connectivity Card (CIOv)	43W4068	Y	Y	N	Y	Y	Y	Y	Y	N	N	N		
ServeRAID MR10ie (CIOv) ^a	46C7167	Y	Y	N	Y	Y	N	Y	Y	N	N	N		
IBM BladeCenter HS23 7875 IBM BladeCenter HX5 7873														
SAS Connectivity Card (CIOv)	43W4068	Y	Y	N	Y	Y	Y	Y	N	N	N	N		
IBM BladeCenter HS23E 8038														
ServeRAID H1135 (CIOv)	90Y4750	N	N	N	N	N	N	N	N	N	N	N		
IBM BladeCenter PS700/701/702 8406														
SAS Connectivity Card (CIOv)	8246	N	N	N	Y	Y	Y	Y	N	Y	N	N		
IBM BladeCenter PS703/704 7891														
SAS Connectivity Card (CIOv)	8246	N	N	N	Y	Y	N	Y	N	Y	Y	Y		

a. Not supported with TS3100 and TS3200.

For additional IBM System Storage compatibility information, see the IBM System Storage Interoperability Center:

<http://www.ibm.com/systems/support/storage/config/ssic>

Table 5-9 lists the compatibility information for System x LTO3, LTO4, and LTO5 SAS-based tape drives.

Consider the following information regarding Table 5-9:

- ▶ SAS connectivity is provided by SAS Connectivity Modules (part number 39Y9195). SAS RAID Controller Modules (part number 43W3584) are not supported.
- ▶ The external SAS cables that are used for attachments of System x tape drives are included with the SAS Enclosure Adapter Kit Option, part number 40K2599.

Table 5-9 IBM System x Tape Drives - SAS compatibility matrix

Tape Drive →		HH LTO3, HH LTO4, HH LTO5									
Operating system →		Microsoft Windows Server				Novell			Red Hat		
Fibre Channel expansion card		2012	2008 R2	2008	2003 SP2	SLES 11	SLES 10	SLES 9	RHEL 6	RHEL 5	RHEL 4
Description	Part number										
IBM BladeCenter HS12 (8028)											
SAS Expansion Card (CFFv)	39Y9190	N	N	Y	Y	N	Y	Y	N	Y	Y
ServeRAID MR10ie (CIOv)	46C7167	N	N	Y	Y	N	Y	Y	N	Y	Y
IBM BladeCenter HS22 (7870)											
SAS Connectivity Card (CIOv)	43W4068	N	Y	Y	Y	Y	Y	N	Y	Y	Y
ServeRAID MR10ie (CIOv)	46C7167	N	N	Y	Y	N	Y	N	N	Y	Y
IBM BladeCenter HS22V (7871) ^a											
SAS Connectivity Card (CIOv)	43W4068	N	Y	Y	N	Y	Y	N	Y	Y	Y
ServeRAID MR10ie (CIOv)	46C7167	N	N	N	N	N	N	N	N	N	N
IBM BladeCenter HS23 (7875) IBM BladeCenter HX5 (7873)											
SAS Connectivity Card (CIOv)	43W4068	N	Y	Y	N	Y	Y	N	Y	Y	N
IBM BladeCenter HS23E (8038)											
ServeRAID H1135 (CIOv)	90Y4750	N	Y	Y	N	Y	Y	N	Y	Y	N

a. LTO3 and LTO4 drives are not supported.

To successfully implement a backup solution that is based on SAS tape devices, consider the following information:

- ▶ The solution components must be compatible with each other (for more information, see Table 5-8 on page 469 and Table 5-9).
- ▶ 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.
- ▶ Do not mix both tape devices and disk storage devices on the same HBA port because this configuration might lead to significant performance degradation of the backup operations.
- ▶ If the SAS module is used for both disk and tape operations for different servers, then plan to implement zoning to isolate different types of traffic.
- ▶ 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.

Additional compatibility information is at the IBM System Storage Interoperation Center:

<http://www.ibm.com/systems/support/storage/config/ssic>

5.9 External storage expansion (JBOD) attachments

With the ServeRAID MR10ie (CIOv) RAID expansion card for BladeCenter servers, you can attach external EXP2512 and EXP2524 enclosures to the supported BladeCenter chassis by using SAS Connectivity Modules. This setup 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.

Important: The EXP2500 ESM firmware level must be 60BG or later.

IBM System Storage EXP2512, part number 174712X, and IBM System Storage EXP2524, part number 174724X, are external expansion enclosures that are used with supported ServeRAID controllers to increase the number of drives and disk space available for attached systems. Figure 5-5 shows the EXP2524.



Figure 5-5 IBM System Storage EXP2524

IBM System Storage EXP2500 enclosures have the following characteristics:

- ▶ Two upstream and one downstream 6 Gb SAS ports on a single ESM. Currently, only one SAS port is supported to connect the EXP2500 unit to the SAS Connectivity Module.
- ▶ One 10/100 Mb Ethernet enclosure management port.
- ▶ EXP2512: Up to 12 hard disk drives.
- ▶ EXP2524: Up to 24 hard disk drives.
- ▶ Support for SAS and SATA HDDs, including intermixing of disk types.
- ▶ Dual redundant power supplies and cooling units.

Table 5-10 lists the solution components that are used in external “just a bunch of disks” (JBOD) attachments.

Table 5-10 External storage expansion solution components

Item	Component description	Part number	Feature code	Min/Max quantity
1	System Storage EXP2500	Varies	None	1/1 per blade server 1/8 per chassis
2 ^a	IBM 3 m SAS Cable	39R6531	3707	1 per EXP
	IBM 1 m SAS Cable	39R6529	3708	1 per EXP
3	BladeCenter E or other supported chassis	Varies	None	1
4	SAS Connectivity Module	39Y9195	3267	1/2 per chassis
5	BladeCenter HS22 or other supported blade server	Varies	None	1/chassis dependent
6	ServeRAID-MR10ie (CIOv)	46C7167	5490	1 per blade server

a. Only one of these two cables is required per expansion unit, depending on cable length requirements.

Supported chassis are as follows:

- ▶ BladeCenter E (8677)
- ▶ BladeCenter H (8852)
- ▶ BladeCenter S (8886)

Supported blade servers are as follows:

- ▶ HS12 (8028)
- ▶ HS22 (7870)
- ▶ HS22V (7871)

Connections: The SAS Connectivity Card (SBB 46C5068) and the ServeRAID-MR10ie card are both required for the HS12 to connect to the EXP3000 or EXP2500 enclosure.

Table 5-11 shows the hard disk drives that are supported in the EXP2512.

Table 5-11 Hard disk drives that are supported in the EXP2512

Description	Part	Feature	Max quantity
3.5-inch Hot-Swap SAS			
300 GB 15000 rpm 6 Gb SAS 3.5-inch HDD	49Y1899	5300	12
450 GB 15000 rpm 6 Gb SAS 3.5-inch HDD	49Y1900	5305	12
600 GB 15000 rpm 6 Gb SAS 3.5-inch HDD	49Y1901	5310	12
3.5-inch Hot-Swap SAS NL			
1 TB 7200 rpm 6 Gb SAS NL 3.5-inch HDD	49Y1903	5365	12
2 TB 7200 rpm 6 Gb SAS NL 3.5-inch HDD	49Y1902	5375	12
3 TB 7200 rpm 6 Gb SAS NL 3.5-inch HDD	90Y8720	5385	12

Important: The 3 TB 7200 rpm 6 Gb SAS NL 3.5-inch HDD (feature code 5385, part number 90Y8720) is not supported in an EXP2512 that is connected to the IBM BladeCenter. For more information, see the following link:

<http://www.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5088976>

Table 5-12 shows the hard disk drives that are supported in EXP2524.

Table 5-12 Hard disk drives that are supported in EXP2524

Description	Part number	Feature	Maximum quantity
2.5-inch Hot-Swap SAS NL			
500 GB 7,200 rpm 6 Gb SAS NL 2.5-inch HDD	49Y1898	5465	24
1 TB 7,200 rpm 6 Gb SAS NL 2.5-inch HDD	81Y9952	5470	24
2.5-inch Hot-Swap SAS			
300 GB 10,000 rpm 6 Gb SAS 2.5-inch HDD	49Y1895	5410	24
600 GB 10,000 rpm 6 Gb SAS 2.5-inch HDD	81Y9596	5420	24
900 GB 10,000 rpm 6 Gb SAS 2.5-inch HDD	81Y9948	5425	24
146 GB 15,000 rpm 6 Gb SAS 2.5-inch HDD	49Y1896	5405	24
300 GB 15,000 rpm 6 Gb SAS 2.5-inch HDD	81Y9944	5406	24
2.5-inch Hot-Swap SAS SSD			
200 GB 2.5-inch SAS SSD	81Y9956	5550	24
400 GB 2.5-inch SAS SSD	81Y9960	5555	24

Figure 5-6 shows a sample solution with the HS22 and EXP2524.

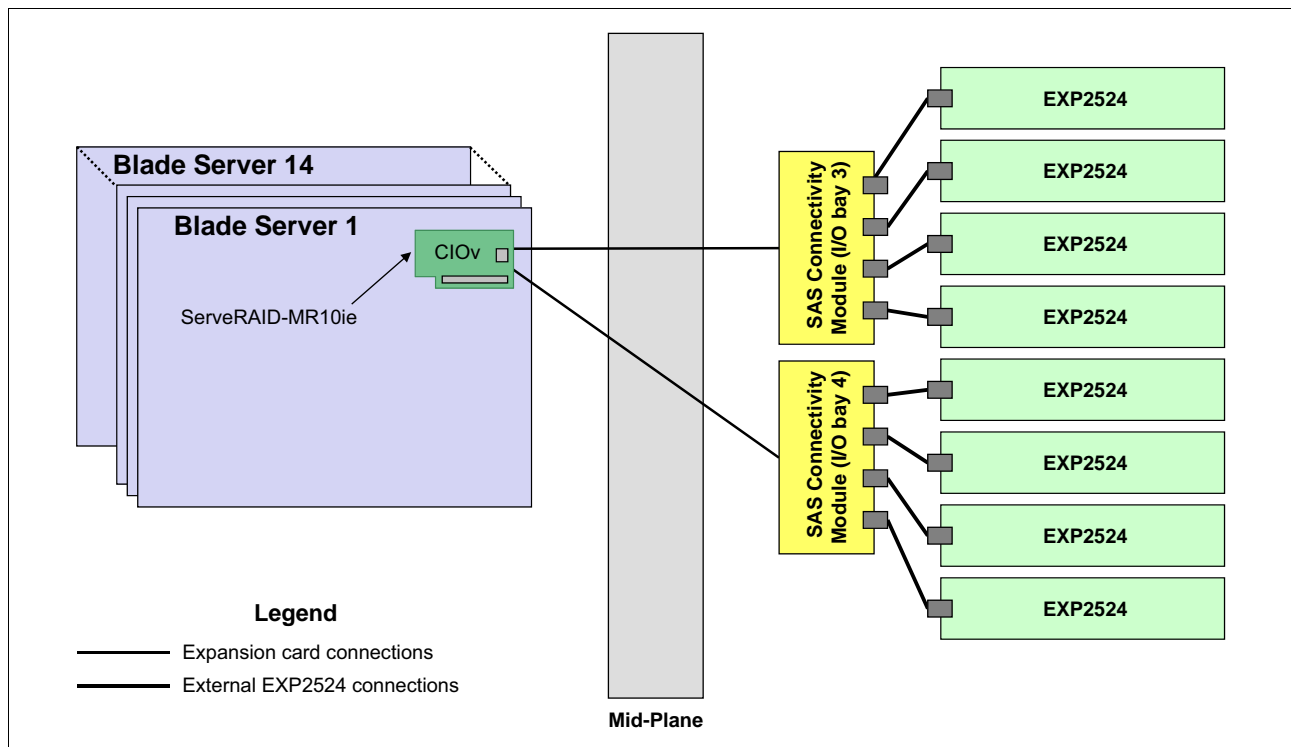


Figure 5-6 Solution topology with HS22 and EXP2524 attachments

When you plan EXP3000 or EXP2500 attachments to the BladeCenter, consider the following information:

- ▶ Every blade that requires 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 EXP2500 units (up to eight EXP2500s, one per blade, with two SAS Connectivity Modules) can be attached to the BladeCenter chassis to increase the number of drives per blade.
- ▶ A maximum of one expansion unit is supported per single blade server.
- ▶ A daisy-chain connection of multiple enclosures to a single SAS Connectivity Module is not supported.
- ▶ There is no redundancy between SAS Connectivity Modules. Each module supports its own set of attached EXPs.
- ▶ An EXP2500 can be connected to only one SAS connectivity Module, even though the environmental service module has two upstream SAS ports. The EXP2500 does not support zoning of the disk drives within the enclosure and cannot be shared among multiple SAS Connectivity Modules.
- ▶ All disks in a single unit can belong only 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 SAS Connectivity Module controls disk assignments by using zoning-like techniques. This situation results in the SAS module maintaining a single, isolated dedicated path between physical HDDs and the blade server. The administrator configures the HDD

assignments and uses predefined templates or creates custom configurations. Custom zoning *must* be implemented to ensure that only one blade server has access to one expansion unit.

- ▶ RAID functionality is supplied by the ServeRAID-MR10ie controllers installed into blade servers. RAID levels supported by ServeRAID-MR10ie are as follows:
 - 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 26 (two internal drives plus 24 drives in one EXP2524).
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 also. However, each volume must have hard disks of the same type, that is, SAS or SATA.

For more information, see the *ServeRAID-MR10ie Controller Installation and User Guide*:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-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 two ways to implement the integrated storage solution for BladeCenter S:

- ▶ Basic local storage using the SAS Connectivity Module
- ▶ Advanced shared storage using the SAS RAID Controller Module

Feature codes: Feature codes for I/O modules listed in this section are used when ordering the IBM BladeCenter S chassis.

5.10.1 Basic local storage using the 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. It also creates 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 in the blade server.

Table 5-13 on page 476 lists the basic local storage solution components for BladeCenter S.

Table 5-13 Basic local storage solution components for BladeCenter S

Item	Component description	Part number	Feature code	Min/max quantity
1	6-Disk Storage Module (DSM)	43W3581	4545	1 / 2
2	SAS Connectivity Module	39Y9195	3267	1 / 2
3 ^a	SAS Expansion Card (CFFv)	39Y9190	2979	1 per blade server
	SAS Connectivity Card (CIOv)	43W4068	1041	1 per blade server
	ServeRAID H1135 (CIOv) ^b	90Y4750	A1XJ	1 per blade server
4	Supported blade server	Varies	None	1 / 6

a. Only one of these cards is required per blade server, depending on the blade server itself.

b. ServeRAID H1135 is used with HS23E (8038) only.

Important: The SAS Connectivity Card (SBB 46C5068) can also be used with an HS12 server to support connections to integrated Disk Storage Modules (DSMs). This card is only available through CTO or special bid.

For more information, see the following sections:

- ▶ For DSMs, see 2.4.5, “BladeCenter S disk storage modules” on page 75.
- ▶ For SAS Connectivity Modules, see 2.10, “SAS I/O modules” on page 168.
- ▶ For the SAS Expansion Card (CFFv), see 3.1.19, “SAS Expansion Card (CFFv)” on page 248.
- ▶ For the SAS Connectivity Card (CIOv), see 3.1.20, “SAS Connectivity Card (CIOv)” on page 249.
- ▶ For the ServeRAID H1135, see 3.1.21, “ServeRAID H1135 (CIOv) Controller” on page 250

Table 5-14 lists the hard disk drives supported in 6-Disk Storage Modules by SAS Connectivity Modules.

Table 5-14 Hard disk drives supported in 6-disk DSMs by SAS Connectivity Modules

Description	Part number	Feature code	Maximum quantity
3.5-inch Hot-Swap SATA			
1000 GB Dual-Port Hot-Swap SATA HDD	43W7630	5561	12 (6 per one DSM)
3.5-inch Hot-Swap NL SAS			
IBM 1 TB 7200 NL SAS 3.5-inch HS HDD	42D0547	5423	12 (6 per one DSM)
IBM 1 TB 7.2 K 6 Gbps NL SAS 3.5-inch HDD	42D0777	5418	12 (6 per one DSM)
IBM 2 TB 7.2 K 6 Gbps NL SAS 3.5-inch HDD	42D0767	5417	12 (6 per one DSM)
3.5-inch Hot-Swap SAS			
IBM 300 GB 15 K 6 Gbps SAS 3.5-inch Hot-Swap HDD	44W2234	5311	12 (6 per one DSM)
IBM 450 GB 15 K 6 Gbps SAS 3.5-inch Hot-Swap HDD	44W2239	5312	12 (6 per one DSM)
IBM 600 GB 15 K 6 Gbps SAS 3.5-inch Hot-Swap HDD	44W2244	5313	12 (6 per one DSM)

Figure 5-7 shows a sample connection topology for basic local storage with one SAS Connectivity Module installed.

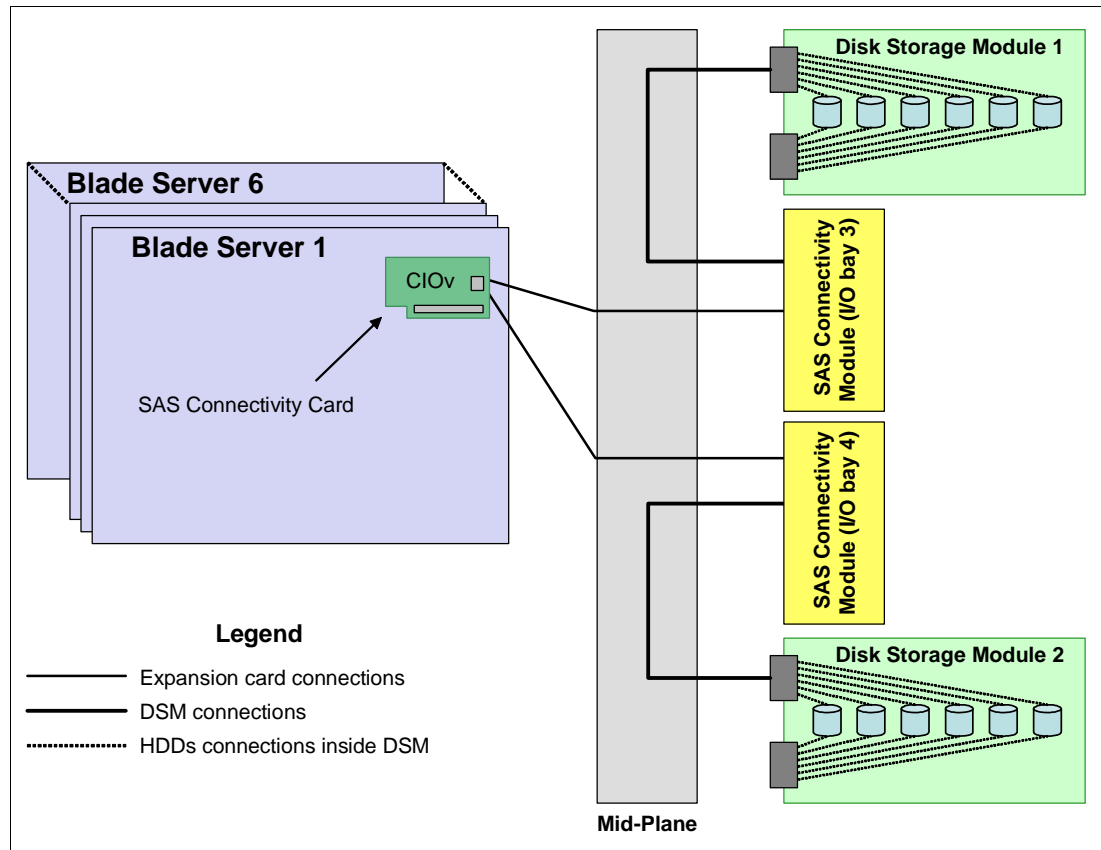


Figure 5-7 SAS I/O connections with one SAS Connectivity Module installed

Figure 5-8 shows a sample connection topology for basic local storage with two SAS Connectivity Modules installed.

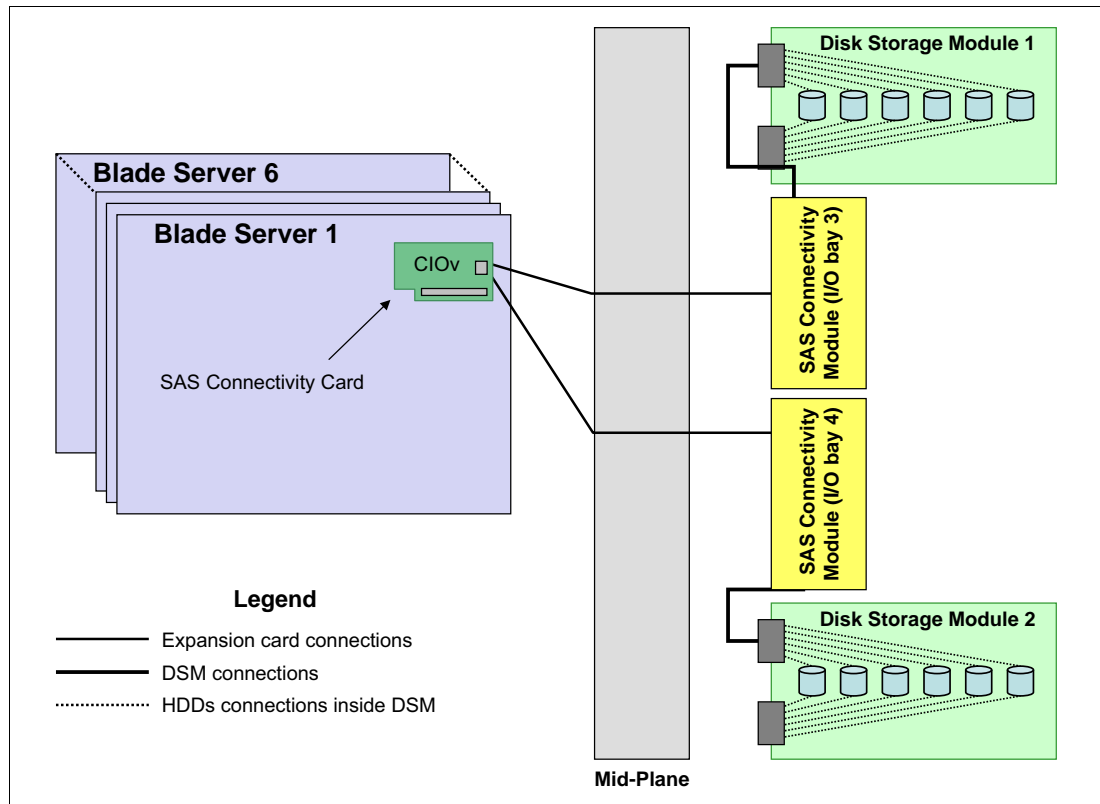


Figure 5-8 SAS I/O connections with two SAS Connectivity Modules installed

You must consider the following information when you plan BladeCenter S basic local storage implementations:

- ▶ Every blade that requires integrated storage connectivity must have one SAS Expansion Card, SAS Connectivity Card, or ServeRAID H1135 (for HS23E only) 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 include, as standard, any DSMs.
- ▶ 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 through an SAS expansion card. There are no LUNs, storage partitions, and so on, as in shared storage systems. Instead, each blade has its own set of physical disks in a DSM (or DSMs).

- ▶ The SAS Connectivity Module controls disk assignments by using zoning-like techniques, which results in the SAS module maintaining a single, isolated dedicated path between physical HDDs and the blade server. The administrator configures the HDD assignments and uses predefined templates or creates custom configurations.
- ▶ The RAID functionality is supplied either by the SAS Expansion Card or onboard SAS controller when SAS Connectivity Cards are used in the blade servers. The following RAID levels are supported by the SAS Expansion Card or by the onboard RAID controller on the blade itself:
 - 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 drive. The IME volume supports up to 10 HDDs plus two optional hot spares. The IS volume supports up to 10 HDDs. The IS volume does not support hot-spare drives.
- ▶ When creating a RAID 1 array, span both disk storage modules. This configuration 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 5-8 on page 478.
- ▶ 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 volumes 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 use 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. 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 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.
- ▶ Map a single logical volume to several blade servers simultaneously.

Table 5-15 lists the advanced local storage components for BladeCenter S.

Table 5-15 Advanced local storage solution components for BladeCenter S

Item	Component description	Part number	Feature code	Min/max ^a quantity
1	6-Disk Storage Module (DSM)	43W3581	4545	1 / 2
	12-Disk Storage Module (DSM)	49Y3234	A3KS	1 / 2
2	SAS RAID Controller Module	43W3584	3734	1 ^b / 2
3 ^c	SAS Expansion Card (CFFv)	39Y9190	2979	1 per blade server
	SAS Connectivity Card (CIOv)	43W4068	1041	1 per blade server
	ServeRAID H1135 (CIOv) ^d	90Y4750	A1XJ	1 per blade server
4	Ethernet Switch in I/O bay 1 ^e	Varies	Varies	1 / 1
5	Supported blade server	Varies	None	1 / 6

a. Minimum/maximum

b. Single Controller configuration is supported.

c. Only one of these cards is required per blade server, depending on the blade server itself.

d. ServeRAID H1135 is used with HS23E (8038) only.

e. For a list of compatible Ethernet switch modules, see Table 1-2 on page 3.

Connectivity: The SAS Connectivity Card (SBB 46C5068) can also be used with HS12 or HS21XM to support connections to SAS RAID Controller Modules. This card is only available through CTO or special bid.

Single controller configuration: Starting with SAS RAID Controller Firmware Version 1.2.2.007, a single controller configuration is supported.

The following Ethernet switch modules are supported:

- ▶ IBM 1/10 Gb Uplink Ethernet Switch Module, 44W4404, feature code 6980
- ▶ IBM 2/3 Copper Gb Ethernet Switch, 32R1860, feature code 3212
- ▶ IBM 2/3 Fiber Gb Ethernet Switch, 32R1861, feature code 3213
- ▶ IBM Layer 2/7 Copper Gb Ethernet Switch, 32R1859, feature code 1494
- ▶ Cisco Catalyst Switch Module 3012, 46C9272, feature code A3FE
- ▶ Server Connectivity Module, 39Y9324, feature code 3220
- ▶ Intelligent Copper Pass-thru Module, 44W4483, feature code 5865

Copper Pass-thru Module support: 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 a standard Ethernet cable.

Table 5-16 lists the hard disk drives supported by SAS RAID Controller Modules in 6-disk storage modules.

Table 5-16 Hard disk drives supported in 6-disk DSMs by SAS RAID Controller Modules

Description	Part number	Feature code	Maximum quantity
3.5-inch Hot-Swap NL SAS			
IBM 1 TB 7.2 K 6 Gbps NL SAS 3.5-inch HDD	42D0777	5418	12 (6 per one DSM)
IBM 2 TB 7.2 K 6 Gbps NL SAS 3.5-inch HDD	42D0767	5417	12 (6 per one DSM)
3.5-inch Hot-Swap SAS			
IBM 300 GB 15 K 6 Gbps SAS 3.5-inch Hot-Swap HDD	44W2234	5311	12 (6 per one DSM)
IBM 450 GB 15 K 6 Gbps SAS 3.5-inch Hot-Swap HDD	44W2239	5312	12 (6 per one DSM)
IBM 600 GB 15 K 6 Gbps SAS 3.5-inch Hot-Swap HDD	44W2244	5313	12 (6 per one DSM)

Table 5-17 lists the hard disk drives supported by SAS RAID Controller Modules in 12-disk storage module.

Table 5-17 Hard disk drives supported in 12-disk DSMs by SAS RAID Controller Modules

Description	Part number	Feature code	Maximum quantity
2.5-inch Hot-Swap SAS			
IBM 300GB 2.5in SFF 10K 6Gbps HS SAS HDD	90Y8877	A2XC	24 (12 per one DSM)
IBM 600GB 2.5in SFF 10K 6Gbps HS SAS HDD	90Y8872	A2XD	24 (12 per one DSM)
IBM 900GB 2.5in SFF HS 10K 6Gbps SAS HDD	81Y9650	A282	24 (12 per one DSM)
IBM 300GB 2.5in SFF HS 15K 6Gbps SAS HDD	81Y9670	A283	24 (12 per one DSM)

For more information, see the following sections:

- For DSMs, see 2.4.5, “BladeCenter S disk storage modules” on page 75.
- For SAS RAID Controller Modules, see 2.10.2, “SAS RAID Controller Module” on page 171.
- For SAS expansion cards, see 3.1.19, “SAS Expansion Card (CFFv)” on page 248.
- For the SAS Connectivity Card (CIOv), see 3.1.20, “SAS Connectivity Card (CIOv)” on page 249.

Figure 5-9 shows a sample topology for a BladeCenter S server with two SAS RAID Controller Modules.

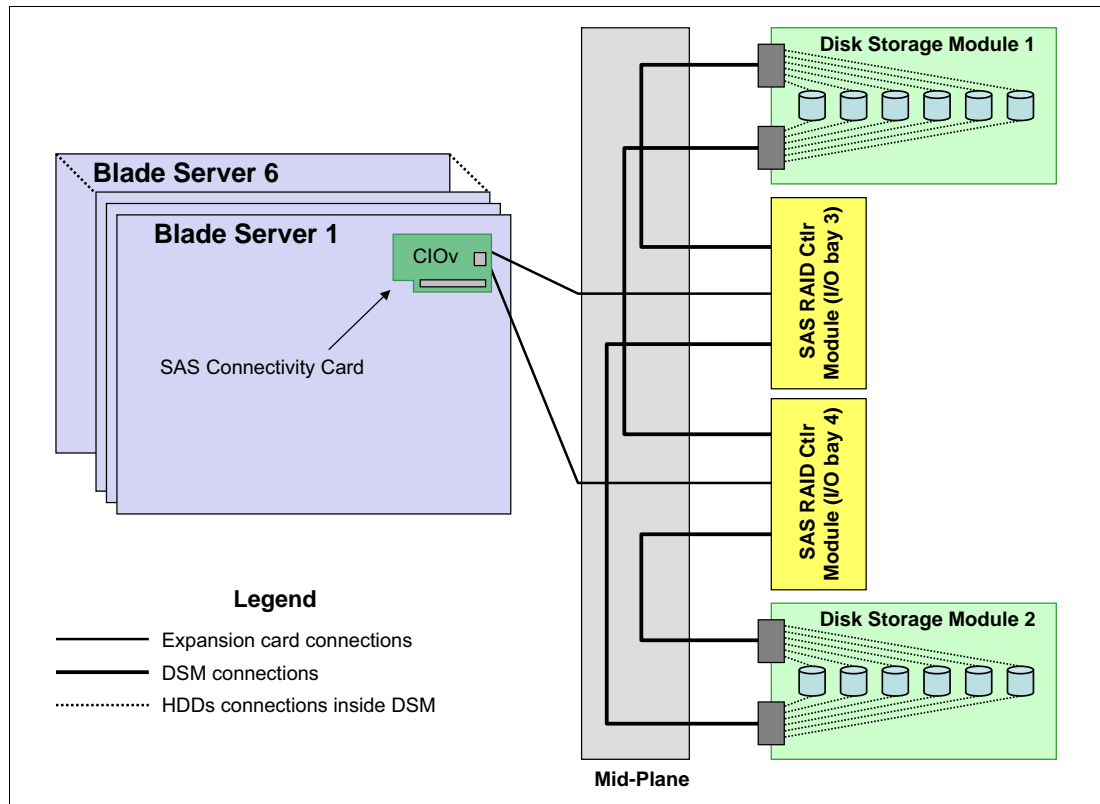


Figure 5-9 BladeCenter S SAS RAID Controller connections topology

Upgrades: Starting with the firmware release 1.2.0, SAS RAID Controller Module now supports online volume expansion, online storage pool expansion, concurrent code update, and online DSM replacement for RAID 1 and 10 configurations. Before this firmware level, the system must be stopped to perform capacity expansion, firmware upgrade, or DSM servicing.

Take several considerations into account when planning BladeCenter S advanced shared storage implementations:

- ▶ Every blade that requires 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 include, as standard, any DSMs.
- ▶ One or 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, that is, 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, and 10.
 - The maximum volume size is limited by the size of the storage pool.
 - The maximum number of volumes is 16 per blade server (maximum of 128 volumes per chassis).
 - One volume can be mapped to all six 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 also. 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 and Novell SUSE Linux.
- ▶ External tape attachments are supported. For more information, see 5.8, “SAS tape storage for IBM BladeCenter” on page 468.

Table 5-18 lists BladeCenter S SAS RAID Controller solution compatibility information.

Table 5-18 SAS RAID Controller compatibility matrix - blade servers by expansion cards

Blade server →			HS12 (8028)	HS21 (8853)	HS21 XM (7995)	HS23 (7875)	HS23E (8038)	HX5 (7872)	HX5 (7873)	PS700/701/702	PS703/704 (7981)
Description	Part number	Feature code									
SAS expansion cards											
SAS Expansion Card (CFFv)	39Y9190	2979	Y	Y	Y	N	N	N	N	N	N
SAS Connectivity Card (CFFv)	None ^a	1591	Y	N	Y	N	N	N	N	N	N
SAS Connectivity Card (CIOv)	43W4068	1041	N	N	N	Y	N	Y ^b	Y ^b	Y	N
ServeRAID H1135	90Y4750	A1XJ	N	N	N	N	Y	N	N	N	N

a. This card is only available through CTO or special bid.

b. Requires solid-state drive (SSD) Expansion Card, part number 46M6908.

Important:

- ▶ Microsoft Cluster Services are only supported with up to six nodes for Windows Server 2003 and 2008 clusters at the time of writing.
- ▶ Linux native HA clustering is supported.
- ▶ VMware VMotion is supported.

For more information about the SAS RAID Controller solution, see the SAS RAID Controller Module documentation:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5078491>

5.11 Converged networks

The term *converged networking* refers to the ability to carry both user data (LAN) and storage data (SAN) network traffic over the same unified fabric. This configuration reduces hardware, electrical power, cooling, and management costs by simplifying the data center infrastructure, which includes reducing the number of ports (on both server adapters and switches), adapters, switches, and the cabling required to build such infrastructure.

There are two key components of converged networks:

- ▶ Converged network adapters (CNAs)
- ▶ Converged switches

CNAs provide host connections with support for both LAN and SAN data transfer protocols. Converged switches manage flows for these two types of traffic, such as forwarding packets, ensuring reliable in-order delivery, guaranteeing SLA for certain types of traffic with quality of service, congestion control, and so on.

The simplest example of a converged network is Ethernet with iSCSI storage targets. In this case, an Ethernet adapter that is installed into the server processes both user LAN and iSCSI storage traffic, and the Ethernet switches play the role of the unified fabric.

In recent years, InfiniBand technology has been adopted by some customers as a unified fabric. This scenario is a more complex because, in addition to InfiniBand host adapters and switches, special devices called *bridges* must be used to connect the InfiniBand fabric to the standard Ethernet (InfiniBand-to-Ethernet bridges) or Fibre Channel (InfiniBand-to-Fibre Channel bridges) networks. The benefits of using such a fabric are offset by the upfront costs of the hardware.

The 10 Gb Converged Enhanced Ethernet (CEE) is a technology that is considered to be an affordable and convenient way to build a converged fabric with support for both Ethernet and Fibre Channel over Ethernet (FCoE) protocols to connect to standard LANs and SANs.

Fibre Channel over Ethernet (FCoE) is a standard that specifies how Fibre Channel protocol is carried over a traditional Ethernet network to enable customers to use proven Fibre Channel software stacks, zoning architectures, and management tools in the converged environment. However, FCoE requires enhancements to be implemented in the underlying Ethernet fabric to make it more reliable and responsive to avoid frame losses (also known as lossless Ethernet). With the introduction of 10 Gb Converged Enhanced Ethernet technology, the standard capabilities of 10 Gb Ethernet were enhanced to make it lossless.

The 10 Gb CEE is an enhanced version of traditional 10 Gb Ethernet that has additional protocols implemented to better control a consolidated data center I/O infrastructure. The key enhancements are as follows:

- ▶ Priority-based Flow Control (802.1Qbb) to eliminate frame loss because of congestion from using an 802.3x PAUSE-like mechanism for individual user priorities, as defined by the IEEE 802.1p specification.
- ▶ Enhanced Transmission Selection (802.1Qaz) to share bandwidth among different traffic classes more efficiently. When a particular traffic flow does not use all of the bandwidth available to it per the traffic classification, then the unused bandwidth could be allocated to other traffic flows.
- ▶ Data Center Bridging Exchange Protocol (part of 802.1Qaz coverage) to simplify network deployment and to reduce configuration errors by providing autonegotiation of IEEE 802.1 DCB features between the network interface card (NIC) and the switch and between switches.

FCoE and CEE together form an up-to-date and solid foundation for data center infrastructure consolidation with converged networking by offering Fibre Channel Over Converged Enhanced Ethernet (FCoCEE) technology. The term *FCoCEE* does not represent any formal standard, and is used to enforce the meaning of CEE in an FCoE implementation. FCoCEE means that the FCoE protocol is running over CEE, and not over standard Ethernet. FCoCEE combines two formal standards (FCoE and CEE) into one common term. FCoCEE solutions for BladeCenter are already available to customers.

For more information about FCoE and FCoCEE, see *An Introduction to Fibre Channel over Ethernet, and Fibre Channel over Convergence Enhanced Ethernet*, REDP-4493.

The FCoCEE solution components include switches, adapters, and targets:

- ▶ FCoCEE-capable 10 Gb Ethernet switches
- ▶ FCoCEE-capable converged network adapters used in host systems
- ▶ Storage targets:
 - Native FCoCEE targets (targets that have native FCoCEE interfaces)
 - Standard FC targets (targets that have native FC interfaces and are connected to the converged network infrastructure with FCoCEE-to-FC bridges or gateway modules)
 - Standard iSCSI targets (targets that support iSCSI and have native 10 Gb Ethernet interfaces)

The traditional FC-based SANs built on FC switches and traditional Ethernet LANs built on Ethernet switches can be seamlessly integrated into a converged networking infrastructure. FC SANs can be integrated by using optional or built-in (like in the Brocade Converged 10 GbE Switch Module) FCoCEE-to-FC gateways that have native FC ports. Ethernet LANs can be integrated *natively* without any additional modules. An FCoCEE-capable port operates as a standard Ethernet port when the device attached to it does not support FCoCEE.

When planning converged network implementation, ensure that the following requirements are met (check the compatibility guides from your storage system vendor):

- ▶ IBM BladeCenter server hardware and HBAs are supported by the storage system.
- ▶ The FCoCEE fabric (10 Gb CEE switches, FC switches, FCoE switches, or bridge or gateway modules, if any) proposed 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).
- ▶ Clustering software is supported by the storage system (in case you plan to implement clustering technologies).
- ▶ SAN Boot is supported by all required infrastructure components (if you plan to use diskless blades).

If any of these requirements are not met, consider using 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.

Detailed compatibility information is at the IBM System Storage Interoperability Center:

<http://www.ibm.com/systems/support/storage/config/ssic>

5.11.1 FCoCEE solution using Virtual Fabric Switches and Virtual Fabric Extension Modules

The BladeCenter VFS and VFE FCoCEE solution requires the following components:

- ▶ IBM BladeCenter H or HT chassis
- ▶ Supported CNA (CFFh) installed in supported blade servers
- ▶ IBM Virtual Fabric 10 Gb Switch Module
- ▶ QLogic Virtual Fabric Extension Module

Figure 5-10 shows the FCoCEE solution topology internal to the BladeCenter H chassis. Bridge bays 3 and 5 both have internal connections to the high-speed I/O bays 7 and 9.

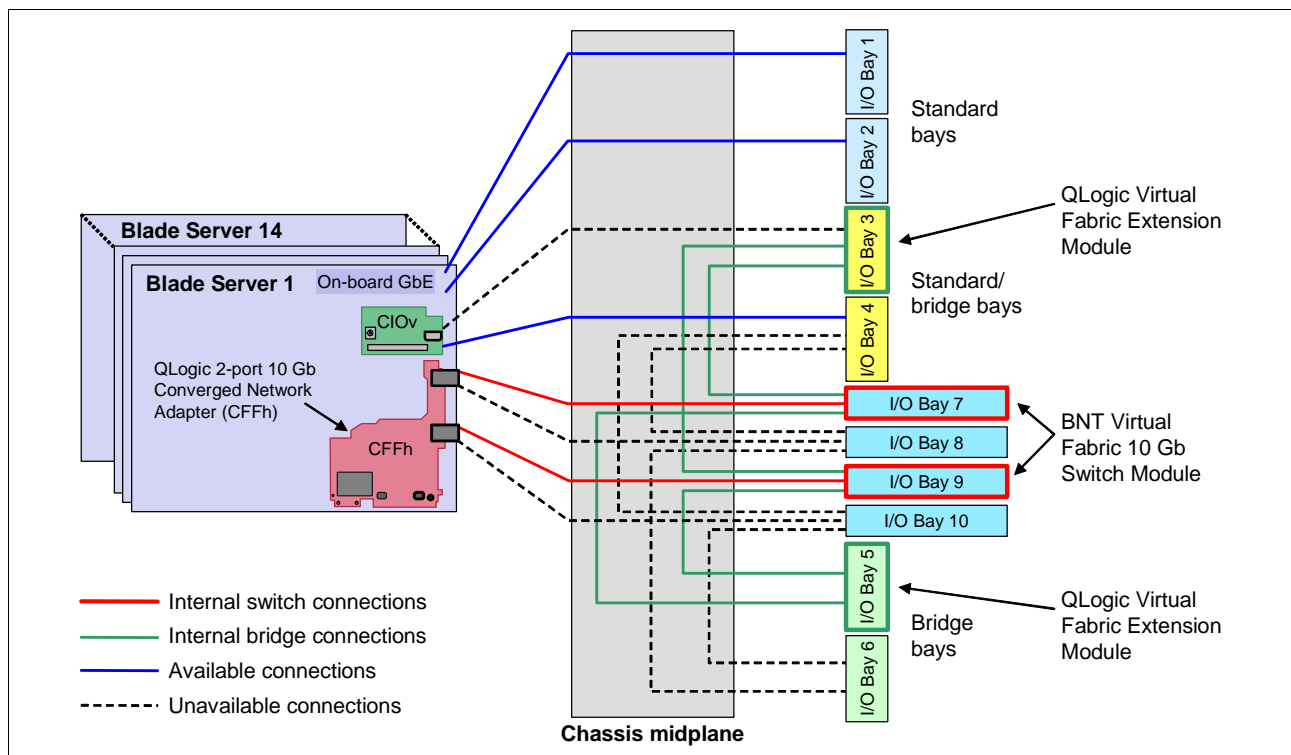


Figure 5-10 BladeCenter H-based FCoCEE solution topology

Three configurations are possible for the BladeCenter based FCoCEE solution:

- ▶ Maximum redundancy and bandwidth
- ▶ Maximum flexibility
- ▶ Entry level

Maximum redundancy and bandwidth

The configuration has the following components:

- ▶ Two switch modules in bays 7 and 9 of the chassis
- ▶ Two extension modules in bays 3 and 5
- ▶ One Converged Network Adapter (CFFh) in each server

As shown in Figure 5-10 on page 486, the switch in bay 7 is connected to the extension module in bay 3, and the switch in bay 9 is connected to the extension module in bay 5. This configuration results in 40 Gbps links between the switch and connected extension module. Four external ports of each switch module (eight total) are disabled (ports 1 - 4 in bay 7 and ports 7 - 10 in bay 9), and those circuits are rerouted to the extension modules. This configuration is best for redundancy and bandwidth. The total available external ports are 12 Ethernet and 12 Fibre Channel.

CIOv card: This configuration precludes the use of bay 3 and bay 4 for I/O modules for use with a CIOv card. Do not install a CIOv card in any blades installed in this chassis.

Maximum flexibility

The configuration has the following components:

- ▶ Two switch modules in bays 7 and 9 of the chassis
- ▶ One extension module in bay 5
- ▶ One Converged Network Adapter (CFFh) in each server

In this configuration, both the switch in bay 7 and the switch in bay 9 are connected to the extension module in bay 5. See Figure 5-10 on page 486 for the connections. This configuration results in 20 Gbps links between each switch and the extension module. Two external ports of each switch module (four in total) are disabled (ports 1 and 2 in bay 7, and ports 9 and 10 in bay 9), and those circuits are rerouted to the extension module. Although the extension module is no longer redundant, it still offers the advantages of two paths from the adapter to the switch and extension module. Unlike the configuration described in “Maximum redundancy and bandwidth” on page 487, you can use the CIOv slot in the blade servers (for example, for SAS connectivity) to connect to I/O modules installed in bays 3 and 4 of the chassis. The total available external ports are 16 Ethernet and 6 Fibre Channel.

Entry level

The configuration has the following components:

- ▶ One switch module in bay 7 (or bay 9) of the chassis
- ▶ One extension module in bay 5 (or bay 3)
- ▶ One Converged Network Adapter (CFFh) in each server

In this configuration, only one port of the Converged Network Adapter is active and one path is established from the adapter to the switch and on to the extension module. This configuration results in a 40 Gbps link between the switch and extension module. Four external ports of the switch module are disabled (ports 1 - 4 if bay 7 is used or ports 7 - 10 if bay 9 is used), and those circuits are rerouted to the extension module. This configuration does not offer any redundancy, but does maximize bandwidth to the extension module. If the extension module is installed in bay 5, you can use of the CIOv slot in the blade servers to

connect to I/O modules installed in bays 3 and 4 of the chassis. The total available external ports are six Ethernet port and six Fibre Channel ports.

Table 5-19 compares the number of available external ports in the configurations described previously.

Table 5-19 Comparing the number of external ports available in supported configurations

Configuration	10 Gb Ethernet ports	8 Gb Fibre Channel ports
Maximum redundancy and bandwidth	12	12
Maximum flexibility	16	6
Entry level	6	6

For more information, see the following sections:

- ▶ 2.7.10, “IBM Virtual Fabric 10 Gb Switch Module” on page 139
- ▶ 2.9.5, “QLogic Virtual Fabric Extension Module” on page 165
- ▶ 3.1.16, “QLogic 2-port 10 Gb Converged Network Adapter (CFFh)” on page 243

5.11.2 FCoCEE solution with the Brocade Converged 10 GbE Switch

The BladeCenter Brocade Converged 10 GbE FCoCEE solution requires the following components:

- ▶ IBM BladeCenter H or HT chassis.
- ▶ Supported Converged Network Adapter (CFFh) installed in supported blade servers
- ▶ Brocade Converged 10 GbE Switch Module

Figure 5-11 shows the Brocade FCoCEE solution topology for the BladeCenter H chassis.

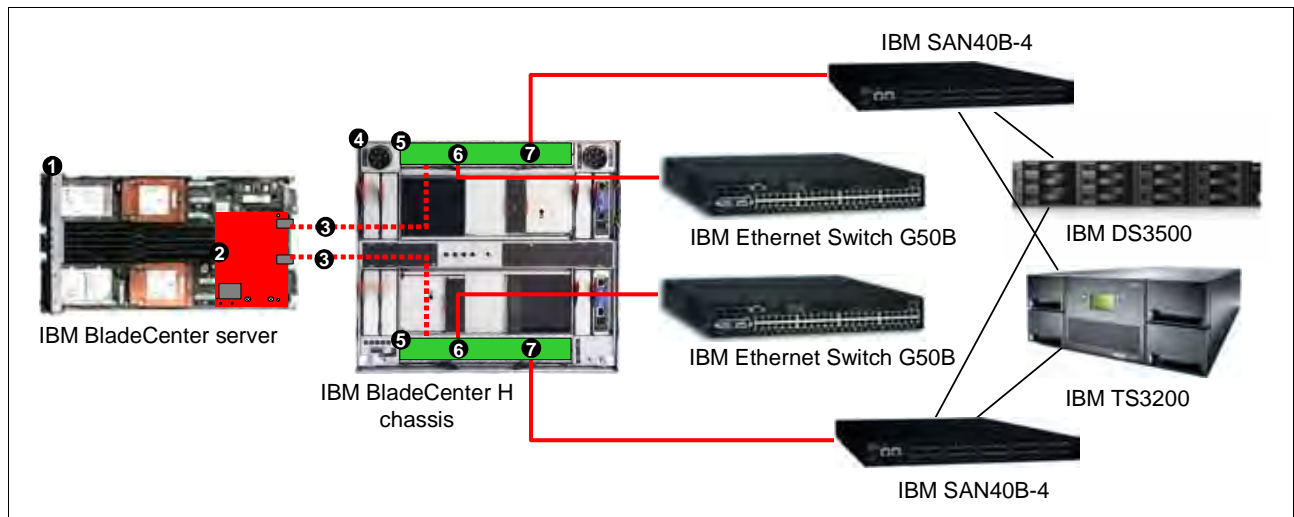


Figure 5-11 BladeCenter H-based Brocade FCoCEE solution topology

The components of the Brocade solution are as follows:

- ▶ IBM BladeCenter HS22 (or other supported blade server)
- ▶ Supported Converged Network Adapter
- ▶ Chassis internal connections from CNA ports to the Brocade Converged 10 GbE Switches
- ▶ BladeCenter H or HT chassis
- ▶ Brocade Converged 10 GbE Switch Module
- ▶ 10 Gb SFP+ optical transceivers to connect to the external Ethernet network
- ▶ 8 Gb SFP+ optical transceivers to connect to the external FC switches

For more information, see the following sections:

- ▶ 2.7.11, “Brocade Converged 10 GbE Switch Module” on page 145
- ▶ 3.1.17, “Brocade 2-port 10 GbE Converged Network Adapter (CFFh)” on page 244
- ▶ 3.1.16, “QLogic 2-port 10 Gb Converged Network Adapter (CFFh)” on page 243

5.11.3 Emulex adapters in x86 servers

Table 5-20 lists FCoE support using Fibre Channel targets for Emulex Virtual Fabric Adapters that are installed in supported IBM x86 blade servers.

Table 5-20 FCoE support using FC targets: Emulex Virtual Fabric Adapters

Ethernet adapter	BladeCenter I/O module	FC Forwarder (FCF)	SAN fabric	Operating systems	FC storage targets
<ul style="list-style-type: none"> Emulex Virtual Fabric Adapter II 90Y3550^a Emulex Virtual Fabric Adapter II Advanced, 90Y3566 Emulex 10GbE VFA II for HS23, 81Y3120^{b,c} Emulex 10GbE VFA Advanced II for HS23, 90Y9332^c Integrated 10GbE LOM (HS23), 94Y8550^d 	<ul style="list-style-type: none"> Brocade Converged 10GbE Switch Module, 69Y1909 (pNIC, vNIC2) 	<ul style="list-style-type: none"> Brocade VDX 6730^f 	<ul style="list-style-type: none"> None^e 	<ul style="list-style-type: none"> RHEL 5 RHEL 6 SLES 11 vSphere 4.1 vSphere 5.0 vSphere 5.1 Windows Server 2008 R2 Windows Server 2012 	<ul style="list-style-type: none"> DCS3700 DS3500 DS5000 DS8000 SVC V3500 V3700 V7000 XIV
			<ul style="list-style-type: none"> IBM B-type Brocade 		
	<ul style="list-style-type: none"> IBM 10Gb Ethernet Pass-thru Module, 46M6181 (pNIC, vNIC2) 	<ul style="list-style-type: none"> IBM G8264CS^g (NPIV) 	<ul style="list-style-type: none"> IBM B-type 		
			<ul style="list-style-type: none"> IBM B-type Cisco MDS 		
	<ul style="list-style-type: none"> Cisco Nexus 4001I Switch Module, 46M6071 (pNIC, vNIC2) 	<ul style="list-style-type: none"> Cisco Nexus 5548 Cisco Nexus 5596 	<ul style="list-style-type: none"> Cisco MDS 		
			<ul style="list-style-type: none"> Cisco MDS 		
	<ul style="list-style-type: none"> IBM Virtual Fabric 10Gb Switch Module, 46C7191 (pNIC, vNIC1, vNIC2) 	<ul style="list-style-type: none"> QLogic Virtual Fabric Extension Module, 46M6172^h 	<ul style="list-style-type: none"> Noneⁱ 	<ul style="list-style-type: none"> RHEL 5 RHEL 6 SLES 10 SLES 11 vSphere 4.1 vSphere 5.0 vSphere 5.1^j Windows Server 2008 R2 Windows Server 2012^j 	<ul style="list-style-type: none"> DS3500 DS5000 DS8000 SVC V3500 V3700 V7000 XIV Tape
			<ul style="list-style-type: none"> IBM B-type Cisco MDS 		
		<ul style="list-style-type: none"> Cisco Nexus 5010 Cisco Nexus 5020 Cisco Nexus 5548 Cisco Nexus 5596 	<ul style="list-style-type: none"> Cisco MDS^l 		
			<ul style="list-style-type: none"> IBM B-type^l 		
		<ul style="list-style-type: none"> Brocade VDX 6730^m 	<ul style="list-style-type: none"> IBM B-type^l 		

a. Requires an Advanced Upgrade, part number 49Y4265.

b. Requires an Advanced Upgrade, part number 90Y9350.

c. Not supported with the Brocade Converged 10GbE Switch Module.

d. Requires an Advanced Upgrade, part number 90Y9310.

e. When in Full Fabric mode, Brocade Converged 10GbE Switch Module supports direct connections to the DS3500 and DS5000 FC storage targets.

f. The configuration with the Brocade VDX supports RHEL 6, SLES 11, vSphere 5.0, and Windows Server 2008 R2.

g. The configuration with the IBM G8264CS supports RHEL 6, SLES 11, vSphere 5.0, vSphere 5.1, and Windows Server 2012.

h. Supports external IBM B-type, Brocade, or Cisco MDS SANs only in Intelligent Pass-thru (N-Port ID Virtualization (NPIV)) mode.

i. When in Full Fabric mode, the QLogic VFE supports direct connections to the FC storage targets.

j. Not supported in the configurations with the QLogic VFE.

k. Not supported in the configurations with the QLogic VFE operating in Full Fabric mode.

l. The configuration does not support tapes.

m. The configuration with the Brocade VDX supports RHEL 6, SLES 11, vSphere 5.0, and Windows Server 2008 R2.

5.11.4 Brocade and QLogic adapters in x86 servers

Table 5-21 lists FCoE support using Fibre Channel targets for Brocade and QLogic Converged Network Adapters that are installed in supported x86 blade servers.

Table 5-21 FCoE support using FC targets: Brocade and QLogic Converged Network Adapters

Ethernet adapter	BladeCenter I/O module	FC Forwarder (FCF)	SAN fabric	Operating systems	FC storage targets
▶ Brocade 2-port 10GbE Converged Network Adapter, 81Y1650	▶ Brocade Converged 10GbE Switch Module, 69Y1909	▶ Brocade 8000 ▶ IBM Converged Switch B32	▶ IBM B-type ▶ Brocade	▶ RHEL 5 ▶ RHEL 6 ▶ SLES 10 ▶ SLES 11 ▶ Windows Server 2008 R2 ▶ Windows Server 2012	▶ DS3500 ▶ DS5000
	▶ IBM 10Gb Ethernet Pass-thru Module, 46M6181	▶ Cisco Nexus 5010 ▶ Cisco Nexus 5020	▶ IBM B-type ▶ Brocade		
			▶ Cisco MDS		
▶ QLogic 2-port 10Gb Converged Network Adapter, 42C1830	▶ IBM 10Gb Ethernet Pass-thru Module, 46M6181	▶ Brocade 8000 ▶ IBM Converged Switch B32	▶ IBM B-type ▶ Brocade	▶ RHEL 5 ▶ RHEL 6 ▶ SLES 10 ▶ SLES 11 ▶ vSphere 4.1 ^b ▶ vSphere 5.0 ^b ▶ Windows Server 2008 / 2008 R2 ▶ Windows Server 2012 ^a	▶ DCS3700 ^b ▶ DS3500 ▶ DS5000 ▶ DS8000 ▶ SVC ▶ V3500 ^b ▶ V3700 ^b ▶ V7000 ▶ XIV ▶ Tape ^c
		▶ Cisco Nexus 5010 ▶ Cisco Nexus 5020	▶ Cisco MDS		
	▶ Brocade Converged 10GbE Switch Module, 69Y1909		▶ IBM B-type ▶ Brocade		
	▶ Cisco Nexus 4001I Switch Module, 46M6071	▶ Cisco Nexus 5010 ▶ Cisco Nexus 5020 ▶ Cisco Nexus 5548 ▶ Cisco Nexus 5596	▶ Cisco MDS		
	▶ IBM Virtual Fabric 10Gb Switch, 46C7191	▶ QLogic Virtual Fabric Extension Module, 46M6172 ^d	▶ IBM B-type ▶ Brocade ▶ Cisco MDS	▶ RHEL 5 ▶ RHEL 6 ^e ▶ SLES 10 ▶ SLES 11 ▶ vSphere 4.1 ^e ▶ Windows Server 2008 R2	▶ DS3500 ▶ DS5000 ▶ DS8000 ▶ SVC ▶ V7000 ▶ XIV ▶ Tape ^f
		▶ Cisco Nexus 5010 ▶ Cisco Nexus 5020 ▶ Cisco Nexus 5548 ▶ Cisco Nexus 5596	▶ Cisco MDS	▶ RHEL 5 ^g ▶ RHEL 6 ▶ SLES 10 ^g ▶ SLES 11 ▶ vSphere 4.1 ^g ▶ vSphere 5.0 ▶ Windows Server 2008 R2	▶ DCS3700 ▶ DS3500 ▶ DS5000 ▶ DS8000 ▶ SVC ▶ V3500 ▶ V3700 ▶ V7000 ▶ XIV
		▶ Brocade VDX 6730	▶ IBM B-type		

a. Supported in the configurations with the Ethernet Pass-thru Module and DS3500 and DS5000 FC storage targets.

b. Only supported in the configurations with the Cisco Nexus 4001I Switch Module.

c. Tapes are not supported with vSphere 4.1 and vSphere 5.0.

d. Supported in Intelligent Pass-thru (NPIV) mode only.

e. Not supported with the DS3500 and DS5000 FC storage targets.

f. Tapes are not supported with vSphere 4.1.

g. Not supported in the configurations with the Brocade VDX 6730.

Table 5-22 lists FCoE support using Fibre Channel targets for QLogic Virtual Fabric Converged Network Adapters that are installed in supported x86 blade servers.

Table 5-22 FCoE support using FC targets: QLogic Virtual Fabric Converged Network Adapters

Ethernet adapter	BladeCenter I/O module	FC Forwarder (FCF)	SAN fabric	Operating systems	FC storage targets
<ul style="list-style-type: none"> QLogic 10Gb Virtual Fabric Adapter, 00Y3332^a QLogic 10Gb Virtual Fabric CNA, 00Y5618 	<ul style="list-style-type: none"> IBM 10Gb Ethernet Pass-thru Module, 46M6181 	<ul style="list-style-type: none"> IBM G8264CS 	<ul style="list-style-type: none"> IBM B-type Brocade Cisco MDS 	<ul style="list-style-type: none"> RHEL 5 RHEL 6 SLES 11 vSphere 4.1^b vSphere 5.0 vSphere 5.1 Windows Server 2008 R2 Windows Server 2012 	<ul style="list-style-type: none"> DCS3700 DS3500 DS5000 DS8000 SVC V3500 V3700 V7000 XIV
		<ul style="list-style-type: none"> Cisco Nexus 5548 Cisco Nexus 5596 	<ul style="list-style-type: none"> Cisco MDS 		
	<ul style="list-style-type: none"> Brocade Converged 10GbE Switch Module, 69Y1909 		<ul style="list-style-type: none"> IBM B-type Brocade 		
	<ul style="list-style-type: none"> IBM Virtual Fabric 10Gb Switch, 46C7191 	<ul style="list-style-type: none"> Cisco Nexus 5548 Cisco Nexus 5596 	<ul style="list-style-type: none"> Cisco MDS 		

a. Requires QLogic 10Gb Virtual Fabric Advanced FoD Upgrade, part number 00Y5622.

b. Not supported in the configurations with the IBM 10Gb Ethernet Pass-thru Module, part number 46M6181.

5.11.5 Power Systems servers

Table 5-23 lists FCoE support using Fibre Channel targets for Converged Network Adapters that are installed in supported Power Systems blade servers.

Table 5-23 FCoE support using FC targets: Power Systems blade servers

Ethernet adapter	BladeCenter I/O module	FC Forwarder (FCF)	SAN fabric	Operating systems	FC storage targets
▶ QLogic 2-port 10Gb Converged Network Adapter, FC 8275	▶ Brocade Converged 10GbE Switch Module, 69Y1909 ^a		▶ None ^b	▶ AIX 5.3 ▶ AIX 6.1 ▶ AIX 7.1 ▶ RHEL 5 ▶ RHEL 6 ▶ SLES 10 ▶ SLES 11 ▶ VIOS 2.1 ^c ▶ VIOS 2.2	▶ DS3500 ▶ DS5000 ▶ DS8000 ▶ SVC ▶ V7000 ▶ XIV ▶ Tape
			▶ IBM b-type ▶ Brocade		
	▶ IBM 10Gb Ethernet Pass-thru Module, 46M6181	▶ IBM Converged Switch B32 ▶ Brocade 8000	▶ IBM b-type ▶ Brocade		
		▶ Cisco Nexus 5010 ▶ Cisco Nexus 5020	▶ Cisco MDS		
	▶ Cisco Nexus 4001I Switch Module, FC 2241	▶ Cisco Nexus 5010 ▶ Cisco Nexus 5020 ▶ Cisco Nexus 5548 ▶ Cisco Nexus 5596	▶ Cisco MDS	▶ AIX 6.1 ▶ AIX 7.1 ▶ RHEL 5 ^d ▶ RHEL 6 ▶ SLES10 ^{d,e} ▶ SLES 11 ▶ VIOS 2.2 ^e	▶ DS3500 ▶ DS5000 ▶ DS8000 ▶ SVC ▶ V3000 ▶ V7000 ▶ XIV ▶ Tape
	▶ IBM Virtual Fabric 10Gb Switch Module, FC 3248	▶ QLogic Virtual Fabric Extension Module, 46M6172 ^f	▶ None ^g	▶ AIX 5.3 ▶ AIX 6.1 ▶ AIX 7.1 ▶ SLES 10 ▶ SLES 11 ▶ RHEL 5 ▶ VIOS 2.2	▶ DS3500 ▶ DS5000 ▶ SVC
			▶ IBM b-type ▶ Brocade ▶ Cisco MDS		
		▶ Cisco Nexus 5010 ▶ Cisco Nexus 5020 ▶ Cisco Nexus 5548 ▶ Cisco Nexus 5596	▶ Cisco MDS	▶ AIX 6.1 ▶ AIX 7.1 ▶ RHEL 5 ^{d,h} ▶ RHEL 6 ▶ SLES10 ^{d,h} ▶ SLES 11 ▶ VIOS 2.2	▶ DS3500 ▶ DS5000 ▶ DS8000 ▶ SVC ▶ V3700 ▶ V7000 ▶ XIV ^h
		▶ Brocade VDX 6730	▶ IBM B-type		

a. Supports NPIV or Full Fabric mode when connected to the external SAN.

b. When in Full Fabric mode, Brocade Converged 10GbE Switch Module supports direct connections to the FC storage targets.

c. Not supported in the configurations with the Brocade Converged 10GbE Switch Module, part number 69Y1909.

d. Not supported with the DS3500 and DS5000 FC storage targets.

e. Not supported with the tapes.

f. Supports external IBM B-type, Brocade, or Cisco MDS SANs only in Intelligent Pass-thru (NPIV) mode.

g. When in Full Fabric mode, QLogic Virtual Fabric Extension Module supports direct connections to the FC storage targets.

h. Not supported in the configurations with the Brocade VDX 6730.

6

IBM zEnterprise Extension Model

IBM has extended the role of the mainframe by adding new infrastructure to the IBM zEnterprise 196 based on the IBM BladeCenter. This new infrastructure is called the zEnterprise BladeCenter Extension (zBX) Model 002.

The zBX brings the computing capacity of systems in a blade form-factor to the zEnterprise system. It is designed to provide a redundant hardware infrastructure that supports the multiplatform environment of the zEnterprise System in a seamless integrated way.

zBX provides, within the System zEnterprise infrastructure, a solution for running AIX workloads with IBM POWER based blades. Also, zBX provides an optimized solution for running Data Warehouse and Business Intelligence queries against IBM DB2® for IBM z/OS®, with fast and predictable response times, while retaining the data integrity, data management, security, availability, and other qualities of service of IBM System z®. This solution is known as the IBM Smart Analytics Optimizer.

The zBX components are configured, managed, and serviced the same way as the other components of the System z server. Despite the fact that the zBX processors might not be System z Processor Units (PUs), the zBX is handled by the System z management firmware known as the IBM zEnterprise Unified Resource Manager.

In this section, we introduce the zBX Model 002 and describe its hardware components. We also explain the basic concepts and building blocks for zBX connectivity.

For more information about the IBM zEnterprise 196, see *IBM zEnterprise 196 Technical Guide*, SG24-7833.

This chapter includes the following topics:

- ▶ 6.1, “zBX hardware description” on page 496
- ▶ 6.2, “zBX racks” on page 496
- ▶ 6.3, “Top-of-rack switches” on page 498
- ▶ 6.4, “zBX BladeCenter chassis” on page 499
- ▶ 6.5, “zBX blades” on page 500
- ▶ 6.6, “Power distribution unit” on page 502
- ▶ 6.7, “zBX management” on page 502
- ▶ 6.8, “zBX connectivity” on page 503

6.1 zBX hardware description

The zBX has a machine type of 2458-002 and is exclusive to the z196. It can host integrated multiplatform systems and heterogeneous workloads, with integrated advanced virtualization management. The zBX Model 002 is configured with the following key components:

- ▶ One to four standard 19-inch 42U IBM zEnterprise racks with the required network and power infrastructure
- ▶ One to eight BladeCenter chassis¹ with a combination of up to 112 different blades

Tip: The zBX design is based on the BladeCenter H chassis.

- ▶ A redundant infrastructure for fault tolerance and higher availability
- ▶ Management support through the z196 Hardware Management Console (HMC) and Support Element (SE)

The zBX can be ordered with a new z196 or as an MES to an existing z196. Either way, the zBX is treated as an extension to a z196 and cannot be ordered as a stand-alone feature.

Figure 6-1 shows a z196 with a maximum zBX configuration. The first rack (rack B) in the zBX is the primary rack where one or two BladeCenter chassis and four top of rack (TOR) switches are. The other three racks (C, D, and E) are expansion racks with one or two BladeCenter chassis each.




Figure 6-1 z196 with a maximum zBX configuration

6.2 zBX racks

The zBX Model 002 (2458-002) hardware is housed in up to four IBM zEnterprise racks. Each rack is 42U high (industry-standard 19 in.) and has four sidewall compartments to support installation of power distribution units (PDUs) and switches, with additional space for cable management.

¹ The IBM Smart Analytics Optimizer solution has a maximum of two zBX racks (B and C) and up to four BladeCenter chassis.

Figure 6-2 shows the rear view of a two-rack zBX configuration. The racks include the following items:

- ▶ Two TOR 1000BASE-T switches (rack B only) for the intranode management network (INMN)
- ▶ Two TOR 10 GbE switches (rack B only) for the intraensemble data network (IEDN)
- ▶ Up to two BladeCenter chassis in each rack with:
 - Up to 14 blades² (POWER7 or x86)
 - Advanced Management Modules (AMM)
 - Ethernet switch modules (ESM)
 - High speed switch (HSS) modules
 - 8 Gbps Fibre Channel switches for connectivity of client-supplied disks
- ▶ Power Distribution Units (PDUs)




Figure 6-2 zBX racks: rear view with BladeCenter chassis

² Depending on the IBM Smart Analytics Optimizer configuration, this number can be either 7 or 14 blades in the first BladeCenter chassis.

A zBX rack can support a maximum of two BladeCenter chassis. Each rack is designed for enhanced air flow and is shipped loaded with the initial configuration. It can be upgraded on site to a larger configuration.

The zBX racks are shipped with lockable standard non-acoustic doors and side panels. The following optional features are also available:

- ▶ The IBM Rear Door Heat eXchanger (FC 0540) reduces the heat load of the zBX emitted into the ambient air. The rear door heat eXchanger is an air-to-water heat exchanger that diverts the waste heat of the zBX to chilled water. The rear door heat eXchanger requires external conditioning units for its use.
- ▶ The IBM acoustic door (FC 0543) can be used to reduce the acoustical noise from the zBX.
- ▶ Height reduction (FC 0570) reduces the rack height to 36U high and accommodates doorway openings as low as 1832 mm (72.1 in.). Order this choice if you have doorways with openings less than 1941 mm (76.4 in.) high.

6.3 Top-of-rack switches

The four top-of-rack (TOR) switches are installed in the first rack (rack B). Adding expansion racks (racks C, D, and E) does not require additional TOR switches.

The TOR switches are located near the top of the rack and are mounted from the rear of the rack. From top to bottom are two 1000BASE-T switches for the intranode management network (INMN) and two 10 GbE switches for the intraensemble data network (IEDN).

A zBX can be managed only by one z196 through the INMN connections. Each VLAN-capable 1000BASE-T switch has 48 ports, which are reserved as follows:

- ▶ One port for each of the two bulk power hubs (BPH) on the controlling z196
- ▶ One port for each of the AMM and Ethernet switch modules (ESM), in each zBX BladeCenters chassis
- ▶ Two management ports, one port for each of the two IEDN 10 GbE TOR switches
- ▶ One port each for interconnecting the two switches

Both switches have the same connections to the corresponding redundant components (the BPH, AMM, ESM, and IEDN TOR switch) to avoid any single point of failure.

Physical placement tip: Although IBM provides a 26 m cable for the intranode management network connection, it is best to have the zBX installed next to or near the *controlling* z196 server for easy access to the zBX for service-related activities or tasks.

Each VLAN-capable 10 GbE TOR switch has 40 ports dedicated to the IEDN. The switch ports have the following connections:

- ▶ Up to eight ports are used for connections to the high-speed switch module of each BladeCenter chassis in the same zBX (as part of IEDN) to provide data paths to blades.
- ▶ Up to eight ports are used for OSA-Express3 10 GbE (LR or SR fiber) connections to the ensemble CPCs (as part of IEDN) to provide data paths between the ensemble CPCs and the blades in a zBX.
- ▶ Up to seven ports are used for zBX-to-zBX connections within a same ensemble (as part of the IEDN).

- ▶ Up to eight ports are used for the customer-managed data network. These connections are not part of IEDN and cannot be managed or provisioned by the Unified Resource Manager. The Unified Resource Manager recognizes them as migration connections and provides access control for their connection to the 10 GbE TOR switches.
- ▶ Two ports are used for interconnections between two switches as a failover path.
- ▶ One port is the management port connected to the INMN 1000BASE-T TOR switch.
- ▶ Two Direct Attached Cables (DAC) interconnect both switches.

For more information about the connectivity options for the INMN and the IEDN and the connectivity rules, see 6.8, “zBX connectivity” on page 503.

6.4 zBX BladeCenter chassis

The zBX BladeCenter chassis is based on the IBM BladeCenter H chassis and is preconfigured with the following components:

- ▶ Two AMMs are connected to the INMN through the 1000BASE-T TOR switches. Thus, firmware and configuration for the AMM is controlled by the SE on the controlling z196, with all the service management and reporting function of AMMs.
- ▶ Two 1000BASE-T (1 Gbps) Ethernet switch modules (ESM1 and ESM2) are installed in switch bays 1 and 2 in the chassis. Each ESM has 14 internal full-duplex Gigabit ports, one connected to each of the blade servers in the BladeCenter chassis, two internal full-duplex 10/100 Mbps ports connected to the AMM modules, and six 1000BASE-T copper RJ-45 connections for INMN connections to the TOR 1000BASE-T switches.

The ESM port 00 is connected to one of the 1000BASE-T TOR switches. As part of the INMN, configuration and firmware of ESM is controlled by the controlling z196 Support Element (SE).

- ▶ Two high-speed switch modules (HSS1 and HSS2) are installed in switch bays 7 and 9. The HSS modules provide 10 GbE uplinks to the 10 GbE TOR switches and 10 GbE downlinks to the blades in the chassis.

Port 01 is connected to one of the 10 GbE TOR switches. Port 10 is used to interconnect HSS in bays 7 and 9 as a failover path.

- ▶ Two 8 Gbps Fibre Channel (FC) switches (FC SW1 and 2) are installed in switch bays 3 and 4. Each switch has 14 internal ports reserved for the blade servers in the chassis, and six external Fibre Channel ports to provide connectivity to a client-supplied external disk for Smart Analytics Optimizer or blade solutions.

For the Smart Analytics Optimizer solutions, ports 00 and 15 of FC SW1 and FC SW2 in the first two Smart Analytics Optimizer chassis are used to connect the client-provided FC disk. Ports 16 - 19 are reserved for cascade connections from other Smart Analytics Optimizer chassis to access the Smart Analytics Optimizer FC disk.

For blade solutions, ports 00 and 15 of FC SW1 and FC SW2 can provide connections to client-provided FC storage disks for the blade servers in the same chassis.

6.5 zBX blades

The zBX Model 002 currently supports three types of blades:

- ▶ Up to 56 IBM Smart Analytics Optimizer blades

The number of Smart Analytics Optimizer blades depends on the Smart Analytics Optimizer solution size. Five Smart Analytics Optimizer solution sizes are available, corresponding to quantities of 7, 14, 28, 42, and 56 Smart Analytics Optimizer blades.

- ▶ Up to 56 HX5 models for zBX

Four configurations of HX5 blades are supported, depending on their memory size (see Table 6-1 on page 501).

- ▶ Up to 112 POWER7 blades

Three configurations of POWER blades are supported, depending on their memory size (Table 6-2 on page 502). The number of blades can be from 1 to 112.

For a zBX configuration, Smart Analytics Optimizer blades can be in nonadjacent chassis, but cannot share a chassis with other types of blades. However, a mixed installation of POWER7 and x86 blades in a chassis is supported.

zBX blade expansion cards

The 10 GbE CFFh expansion cards and 8 Gbps Fibre Channel CIOv expansion cards provide I/O connectivity to the IEDN, the INMN, or client-supplied FC storage disks.

IBM Smart Analytics Optimizer blade

IBM Smart Analytics Optimizer blade uses x86 technology and has these specifications:

- ▶ Two 2.93 GHz quad core processors
- ▶ Twelve 4 GB DIMMs, giving a total of 48 GB (24 GB for each quad core processor)
- ▶ A 32 GB solid-state drive (SSD)
- ▶ A 2-port 10 GbE CFFh expansion card for IEDN connection
- ▶ An 8 Gbps Fibre Channel CIOv expansion card for a client provided DS5020 connection

Tip: The IBM Smart Analytics Optimizer blade is based on the HS22.

IBM BladeCenter HX5 (7873) Models for zBx

The IBM BladeCenter HX5 is a scalable blade server designed to provide new levels of utilization, performance, and reliability for compute and memory-intensive workloads, such as databases, virtualization, business intelligence, modeling and simulation, and other enterprise applications.

Orders for the blades and operating systems can be made either through your IBM Business Partner, distributor, IBM marketing representative, or through the web.

Table 6-1 on page 501 shows the feature code and parts information for the IBM BladeCenter HX5 that works with zBX. The part numbers listed are for pricing and planning purposes only. Unsupported blade configurations do *not* work in the zBX.

Table 6-1 Feature code and parts information for the IBM BladeCenter HX5 that works with zBX

Blade	Config 0 (7873-AAx)	Config 1 (7873-ABx)	Config 2 (7873-ACx)	Config 3 (7873-ADx)
Processor	2 x Xeon 2.13 GHz 105 W (E7-2830 8C)	2 x Xeon 2.13 GHz 105 W (E7-2830 8C)	2 x Xeon 2.13 GHz 105 W (E7-2830 8C)	2 x Xeon 2.13 GHz 105 W (E7-2830 8C)
Memory	8x 8 GB	16x 8 GB	8x 8 GB + 8x 16 GB	16x 16 GB
Speed Burst Card	1	1	1	1
Internal Disks	2x 100 GB MLC SSD	2x 100 GB MLC SSD	2x 100 GB MLC SSD	2x 100GB MLC SSD
CFFh I/O Expansion	10 GbE	10 GbE	10 GbE	10 GbE
CIOv I/O Expansion	1	1	1	1

Do not make a selection for the following additional features available on the HX5 (7873):

- ▶ Memory mirroring (FC 9017).
- ▶ Memory sparing (FC 9016).
- ▶ Expansion Unit (FC 9295).
- ▶ USB compact flash (FC A1NP).
- ▶ SAS Connectivity Card - CIOv (FC 1593).
- ▶ Do not select integrated RAID - no Integrated Solid State Mirroring (7859) or Integrated Solid State Striping (7860). You must select No Internal RAID (FC 9012).
- ▶ An Operating system software preinstallation should not be ordered. Any preinstalled software is overlaid by the hypervisor.
- ▶ Do not order a hypervisor. Unified Resource Manager installs an integrated hypervisor on the blade when installed in the zBX, and any preinstalled software is overlaid by the hypervisor.
- ▶ No System Director software.
- ▶ No additional maintenance.
- ▶ Additional feature codes, such as FC#0785, FC#7493, and so on, are standard feature codes based on the configuration. They are not selected for ordering, but are base features for the blade.

The following operating systems are supported:

- ▶ For Linux
RHEL 5.5, 5.6 or RHEL 6 and Novell SLES 10 (SP4) or SLES 11 SP1 (Linux must be 64-bit support only)
- ▶ For Microsoft Windows
Microsoft Windows Server 2008 R2 Datacenter Edition 64-bit version (beginning 16 December 2011.)

POWER7 blade

The POWER7 blade is a single-width blade, which includes a POWER7 processor, with up to 16 DIMMs, and an HDD. The POWER7 blade supports 10 GbE connections to IEDN, and 8 Gbps FC connections to the client provide Fibre Channel storage through the FC SW1 and FC SW2 in the chassis.

The POWER7 blade is loosely integrated with a zBX, so that you can acquire supported blades through existing channels or IBM. The primary HMC and SE of the controlling z196 perform entitlement management for installed POWER7 blades on a one-blade basis. Supported configurations of POWER7 blades are listed in Table 6-2.

Table 6-2 Supported configuration of POWER7 blades

	Configuration 1	Configuration 2	Configuration 3
Processors	3.0 GHz@150 W	3.0 GHz@150 W	3.0 GHz@150 W
DIMMs	32 GB (8x 4 GB)	64 GB (16x 4 GB)	128 GB (16x 8 GB)
Internal disks	1x 300 GB HDD	1x 300 GB HDD	1x 300 GB HDD
CFFh I/O expansion	10 GbE	10 GbE	10 GbE
CIOv I/O expansion	8 Gbps FC	8 Gbps FC	8 Gbps FC

Tip: The POWER7 blade used in zBX is based on the PS701.

6.6 Power distribution unit

The power distribution units (PDUs) provide a connection to the main power source, the power connection to the intranode management network and intraensemble data network top of rack switches, and the power connection to the BladeCenter. The number of power connections needed is based on the zBX configuration. A rack contains two PDUs if one BladeCenter is installed or four PDUs if two BladeCenters are installed.

6.7 zBX management

One key feature of the zBX is its integration under the System z management umbrella. Thus, initial firmware installation and updates and patches follow the already familiar pattern of System z. The same reasoning applies to the configuration and definitions.

Similar to channels and processors, the SE has a view for the zBX blades. This view shows icons for each of the zBX component's objects, including an overall status (power, operational, and so on).

The following functions and actions are managed and controlled from the z196 HMC/SE:

- ▶ View firmware information for the BladeCenter and blades.
- ▶ Retrieve firmware changes.
- ▶ Change the firmware level.
- ▶ Back up and restore critical data.

zBX configuration data is backed up as part of zEnterprise 196 SE backup. It is restored with the replacement of a blade.

6.8 zBX connectivity

The types of LANs, each with redundant connections, that attach to the zBX are as follows:

- ▶ Intranode Management Network (INMN)
- ▶ Intraensemble Data Network (IEDN)
- ▶ Customer managed data network

The INMN is fully isolated and only established between the controlling z196 server and the zBX. The IEDN connects the zBX to a maximum of eight z196 servers.

Each z196 server must have a minimum of two connections to the zBX. The IEDN is used to connect a zBX to a maximum of seven other zBXs. The IEDN is a VLAN-capable network that allows enhanced security by isolating data traffic between virtual servers.

Figure 6-3 shows a high-level summary of the connectivity required for the zBX environment. It shows the z196 connections through two OSA-Express3 1000BASE-T features (CHPID type OSM) to the INMN TOR switches. The OSA-Express3 10 GbE features (CHPID type OSX) connect to the two IEDN TOR switches. Depending on workload requirements, any OSA-Express2 or OSA-Express features (CHPID type OSD) can connect to the customer-managed data network.

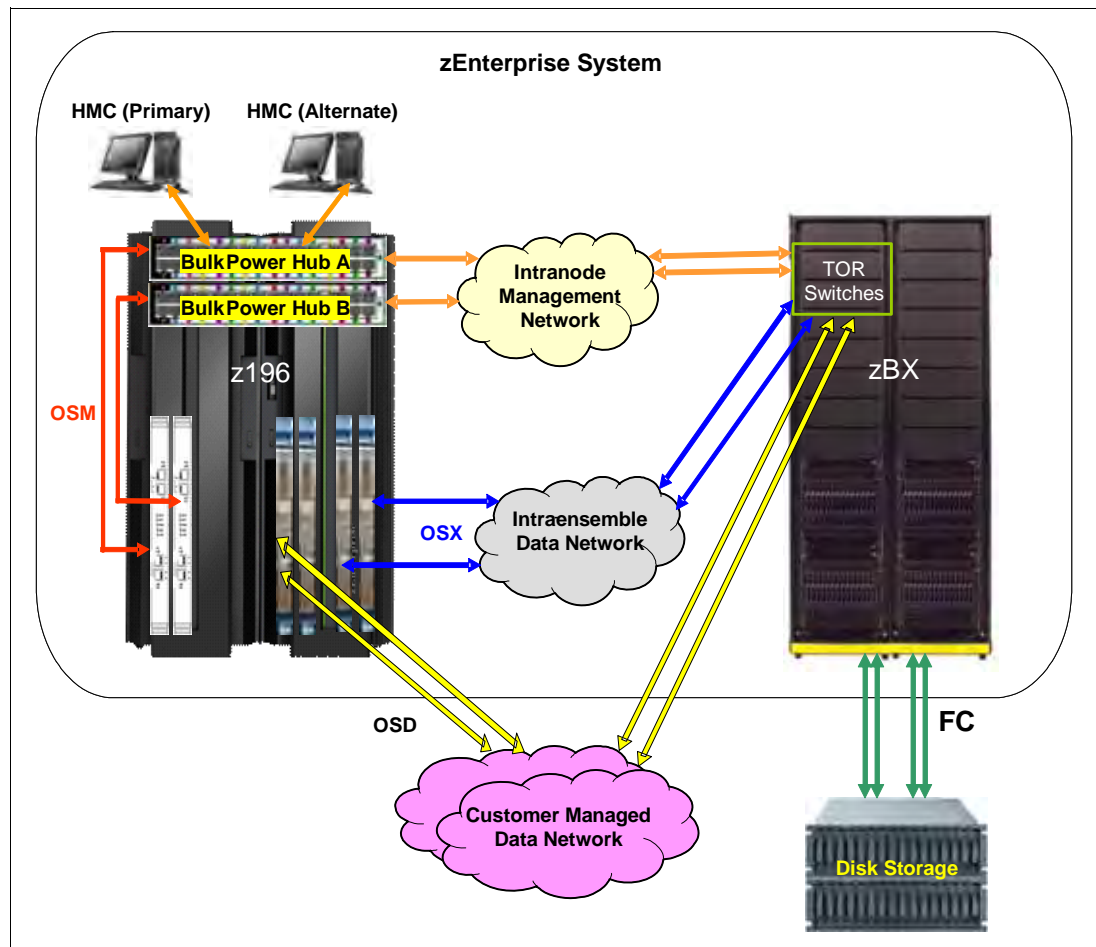


Figure 6-3 INMN, IEDN, customer-managed data network, and Fibre Channel connectivity

The Fibre Channel (FC) connections are required only between the zBX and the attached Fibre Channel disk or storage area network (SAN).

Cables: It is the client's responsibility to supply the cables for the IEDN, the customer-managed network, and the connection between the zBX and the FC disk.

The IEDN provides private and secure 10 GbE high-speed data paths between all elements of a zEnterprise ensemble (up to eight z196s with optional zBXs).

The zBX is managed by the HMC through the physically isolated INMN, which interconnects all resources of the zEnterprise System (z196 and zBX components).

Intranode management network

The scope of the intranode management network (INMN) is within an ensemble *node*. A node consists of a z196 CPC and its optional zBX. INMNs in different nodes are not connected to each other. The INMN connects the Support Element (SE) of the z196 to the hypervisor, optimizer, and guest management agents within the node.

Communication across the INMN is exclusively for enabling the Unified Resource Manager of the HMC to perform its various management disciplines (for example, performance management, network virtualization management, or energy management) for the node. The z196 connection to the INMN is achieved through the definition of a CHPID type OSM, which can be defined over an OSA-Express3 1000BASE-T Ethernet feature. There is also a 1 GbE infrastructure within the zBX.

Primary and alternative HMCs

The zEnterprise System Hardware Management Console (HMC) that has management responsibility for a particular zEnterprise ensemble is called a primary HMC. Only one primary HMC is active for an ensemble. This HMC has an alternative HMC to provide redundancy. It is not available for use until it becomes the primary HMC in a failover situation. To manage ensemble resources, the primary HMC for that ensemble must be used. A primary HMC can perform all HMC functions.

An HMC network configuration consists of the following elements:

- ▶ Primary and alternative HMCs must be connected to the same LAN segment.
The z196 can be on a different LAN segment.
- ▶ One pair of HMCs (primary and alternative) is required per ensemble:
 - The primary HMC can control up to eight z196s with one zBX each.
 - The two HMCs must be on the same LAN segment.
- ▶ The primary and alternative HMCs use Category 6 Ethernet cables:
The maximum distance is 100 meters, but it can be extended with an Ethernet switch.
- ▶ An additional 16-port Ethernet switch (FC 0070) can be ordered for the HMC network.
- ▶ A zEnterprise ensemble can be configured without zBX hardware (using only z196 nodes).

Figure 6-3 on page 503 shows the primary and alternative HMC configuration that connects with the two bulk power hubs (BPHs) in the z196 through a customer-managed management network. The 1000BASE-T TOR switches in the zBX are also connected to the BPHs in the z196.

Ports: All ports on the z196 BPH are reserved for specific connections. Any deviations or miscabling affects the operation of the z196 system.

Intraensemble data network

The intraensemble data network (IEDN) is the main application data path provisioned and managed by the Unified Resource Manager of the controlling z196. Data communications for ensemble-defined workloads flow over the IEDN between nodes of an ensemble.

All of the physical and logical resources of the IEDN are configured and managed by the Unified Resource Manager. The IEDN extends from the z196 through the OSA-Express3 10 GbE ports when defined as CHPID type OSX. The minimum number of OSA-Express3 10 GbE features is two per z196 server. Similarly, a 10 GbE networking infrastructure within the zBX is used for IEDN access.

The 10 GbE TOR switches use small form factor pluggable (SFP) optics for the external connections and direct attach cables (DAC) for connections as follows:

- ▶ TOR ports that are reserved for the z196 OSX IEDN connections use Small Form Factor Pluggable (SFP) Modules according to the zBX order:
 - FC 0632 LR SFP to FC3370 OSA-Express3 10 GbE LR
 - FC 0633 SR SFP to FC3371 OSA-Express3 10 GbE SR
- ▶ TOR ports that are reserved for IEDN to BladeCenter attachment use Direct Attached Cables (DAC) that are included with the zBX. These cables are hardwired 10 GbE SFP cables. The feature codes indicate the length of the cable:
 - FC 0626: 1 m for Rack B BladeCenters and IEDN to IEDN
 - FC 0627: 5 m for Rack C BladeCenter
 - FC 0628: 7 m for Racks D and E BladeCenters

The 10 GbE fiber optic cable types and their maximum distance are as follows:

- ▶ The client provides all IEDN cables (except for zBX internal connections).
- ▶ Multimode fiber:
 - 50 micron fiber at 2000 MHz-km: 300 meters (984 feet)
 - 50 micron fiber at 500 MHz-km: 82 meters (269 feet)
 - 62.5 micron fiber at 200 MHz-km: 33 meters (108 feet)
- ▶ Single mode fiber: 10 km (6.2 miles)

zBX storage connectivity

Each BladeCenter chassis in the zBX has two 20-port 8 Gbps Fibre Channel (FC) switch modules. Each switch has 14 internal ports and six shortwave (SX) external ports. The internal ports are reserved for the blades in the chassis.

Those external ports are used to connect the client-provided external Fibre Channel disk and for a cascaded connection between BladeCenter chassis to access the Smart Analytics Optimizer FC disk. Cascaded connections between BladeCenter chassis are only required for IBM Smart Analytics Optimizer solutions with three or four BladeCenter chassis.

Client-provided multi-mode LC duplex cables are used for FC disk connections to support speeds of 8 Gbps, 4 Gbps, or 2 Gbps. (A speed of 1 Gbps is not supported.) The maximum distance depends on the speeds, fiber diameter, and signal frequency.

IBM blade storage connectivity

IBM blades use two ports in both FC switch modules of the BladeCenter chassis to connect either directly to FC disk storage or through a SAN switch to FC disk storage (Figure 6-4). A cascaded connection between BladeCenter chassis in the zBX is not necessary.




Figure 6-4 BladeCenter chassis storage connectivity options

The client provides all cables, FC disk storage, and SAN switches. It is also the client's responsibility to configure and cable the FC disk storage.

Disk storage: FC disk storage cannot be shared with the IBM Smart Analytics Optimizer solution.

IBM Smart Analytics Optimizer storage connectivity

In an IBM Smart Analytics Optimizer configuration that has one or two BladeCenters, both FC switch modules of each BladeCenter chassis connect directly to FC disk storage. The number of FC connections from each BladeCenter is four.

Requirements: The IBM Smart Analytics Optimizer requires an IBM System Storage DS5020. The DS5020 must be attached directly to the BladeCenter chassis.

In an IBM Smart Analytics Optimizer configuration that has three or four BladeCenters, the first and second BladeCenter chassis are connected to the DS5020.

The third and fourth BladeCenter chassis must have cascaded connections between the first and second BladeCenter chassis by using the respective ports of their FC switch modules.

The foregoing rules are also applied when a zBX has a mix of Smart Analytics Optimizer and POWER7 BladeCenter chassis, but the cascaded connections are only required between the IBM Smart Analytics Optimizer BladeCenter chassis.

Supported FC disk storage

Supported FC disk types and vendors with IBM blades are listed on the IBM System Storage Interoperation Center (SSIC) website:

<http://www.ibm.com/systems/support/storage/config/ssic/>

For an IBM Smart Analytics Optimizer solution, a client-supplied IBM System Storage DS5020 with the appropriate configuration and fiber optic cables is required.

The FC 7802 DS5020 Linux/intel license feature is needed to support the Smart Analytics Optimizer.

Table 6-3 lists the storage specifications for each Smart Analytics Optimizer solution size.

Table 6-3 Storage specifications (DS5020) based on Smart Analytics Optimizer solution size

Solution size	A1-7	A1-14	A1-28	A1-42	A1-56
Number of Fibre Channel ports	4	4	8	8	8
Number of disks^a	16	16	16	32	32

a. Each disk is a 1000 GB/7.2 K SATA II E-DDM.

The DS5020 needs to be prepared and connections to the zBX must be in place before it can be used by the Smart Analytics Optimizer application.

Based on the data mart definitions provided by the DB2 administrator, the IBM Smart Analytics Optimizer application sends data from DB2 for z/OS to the zBX. The zBX firmware compresses the data and then stores it in the DS5020. From the DS5020, the compressed data is read into blade memory during blade initialization. Processing for DB2 queries is done entirely in blade memory.

The DS5020 also stores the zBX firmware, which is provided by z196 MCLs and loaded from the z196 HMC/SE. After the firmware is stored, each blade boots from the DS5020.

zBX connectivity example: Single node ensemble with a zBX

Figure 6-5 shows a single node ensemble with a zBX. The necessary components include the controlling z196 CPC1 and the attached zBX, switches, and Fibre Channel disk storage. For simplicity, we do not show the redundant connections in the configuration example.




Figure 6-5 Single node ensemble with zBX

The components are as follows:

- ▶ Client-provided management network:
 - IBM supplies a 50 ft. Ethernet RJ-45 cable with the 1000BASE-T (1 GbE) switch (FC 0070).
 - The switch connects to the reserved client network ports in the z196 Z29BPS11/PS31-J02. A second switch connects to Z29BPS11/PS31-J01.
- ▶ Intranode management network:
 - Two CHPIDs from two different OSA-Express3 1000BASE-T features are configured as CHPID type OSM.
 - IBM supplies two 3.2 m Ethernet Category 6 cables from the OSM CHPIDs (ports) to z196 Bulk Power Hub (BPH) Z29BPS11/PS31-J07. (This connection is a z196 internal connection supplied with feature code 0025.)
- ▶ Intranode management network - extension

IBM supplies two 26 m Category 5 Ethernet cables (chrome gray plenum rated cables) from zBX Rack B INMN-A/B switches port J47 to z196 Bulk Power Hub (BPH) Z29BPS11/PS31-J06.

- ▶ Intraensemble data network:
 - Two ports from two different OSA-Express3 10 GbE (SR Short Reach or LR Long Reach) features are configured as CHPID type OSX.
 - The client supplies the fiber optic cables (single mode or multimode).
- ▶ 8 Gbps Fibre Channel attached disk storage:
 - The client supplies all Fibre Channel cables (multimode) from the zBX to the attached disk storage.
 - The client supplies switches if necessary.
 - The client is responsible for the configuration and management of the FC-attached disk storage.

See the following sources for more information:

- ▶ For installation details:
 - *IBM zEnterprise BladeCenter Extension Model 002 Installation Manual for Physical Planning*, GC27-2611
 - *IBM zEnterprise BladeCenter Extension Model 002 Installation Manual*, GC27-2610
- ▶ For details about the zBX solution components, see *IBM zEnterprise 196 Technical Guide*, SG24-7833.

Abbreviations and acronyms

AC	alternating current	DP	dual processor
ACL	access control list	DSA	Dynamic System Analysis
AES	Advanced Encryption Standard	DSCP	Differentiated Services Code Point
AMM	Advanced Management Module	DSM	disk storage module
API	application programming interface	DTP	Dynamic Trunking Protocol
ARP	Address Resolution Protocol	DVMRP	Distance Vector Multicast Routing Protocol
AS	Australian Standards	ECC	error checking and correcting
ASIC	application-specific integrated circuit	EIGRP	Enhanced Interior Gateway Routing Protocol
BASP	Broadcom Advanced Server Program	EMEA	Europe, Middle East, Africa
BBI	browser-based interface	EOT	enhanced object tracking
BCM	Broadcom	ESD	electrostatic discharge
BGP	Border Gateway Protocol	ESM	Ethernet switch modules
BIOS	basic input/output system	ETSI	European Telecommunications Standard Industry
BMC	Baseboard Management Controller	FAN	Fabric Address Notification
BOFM	BladeCenter Open Fabric Manager	FC	Fibre Channel; feature code
BPDU	bridge protocol data unit	FCP	Flow Control Packet
BSMP	blade system management processor	FCSM	Fibre Channel Switch Module
BTU	British Thermal Unit	FDX	full duplex
CD-ROM	compact-disc read-only memory	FSB	front-side bus
CDP	Cisco Discovery Protocol	FTP	File Transfer Protocol
CLI	command-line interface	GB	gigabyte
CNA	Cisco Network Assistance	GUI	graphical user interface
CNS	Cisco Network Services	HA	high availability
COG	configuration and option guide	HBA	host bus adapter
CPM	Copper Pass-thru Module	HCA	host channel adapter
CPU	central processing unit	HD	high definition
CRU	customer-replaceable units	HDD	hard disk drive
CTS	clear to send	HH	half high
DASD	direct access storage device	HPC	high performance computing
DC	domain controller	HS	hot swap
DDM	deployment and management	HSFF	high-speed form factor
DDR	double data rate	HSRP	Hot Standby Routing Protocol
DHCP	Dynamic Host Configuration Protocol	HT	hyper-threading
DIMM	dual inline memory module	HTTP	Hypertext Transfer Protocol
DIP	destination IP	I/O	Input/output
DMAC	destination MAC address	IB	InfiniBand
DNS	Domain Name System	IBM	International Business Machines
		ICMP	Internet Control Message Protocol

ICPM	Intelligent Copper Pass-thru Module	MSDP	Multicast Source Discovery Protocol
ID	identifier	MSIM	Multi-Switch Interconnect Module
IEC	International Electro-technical Commission	MSTP	Multiple Spanning Tree Protocol
IEEE	Institute of Electrical and Electronics Engineers	MTM	machine-type-model
IGMP	Internet Group Management Protocol	MVR	multicast VLAN registration
IGRP	Interior Gateway Routing Protocol	NAT	Network Address Translation
IM	integrated mirroring	NDCLA	Non-Disruptive Code Load Activation
IME	integrated mirroring enhanced	NEBS	Network Equipment Building System
IOS	Internetwork operating system	NGN	next-generation network
IP	Internet Protocol	NIC	network interface card
IPM	Intelligent Pass-thru Module	NMI	non-maskable interrupt
IPMI	Intelligent Platform Management Interface	NOS	network operating system
IPTV	Internet Protocol Television	NPIV	N_Port ID Virtualization
IRDP	ICMP Router Discovery Protocol	NSF	nonstop forwarding
IS	integrated striping	NTP	Network Time Protocol
ISL	Inter-Switch Link	OBFL	onboard failure logging
ISP	Internet service provider	ODPA	On-Demand Port Activation
IT	information technology	OPM	Optical Pass-thru Module
ITS	IBM Integrated Technology Services	OS	operating system
ITSO	International Technical Support Organization	OSPF	Open Shortest Path First
IVM	Integrated Virtualization Manager	PBR	policy-based routing
KB	kilobyte	PC	personal computer
KVM	keyboard video mouse	PCI	Peripheral Component Interconnect
LACP	Link Aggregation Control Protocol	PDF	Portable Document Format
LAN	local area network	PDU	power distribution unit
LED	light emitting diode	PFA	Predictive Failure Analysis
LLDP	Link Layer Discovery Protocol	POST	power-on self-test
LPAR	logical partitions	PVST	Per-VLAN Spanning Tree
LR	long range	PXE	Preboot Execution Environment
LTO	linear tape-open	RADIUS	Remote Authentication Dial In User Service
LUN	logical unit number	RAID	Redundant Array of Independent Disks
MAC	media access control	RAM	random access memory
MAN	metropolitan area network	RAS	reliability, availability, and serviceability
MB	megabyte	RDIMM	Registered DIMM
MDS	multilayer datacenter switch	RDM	Remote Deployment Manager
MM	Management Module	RDMA	Remote Direct Memory Access
MMF	multi mode fiber	RHEL	Red Hat Enterprise Linux
MP	multiprocessor	RIP	Routing Information Protocol
MPI	Message Passing Interface	RMCP	Remote Management Control Protocol

RMON	remote Monitoring	TCP/IP	Transmission Control Protocol/Internet Protocol
RP	route processor	TDR	Time Domain Reflector
RPF	reverse path forwarding	TFTP	Trivial File Transfer Protocol
RSA	Remote Supervisor Adapter	TOE	TCP offload engine
RSCN	Registered State Change Notification	TX	transmit
RSTP	Rapid Spanning Tree Protocol	UDLD	Unidirectional Link Detection
SAN	storage area network	UDP	User Datagram Protocol
SAS	Serial Attached SCSI	URL	Uniform Resource Locator
SBB	Sales Building Block	USB	Universal Serial Bus
SCM	Storage Configuration Manager	UTF	Universal Telco Frame
SCSI	Small Computer System Interface	UTP	unshielded twisted pair
SDD	Subsystem Device Driver	VBS	Virtual Blade Switch
SDK	software developer kit	VIOS	Virtual I/O Server
SDR	single data rate	VLAN	virtual LAN
SFF	small form factor	VLP	very low profile
SFP	small form-factor pluggable	VM	virtual machine
SIO	Storage and I/O	VOIP	voice over Internet Protocol
SIP	Source IP	VPD	vital product data
SLB	Server Load Balancing	VQP	VLAN Query Protocol
SLES	SUSE Linux Enterprise Server	VRRP	Virtual Router Redundancy Protocol
SM	Subnet Manager	VSAN	Virtual Storage Area Network
SMAC	Source MAC address	VT	Virtualization Technology
SMI-S	Storage Management Initiative - Specification	VTP	VLAN Trunking Protocol
SMP	symmetric multiprocessing	WAN	wide area network
SNMP	Simple Network Management Protocol	WCCP	Web Cache Communication Protocol
SOL	Serial over LAN	WRR	weighted round-robin
SPORE	ServerProven Opportunity Request for Evaluation	WTD	weighted tail drop
SR	short range	WWN	worldwide name
SRP	Storage RDMA Protocol	XM	extended memory
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		

Related publications

We consider the publications that we list in this section suitable for a more detailed discussion of the topics that we cover in this book.

IBM Redbooks publications

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 website:

ibm.com/redbooks

- ▶ *An Introduction to Fibre Channel over Ethernet, and Fibre Channel over Convergence Enhanced Ethernet*, REDP-4493
- ▶ *IBM BladeCenter Interoperability Guide*, REDP-BCIG
- ▶ *IBM BladeCenter PS700, PS701, and PS702 Technical Overview and Introduction*, REDP-4655
- ▶ *IBM BladeCenter PS703 and PS704 Technical Overview and Introduction*, REDP-4744
- ▶ *IBM BladeCenter Virtual Fabric Solutions*, SG24-7966
- ▶ *IBM eX5 Portfolio Overview: IBM System x3850 X5, x3950 X5, x3690 X5, and BladeCenter HX5*, REDP-4650
- ▶ *Implementing the Brocade Access Gateway for IBM BladeCenter*, REDP-4343
- ▶ *Implementing the IBM BladeCenter S Chassis*, SG24-7682
- ▶ *Implementing the QLogic Intelligent Pass-thru Module for IBM BladeCenter*, REDP-4342

In addition, see the following information at the IBM Redbooks website:

- ▶ IBM Redbooks Product Guides for IBM BladeCenter
<http://www.redbooks.ibm.com/Redbooks.nsf/portals/BladeCenter?Open&page=pg>
- ▶ IBM Redbooks - BladeCenter
<http://www.redbooks.ibm.com/portals/bladecenter>

Product publications

This section lists BladeCenter product publications.

- ▶ IBM BladeCenter Information Center
<http://publib.boulder.ibm.com/infocenter/bladectr/documentation>
- ▶ IBM BladeCenter - Blade Server Base Specifications
 - *Blade Server Base Specification for Processor Blade Subsystems*
 - *Blade Server Base Specification for Switch Module Subsystems*
 - *Blade Server Base Specification for I/O Expansion Cards*
 - *Blade Server Base Spec. for SERDES High-Speed Elec. Signaling*
 - *Blade Server Base Specification for VPD*

These publications can be found at:

http://www.ibm.com/systems/bladecenter/resources/open_specs.html

Online resources

These websites are also relevant as further information sources.

Product information

- ▶ Benchmarks for BladeCenter:
<http://www.ibm.com/servers/eserver/bladecenter/benchmarks/>
- ▶ IBM Offering Information:
<http://www.ibm.com/common/ssi>
- ▶ IBM ServerProven - BladeCenter
<http://www.ibm.com/systems/info/x86servers/serverproven/compat/us/eserver.html>
- ▶ System Storage Interoperation Center (SSIC)
<http://www.ibm.com/systems/support/storage/ssic>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



Redbooks

IBM BladeCenter Products and Technology

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5" spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(->Hide:)>Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine:fm still open and **File>Import>Formats** the Conditional Text Settings (ONLY!) to the book files.



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**

Blade servers captured industry focus because of their modular design, which can reduce costs with a more efficient usage of valuable floor space, and their simplified management, which can speed up such tasks as deploying, reprovisioning, updating, and troubleshooting hundreds of blade servers. By integrating resources and sharing key components, costs are reduced and availability is increased.

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

This eleventh edition includes new products that were announced in the first and second quarters of 2013. This book is for clients, IBM Business Partners, and IBM employees who want to know the details about IBM BladeCenter chassis, servers, and options.

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