



QLogic 12800 InfiniBand Switches for the IBM Intelligent Cluster

Product Guide

The QLogic 12800 InfiniBand switch is a new class of data center InfiniBand interconnect that links HPC resources using a scalable, 40 Gbps, low latency fabric. Customers can tailor the modular product to meet present and future needs. The QLogic 12800 is a series of products that deliver unmatched high-speed networking features and functions.

Figure 1 shows the 12800 switch modules.



Figure 1. QLogic 12800-360 (left) and 12800-180 (right) InfiniBand switches

Did you know?

QLogic 12800 InfiniBand switch series is a part of IBM Intelligent Cluster solution (formerly IBM Systems Cluster 1350). IBM Intelligent Cluster is your key to a fully integrated HPC solution. IBM clustering solutions include servers, storage, and industry-leading OEM interconnects that are factory-integrated, fully tested, and delivered to your door, ready to plug into your data center, all with a single point of contact for support.

Part number information

Table 1 shows the part numbers to order these modules and additional options for them.

Table 1. IBM part numbers for ordering

Description	IBM part number
QLogic 12800-180 QDR InfiniBand Switch Bundle	0449-022
QLogic 12800-360 QDR InfiniBand Switch Bundle	0449-023
QLogic 12800 QDR Single Spine Module	0449-024
QLogic 12800 QDR Double Spine Module	0449-025
QLogic 12800 Management Module	49Y8543
QLogic 12800 24-port 4X QDR Ultra-HD InfiniBand Leaf Module	0449-026
QLogic 12800 18-port 4X QDR Ultra-HP InfiniBand Leaf Module	0449-027
QLogic 12800 Hot-swap Power Supply	49Y8504

The QLogic 12800-180 and 12800-360 QDR InfiniBand director switches, based on QLogic TrueScale ASIC technology, deliver the next evolution in switch fabric performance for High Performance Computing (HPC) environments.

These switches deliver high port density and low power per port:

- The 12800-180 provides 324 QDR ports using 18-port Ultra-High Performance (UHP) leaf modules or 432 QDR/DDR ports using 24-port Ultra-High Density (UHD) leaf modules in a 14U chassis.
- The 12800-360 provides 648 QDR ports using 18-port Ultra-High Performance (UHP) leaf modules or 864 QDR/DDR ports using 24-port Ultra-High Density (UHD) leaf modules in a 29U chassis.

The high-availability 12800 design includes hot swappable InfiniBand spine and leaf modules, fully redundant power and cooling, and redundant management processors supporting chassis management (CLI and GUI), as well as embedded InfiniBand Subnet Managers (SMs). The 12800 also supports advanced QLogic fabric features including adaptive routing, Virtual Fabrics (vFabric), Quality of Service (QoS), and management wizards for automated installation, configuration, and monitoring to maximize operational efficiency.

Benefits

The QLogic 12800 InfiniBand switch offers the following benefits:

- **Low Latency:** The QLogic 12800 provides scalable, predictable low latency, even at 90% traffic use. Predictable latency means HPC applications can be scaled easily without worrying about diminished cluster performance or costly system tuning efforts.
- **Flexible Partitioning:** The QLogic 12800 advanced design is based on an architecture that provides comprehensive virtual fabric partitioning capabilities that enable the InfiniBand fabric to support the evolving requirements of an organization. The TrueScale architecture, together with IFS, allows the fabric to be shared by mission critical applications while delivering maximum bandwidth utilization.
- **Modular Design:** InfiniBand port, power, cooling, and management modules are common in the series, giving customers the flexibility to deploy and grow HPC environments in a cost-effective fashion.
- **Investment Protection:** The QLogic 12800 adheres to the InfiniBand Trade Association version 1.2 specification, ensuring the ability to interoperate with all other IBTA-compliant devices.
- **Highly Reliable:** This system is designed for high availability with features that include: port-to-port and module-to-module failover, non-disruptive firmware upgrades, component level diagnostics and alarming, and both in-band and out-of-band management.
- **Easy to Manage:** The 12800 uses QLogic's advanced IFS software for quicker installation and configuration. IFS has advanced tools to verify fabric configuration, topology, and performance. Faults are automatically isolated to the component level and reported.
- **Simple Installation and Configuration:** Using the installation and configuration wizards contained in the IFS package allows end users to bring up fabrics in days instead of weeks.
- **Power Optimized:** Maximum performance is delivered with minimal power and cooling requirements as part of QLogic's Star Power commitment to developing green solutions for the data center.

Features and specifications

The QLogic 12800 InfiniBand switches include the following features and functions:

- Between 18 and 864 ports of InfiniBand QDR (40Gbps) performance with support for DDR and SDR
 - 40/20/10-Gbps auto-negotiation links
 - Supports Quad Small Form Factor Pluggable (QSFP) optical and copper cable specifications
- TrueScale architecture, with scalable, predictable low latency
 - Scales to 51.8-Tbps aggregate bandwidth
 - Switching latency: 140 to 420 ns
- Multiple Virtual Lanes (VLs) per physical port
 - Virtual lanes: 8 plus 1 management
- Maximum MTU size: 4096 bytes
- Maximum multicast table size: 1024 entries
- Supports virtual fabric partitioning
- Fully-redundant system design
- Option to use Ultra High Density (UHD) leafs for maximum connectivity or Ultra High Performance (UHP) leafs for maximum performance
 - UHP Module: 18 QDR ports
 - UHD Module: 24 QDR ports
- Integrated chassis management capabilities for installation, configuration, and ongoing monitoring
- Optional InfiniBand Fabric Suite (IFS) management solution that provides expanded fabric views

and fabric tools

- Complies with InfiniBand Trade Association (IBTA) version 1.2 standard

The QLogic 12800-360 has the following specifics:

- 18 to 864 ports
- 51.8 Tbps capacity
- Supports up to 36 leaf modules

The QLogic 12800-180 has the following specifics:

- 18 to 432 ports
- 25.9 Tbps switching capacity
- Supports up to 18 leaf modules

The QLogic 12800 InfiniBand switch family supports the following management methods:

- Command Line Interface
- Optional external server-based InfiniBand-compliant subnet manager
- Optional embedded fabric management
- IBTA-compliant SMA, PMA, and BMA
- Chassis management GUI
- SNMP support
- Access methods
 - 10/100 Ethernet Base-T
 - Serial port (RS-232 with DB9)

The following LEDs are located on QLogic 12800 InfiniBand switches:

- 2 per leaf module plus 1 per InfiniBand port
- 2 per single-spine module
- 4 per dual-spine module
- 3 per power supply module
- 8 per management module
- 3 per serial port/Ethernet module

Supported System x and BladeCenter servers, adapters, and cables

The QLogic 12800 InfiniBand switch supports IBM Intelligent Cluster nodes with InfiniBand adapters installed in the System x or BladeCenter servers. The tables below show core solution components from the IBM Intelligent Cluster Component Guide Release 10A.

QDR InfiniBand solution with System x servers and QLogic 12800 series switches

Table 2 shows core components to create full-speed QDR (40 Gbps) InfiniBand solution using QLogic 12800 InfiniBand switches and System x servers.

Table 2. QDR InfiniBand solution core components for System x

Description	IBM part number or machine type
InfiniBand Host Channel Adapters	
Mellanox ConnectX-2 VPI Single-port QSFP QDR IB/10GbE PCI-E 2.0 HCA	81Y1531
Mellanox ConnectX-2 VPI Dual-port QSFP QDR IB/10GbE PCI-E 2.0 HCA	81Y1535
Mellanox ConnectX Single-Port 4X QDR InfiniBand x8 PCI-E 2.0 HCA	46M2203
Mellanox ConnectX Dual-Port 4X QDR InfiniBand x8 PCI-E 2.0 HCA	46M2199
Compute Nodes	
System x3450	7948
System x3550 M3	7944
System x3650 M3	7945
System x3655	7943
System x3455	7940, 7941
System x3755	7163
System x3850 X5	7145
System x iDataPlex dx360 M2	7321, 7323
System x iDataPlex dx360 M3	6391
Cables	
3m QLogic Optical QDR InfiniBand QSFP Cable	Feature code 1767
10m QLogic Optical QDR InfiniBand QSFP Cable	Feature code 1768
30m QLogic Optical QDR InfiniBand QSFP Cable	Feature code 1769
3m IBM Optical QDR InfiniBand QSFP Cable	Feature code 5989
10m IBM Optical QDR InfiniBand QSFP Cable	Feature code 5990
30m IBM Optical QDR InfiniBand QSFP Cable	Feature code 5991
0.5m QLogic Copper QDR InfiniBand QSFP 30AWG Cable	Feature code 3725
1m QLogic Copper QDR InfiniBand QSFP 30AWG Cable	Feature code 3726
3m QLogic Copper QDR InfiniBand QSFP 28AWG Cable	Feature code 3727

DDR InfiniBand (20 Gbps) solution with BladeCenter servers and QLogic 12800 series switches

Table 4 shows core components to create a full-speed DDR (20 Gbps) InfiniBand solution using QLogic 12800 InfiniBand switches and System x servers.

Table 4. DDR InfiniBand solution core components for BladeCenter

Description	IBM part number or machine type
InfiniBand Host Channel Adapters	
4X DDR InfiniBand CFFh Expansion Card for BladeCenter	43W4423
Chassis and Compute Nodes	
BladeCenter H chassis	8852
BladeCenter HS21 server	8853
BladeCenter HS21 XM server	7995
BladeCenter HS22 server	7870
BladeCenter LS22 server	7901
BladeCenter LS42 server	7902
BladeCenter JS22 server	7998
BladeCenter QS22 server	0793
Pass-thru Modules and Cables	
4X InfiniBand Pass-thru HSSM for BladeCenter*	43W4419
3m QLogic Optical DDR InfiniBand QSFP-to-CX4 Cable	59Y1908
10m QLogic Optical DDR InfiniBand QSFP-to-CX4 Cable	59Y1912
30m QLogic Optical DDR InfiniBand QSFP-to-CX4 Cable	59Y1916

* The 4X InfiniBand Pass-thru high-speed switch module supports DDR speeds but does not support QDR speeds.

Related publications

For more information, see the following documents:

- IBM US Announcement Letter (May 18, 2010 - Refresh 10A):
<http://ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS110-140>
- Publications for the Cluster 1350
<http://ibm.com/support/entry/portal/docdisplay?Indocid=MIGR-5080504>
- QLogic 12800-360 product page
<http://www.qlogic.com/Products/HPCNetworking/InfiniBandSwitches/Pages/12800360.aspx>
- QLogic 12800-180 product page
<http://www.qlogic.com/Products/HPCNetworking/InfiniBandSwitches/Pages/12800180.aspx>

Related product families

Product families related to this document are the following:

- [Top-of-Rack Switches](#)

Notices

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

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO 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 jurisdictions 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. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result 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.

© Copyright Lenovo 2016. All rights reserved.

This document, TIPS0722, was created or updated on May 24, 2010.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<http://lenovopress.com/TIPS0722>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <http://lenovopress.com/TIPS0722>.

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <http://www.lenovo.com/legal/copytrade.html>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

BladeCenter®

iDataPlex®

System x®

Intelligent Cluster™

X5™

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