

# Planning for Converged Fabrics

## The Next Step in Data Center Evolution

Best practices when migrating to Fibre Channel Over Ethernet

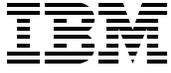
Deployment planning for Layer 1 physical infrastructure

Exploiting products and technology from IBM and Panduit



Srihari Angaluri  
Martin Bachmaier  
Panduit Corp  
David Watts





International Technical Support Organization

**Planning for Converged Fabrics: The Next Step in  
Data Center Evolution**

May 2010

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

**First Edition (May 2010)**

This edition applies to IBM and Panduit products.

This document created or updated on November 30, 2012.

**© Copyright International Business Machines Corporation 2010. 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

<b>Notices</b> .....	v
Trademarks .....	vi
<b>Preface</b> .....	vii
The team who wrote this paper .....	vii
Now you can become a published author, too! .....	viii
Comments welcome .....	viii
Stay connected to IBM Redbooks .....	ix
<b>Chapter 1. Introduction</b> .....	1
1.1 The need for converged fabrics .....	2
1.2 Converged fabrics .....	3
1.3 Fibre Channel over Ethernet .....	4
1.4 Summary .....	5
<b>Chapter 2. Solutions</b> .....	7
2.1 Introduction to FCoE .....	8
2.1.1 FCoE functionality .....	8
2.1.2 FCIA FCoE recommendations .....	9
2.1.3 FCoE physical media considerations .....	9
2.2 FCoE products and technologies .....	11
2.2.1 Converged Network Adapters (CNAs) .....	11
2.2.2 FCoE Switches .....	15
2.2.3 FCoE cabling systems .....	20
2.3 Sample configurations .....	27
2.3.1 BladeCenter H configuration .....	28
2.3.2 BladeCenter HT configuration .....	30
2.3.3 Top-of-rack LAN switching and end-of-row SAN switching (LAN aggregation switching) .....	32
2.3.4 Top-of-rack LAN/SAN switching and end-of-row patching .....	35
2.4 FCoE deployment scenarios .....	38
2.4.1 Scenario 1: Existing data centers with large investment in FC storage .....	38
2.4.2 Scenario 2: New data center installation .....	38
2.4.3 Scenario 3: Existing data centers with mix of storage (DAS, FC, NAS) requiring near-term data center consolidation .....	39
2.4.4 Scenario 4: Small to medium data centers .....	39
<b>Chapter 3. Benefits and business impact</b> .....	41
<b>Related publications</b> .....	45
IBM Redbooks .....	45
Other publications .....	45
Online resources .....	46
How to get Redbooks .....	46
Help from IBM .....	46



# Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

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

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

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

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

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

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

## COPYRIGHT LICENSE:

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

# Trademarks

IBM, the IBM logo, and [ibm.com](http://www.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:

The following terms are trademarks of other companies:

BladeCenter®

IBM®

Redbooks®

Redpaper™

Redbooks (logo) ®

System Storage™

System x®

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

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

# Preface

Over the past decade, almost every industry has witnessed increasing demands for compute capacity and an exponential growth of data, which has resulted in complex and expensive IT infrastructures that are not only cumbersome to deploy, but also to administer and maintain. To counter these increasing demands and complexity in data center infrastructures, as well as to improve utilization, efficiency, and security, there has been greater emphasis on server virtualization and infrastructure consolidation.

The goal of this IBM® Redpaper™ document is to highlight some of the important challenges currently faced in the IT industry with infrastructure growth, and to outline the need for *converged fabrics*. We introduce the emerging IT standards for convergence such as Convergence Enhanced Ethernet (CEE) and Fibre Channel over Ethernet (FCoE). We describe many of the available products and solutions enabling converged fabrics, such as Converged Network Adapters and Converged Fabrics switches that IBM along with business partners such as Panduit are bringing to market.

We continue the discussion by describing the elements of physical deployment planning of convergence technologies, and discuss various scenarios that demonstrate transitioning from traditional to converged infrastructures while following the best practices.

This paper is aimed at data center managers and others who are involved with the planning, implementation, and maintenance of infrastructure in the IT environment.

## The team who wrote this paper

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

**Srihari Angaluri** is a Senior Solutions Architect in the System x® Integrated Solutions group in Research Triangle Park, NC, USA. He has several years of experience in developing and implementing high-performance cluster systems and applications in various industries such as banking, finance, media and entertainment, education and oil and gas. Srihari holds an M.S. in Computer Science and is currently pursuing an MBA degree at Duke University.

**Martin Bachmaier** is an IT Architect in the IBM hardware development lab in Boeblingen, Germany, and leads the Open Systems Design and Development infrastructure team that develops systems using IBM next generation processors. His current responsibilities include the architecture and management of Linux®-based infrastructures and HPC cluster environments. Martin has worked for IBM for more than six years and is deeply involved in IBM's cloud computing activities. He holds a degree in Computer Science from the University of Cooperative Education in Stuttgart, Germany, and a Bachelor of Science from the Open University in London, UK. He is an IBM Certified Systems Expert and holds the CCNA, CCNA Security, and VMware Certified Professional credentials.

**Tom Boucher and the team from the Network Systems Group of Panduit** - Panduit is a leading developer and provider of solutions that help customers optimize the physical infrastructure through simplification, agility, and operational efficiency. Panduit's Unified Physical Infrastructuresm (UPI) based solutions give enterprises the capabilities to connect, manage, and automate communications, computing, power, control, and security systems for a smarter, unified business foundation. Strong relationships with technology leaders

complemented with its global staff and unmatched service and support make Panduit a valuable and trusted partner.

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

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

- ▶ Linda Robinson, ITSO Raleigh
- ▶ Alison Chandler, ITSO Poughkeepsie

## 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](http://ibm.com/redbooks/residencies.html)

## Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks publications in one of the following ways:

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

[ibm.com/redbooks](http://ibm.com/redbooks)

- ▶ Send your comments in an e-mail to:

[redbooks@us.ibm.com](mailto: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/IBMRedbooks>
- ▶ 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>





# Introduction

Over the past decade, almost every industry has witnessed increasing demand for compute capacity and an exponential growth of data, which has resulted in complex and expensive IT infrastructures that are not only cumbersome to deploy, but also to administer and maintain. Today's typical IT infrastructures demand highly skilled personnel to manage various individual subsystems such as servers, networking, and storage.

With the emergence of new computing models and applications such as *cloud computing*, the challenge of dealing with growth becomes paramount for IT managers. As compute, networking, and storage infrastructures scale, the complexity and costs of procurement, deployment, and management of these infrastructures also increase tremendously. Lately, in order to counter these increasing demands and complexity in data center infrastructures, as well as to improve other aspects such as utilization, efficiency, security, and so forth, there has been increased emphasis on server virtualization and infrastructure consolidation.

The goal of this paper is to highlight some of the important challenges currently faced in the IT industry with infrastructure growth and outline the need for *converged fabrics*. In this chapter, we introduce the emerging IT standards for convergence, such as Convergence Enhanced Ethernet (CEE) and Fibre Channel over Ethernet (FCoE). Several products and solutions enabling converged fabrics, such as Converged Network Adapters (CNAs) and Converged Fabrics switches, are introduced in Chapter 2, "Solutions" on page 7.

The standards and technologies for converged fabrics have only been introduced recently, so this market is still immature. Consequently, best practices for the adoption and deployment of converged fabrics technologies are currently lacking in the industry. To facilitate an informed transition to converged networks, we put special emphasis on the physical deployment planning of convergence technologies, and discuss various scenarios in transitioning to FCoE-based infrastructures.

We believe the techniques outlined in this paper will be helpful for the proper deployment as well as efficient use of FCoE in any type of environment – from small scale lab infrastructures all the way to large scale data center environments with tens of hundreds of servers and storage devices.

## 1.1 The need for converged fabrics

A typical data center infrastructure today consists of several distinct network fabrics, to carry three different traffic types:

- ▶ One network fabric for storage I/O (Fibre Channel-based, for example)
- ▶ One network fabric for computing and management (Ethernet-based, for example)
- ▶ One network fabric for interprocess communications

Each fabric is made of the respective network adapters on the host side, the switching devices, and the physical interconnects (for example, fiber optics and Ethernet). This approach is illustrated in Figure 1-1.

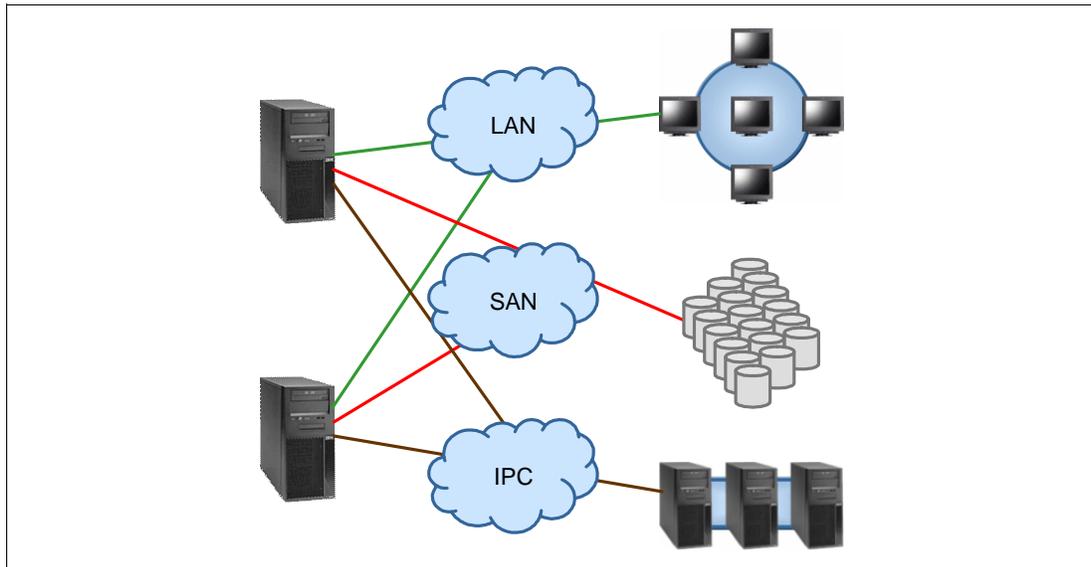


Figure 1-1 Today's segmented fabric approach for networking

As shown in Figure 1-1, there are commonly three types of traffic – LAN, SAN, and IPC:

- ▶ The LAN traffic is typically 1 Gbps Ethernet used for low-bandwidth communication such as server out-of-band management, remote access, PXE booting, DNS, and so forth.
- ▶ The IPC fabric is used for interprocess communication and message passing, such as the process-to-process communication that takes place in a tightly coupled cluster system. Since this is bandwidth-intensive and latency-sensitive communication, a high-bandwidth, low-latency interconnect such as InfiniBand is generally used for this type of traffic.
- ▶ The SAN fabric is used for storage area networking. This network is used for all I/O traffic to and from hosts to the storage devices such as disk and tape systems. The most popular physical medium used for storage I/O today is optic-fibre. Currently, fibre-optic medium supports speeds of 8 Gbps.

Although this traditional approach of segmenting storage and data traffic has certain advantages (such as traffic isolation, independent administration, and manageability), it nevertheless poses several disadvantages, including higher infrastructure costs, complexity of management, and under-utilization of resources. Customers have to invest in separate infrastructures for LAN, SAN, and IPC fabrics, including host adapters, physical media, and network switching, routing, and other device-specific equipment. From a manageability standpoint, each network fabric is managed as a separate entity, with dedicated personnel assigned to manage each fabric.

At any one time, the individual fabrics tend to not be fully utilized because only a percentage of peak bandwidth is used in the respective networks, which results in poor overall utilization of these resources, and consequently lower ROI.

## 1.2 Converged fabrics

Given the complexities and costs associated with the traditional segmented network fabric design in data centers, the networking industry and standards committees over the recent years have been driving discussions and working on new standards that will be necessary to create a unified or converged fabric. Such a fabric will enable multiple traffic types—such as interprocess communication, remote management, clustering, data, as well as storage—all to be carried over a single physical medium, the *converged fabric*.

There have been several attempts to unify and extend the reach of storage networks with protocols such as iSCSI, FCIP, iFCP and InfiniBand<sup>1</sup>; however, these technologies did not successfully meet the goals of converged fabrics due to reasons such as performance, interoperability, industry adoption, cost, or the requirement for deploying a new infrastructure from scratch (for example, InfiniBand).

Because of its ubiquitous nature and wide popularity in the industry, the standards bodies have selected Ethernet as the converged medium. The 10 Gigabit Ethernet standard was specified as the base on which to define new protocols that will enable it to meet the requirements to effectively carry various traffic types, while delivering the performance, reliability, and other familiar functions of the traditional network and storage protocols.

The new protocols designed to enhance Ethernet—Priority-based Flow Control (PFC), Enhanced Transmission Selection (ETS), and Data Center Bridging Capabilities Exchange (DCBX)—are part of the *Convergence Enhanced Ethernet* (CEE) standard, which is defined by the IEEE 802.1 working group.

Figure 1-2 on page 4 illustrates the central idea behind converged fabrics. With a combination of 10 Gigabit Ethernet technology and the new CEE protocols, users will be able to use the 10 Gigabit Ethernet fabric to carry traditional network traffic as well as communicate with storage devices using Fibre Channel protocol. Hence, a single unified fabric that consists of newly introduced Converged Network Adapters (CNAs) that install in the hosts, and network switches that understand the associated CEE protocols, will be sufficient to carry all traffic types.

---

<sup>1</sup> See <http://www.networkworld.com/news/tech/2008/042208-tech-update.html>

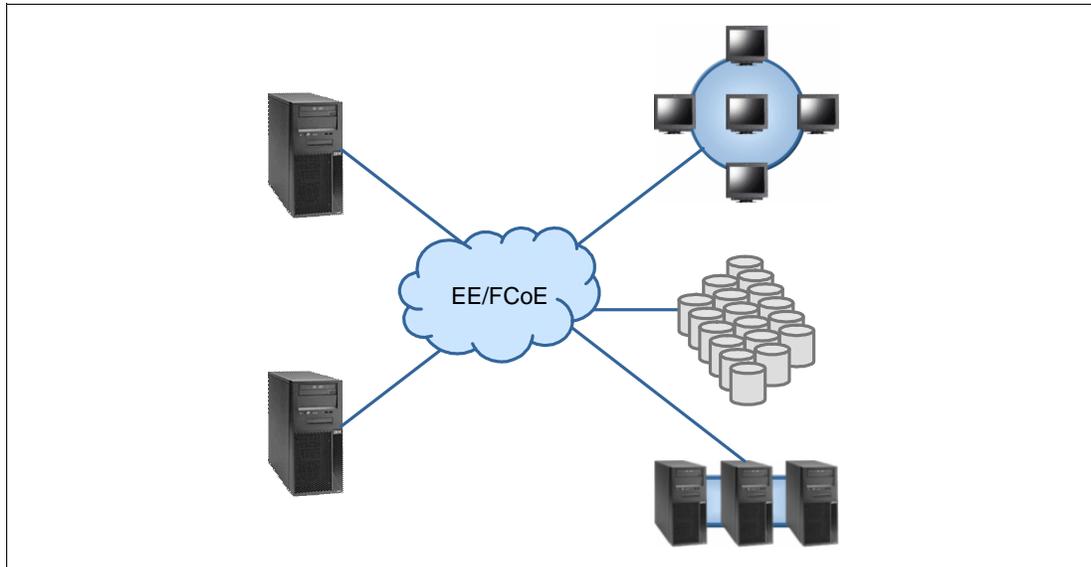


Figure 1-2 Network consolidation using Enhanced Ethernet and FCoE technologies

### 1.3 Fibre Channel over Ethernet

To facilitate convergence and overcome the problems encountered with previous attempts at a convergence standard, a consortium of storage and data network vendors under the INCITS T11 committee worked on a new standard called *Fibre Channel over Ethernet* (FCoE)<sup>2</sup>. FCoE is designed to enable network convergence and cost-effective SAN expansion in the enterprise data center. The FCoE protocol, which will enable Fibre Channel protocol packets to be encapsulated into Ethernet frames, is implemented on top of Convergence Enhanced Ethernet.

With the standardization of converged fabrics via FCoE and related technologies, several vendors, including IBM, have started announcing products and technologies in support of FCoE. Customers are beginning to take great interest in converged fabrics technologies because they promise to deliver better TCO, high cost savings, simplified deployment, manageability, and unified administration, among other advantages.

Because we are still in the early stages of FCoE adoption, best practices for the deployment, management, and use of converged fabrics have not fully evolved. Early customers are using FCoE primarily on an experimental basis and at a small scale, consisting of tens of servers and switching infrastructure. Scaling the infrastructure beyond a small set of servers and switches will pose new challenges and require detailed planning. Consequently, the full potential of FCoE may not be realized by customers who do not fully understand various implications and approaches involved in properly deploying and utilizing the technology.

Figure 1-3 on page 5 shows the use of 10 Gbps Ethernet medium for carrying LAN, IPC, and SAN traffics. Essentially, the Fibre Channel protocol packets are encapsulated in the Ethernet frames and transmitted to a new FCoE-capable switch, which can also incorporate a data center bridging capability. The switch device will extract the Fibre Channel frame from the Ethernet frame and forward the packet to the storage device through the data center bridge. Furthermore, new storage products that are FCoE-aware are being introduced to the market.

<sup>2</sup> See <http://www.fcoe.com/> for more information.

These devices can directly communicate with FCoE switches, without the need for an intermediate bridge.

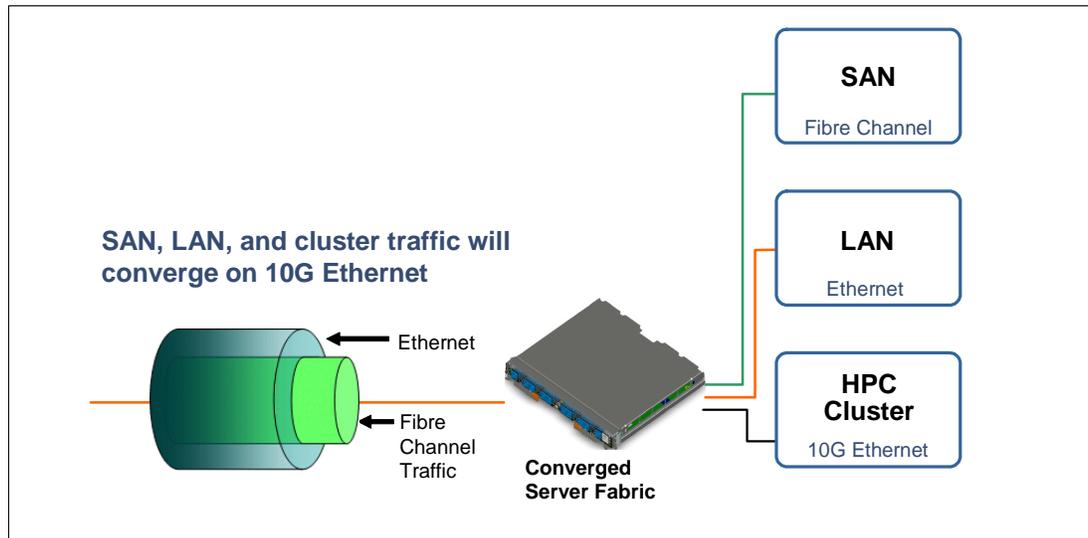


Figure 1-3 Convergence of SAN, LAN and Cluster Traffic over 10GbE Fabric

## 1.4 Summary

As the need to consolidate server infrastructures and simplify the networking in data centers emerges, new standards related to the Convergence Enhanced Ethernet (CEE) are coming to the rescue of IT managers and owners.

The goal of this paper is to introduce these standards and describe the new products and solutions from IBM and partner companies in support of the converged fabrics. The physical deployment aspects related to FCoE and the key planning guidelines for proper deployment and operation are discussed in the paper with example deployment scenarios and case studies.

The next chapter discusses FCoE protocol in more depth and introduces various new FCoE products and solutions from IBM such as CNA adapters and FCoE switches. Chapter 2 also discusses concepts related to the physical deployment of FCoE technology and highlights some of the key aspects involved in planning for FCoE deployment. Products from IBM partner company Panduit, Inc., which facilitate the physical deployment of FCoE products and enable the efficient management of FCoE infrastructures in large-scale settings are also described in the next chapter.

In Chapter 3, we discuss the business benefits of FCoE technology: energy savings, data center simplification, preparing for future IT requirements, and how FCoE can reduce your total cost of ownership.





# Solutions

This chapter provides a brief introduction to the technical details of Fibre Channel over Ethernet (FCoE) and shows you how to use it. It describes the products (switches and host adapter) currently available in the IBM portfolio as well as Panduit's offerings for structured FCoE data center cabling. It closes with a discussion of different FCoE deployment scenarios in today's data centers and examples of how to deploy FCoE from a physical perspective.

Topics in this chapter are:

- ▶ 2.1, "Introduction to FCoE" on page 8
- ▶ 2.2, "FCoE products and technologies" on page 11
- ▶ 2.3, "Sample configurations" on page 27
- ▶ 2.4, "FCoE deployment scenarios" on page 38

## 2.1 Introduction to FCoE

Fibre Channel Over Ethernet defines a new standard for the traditional Fibre Channel storage protocols, which allows Ethernet as the physical transmission technology for the Fibre Channel storage traffic. The final draft was released in June 2009 and is currently in the process of publication as an ANSI standard. FCoE combines Fibre Channel and Ethernet technologies to provide a “converged” network fabric that can be used as a multi-purpose medium to carry typical data center storage and network traffic types such as systems management, server-to-server messaging, client/server messaging, and storage (SAN) connectivity.

The biggest advantage of FCoE is the single converged fabric that carries both traditional Fibre Channel SAN traffic and Ethernet traffic. For existing SAN environments using Fibre Channel, FCoE allows interoperability with the OS and management tools because the underlying interfaces do not change. Hence, the investments for FC-SAN owners in their existing OS and management tools will be preserved when they transition to FCoE technology. In addition, because FCoE consolidates the physical medium into a single network to carry both storage and Ethernet traffic, it reduces infrastructure costs like power or cooling and simplifies management.

### 2.1.1 FCoE functionality

With the emergence of 10 Gbps Ethernet, fabric bandwidth is no longer the bottleneck for today’s data center. Those in need of more than 10 Gbps can aggregate multiple links, or they can upgrade to 40 Gbps or 100 Gbps Ethernet in the near future.

Figure 2-1 shows how FC packets are sent over Ethernet. The FC packet consists of the FC payload (the actual data) and the original FC header. This FC packet then gets a trailing FCoE header and is encapsulated in a standard Ethernet frame.

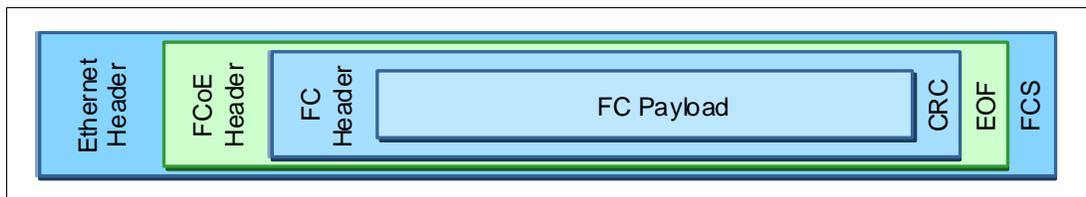


Figure 2-1 Encapsulation of Fibre Channel payload in Ethernet frames

Ethernet is not lossless, which means it is within specification if Ethernet frames get lost during transmission. Higher layer protocols (for example, TCP) ensure detection and retransmission of lost packets. Fibre Channel, on the other hand, is lossless. To enable the use of Ethernet as the medium for FC communication, several enhancements had to be developed that fix Ethernet shortcomings like the lossy behavior, and that allow for improved flow control. The Data Center Bridging (DCB) Task Group of the IEEE 802.1 Working Group is currently working on standards to satisfy requirements when using Ethernet as a single converged fabric. The standards cover the following major areas:

- ▶ Congestion Notification (CN, IEEE 802.1Qau) provides sophisticated end-to-end congestion management for all upper-layer protocols, like FCoE.
- ▶ Priority-Based Flow Control (PFC, IEEE 802.1Qbb) provides a link level flow control mechanism that enables independent and granular control of different Class of Service (802.1p) streams.

- ▶ Enhanced Transmission Selection (ETS, IEEE 802.1Qaz) provides a common management framework for dynamic assignment of bandwidth to different Class of Services.
- ▶ Data Center Bridging Capabilities Exchange Protocol (DCBCXP, IEEE 802.1AB) is a discovery and capability exchange protocol for Enhanced Ethernet to ensure a consistent configuration across multiple devices. It is based on the Link Layer Discovery Protocol (LLDP).

For more detailed technical information about FCoE, refer to the IBM Redpaper document *An Introduction to Fibre Channel over Ethernet, and Fibre Channel over Convergence Enhanced Ethernet*, REDP-4493, available from the following Web site:

<http://www.redbooks.ibm.com/abstracts/redp4493.html>

## 2.1.2 FCIA FCoE recommendations

The Fibre Channel Industry Association (FCIA) has published a guideline specifying what technologies and standards to consider when deploying Fibre Channel over Ethernet<sup>1</sup>. Table 2-1 shows the recommendations.

Table 2-1 FCIA FCoE recommendations for a converged network

Category	Recommendations
10GbE physical protocol	<ul style="list-style-type: none"> <li>▶ Optical connectivity: IEEE 802.3 10 Gbps serial (<math>\leq 100m</math> MM OM3 recommended to new installations for compatibility with 40/100gE and 16/32GFC)</li> <li>▶ Copper connectivity: SFP+ copper definition for external copper connection (10GBase10-KX4/KR internal backplanes and blades)</li> <li>▶ Connector for optics and copper: SFP+</li> <li>▶ 802.3X Link Level Flow Control and Baby Jumbo Frames</li> </ul>
Ethernet logical protocol	Data Center Bridging (DCB): <ul style="list-style-type: none"> <li>▶ IEEE 802.1Qaz Enhanced Transmission Selection (Priority Groups)</li> <li>▶ IEEE 802.1Qbb Priority-based Flow Control</li> <li>▶ DCB Capability Discovery and Exchange Protocol (DCBX)</li> </ul>
Fibre Channel logical protocol	INCITS T11.3 BB-5 <ul style="list-style-type: none"> <li>▶ FCoE Frame Format</li> <li>▶ FCoE Initialization Protocol (FIP)</li> </ul>
Server	PCI Express 2.0

It starts with the physical layer (top layer) that shows the supported media and protocols for FCoE. Then, the logical protocols (second layer) show the new enhancements to traditional Ethernet. The third layer defines the FCoE format, which is responsible for the encapsulation of Fibre Channel packets into FCoE frames. All this is based on PCI Express 2.0 as the hardware interface to support the required throughput for 10 Gigabit Ethernet.

## 2.1.3 FCoE physical media considerations

As with any new generation of networks, the supported physical media becomes important. The end user wants to continue to use existing media (like twisted pair) because they are cheap and quite often already installed. The problem is to find a way to map all the features and improvements of the new protocol onto the old existing media. Sometimes, only a subset of the features can be supported, or the old media has power or length restrictions. This is

<sup>1</sup> It is available from [http://www.fibrechannel.org/documents/doc\\_download/6-fcia-fcoe-guideline](http://www.fibrechannel.org/documents/doc_download/6-fcia-fcoe-guideline)

why additional, new media get specified. For FCoE using 10 Gbps Ethernet, two media are specified at the time of this writing: Fibre optic cable and a special Direct Attach Copper (DAC) Cable where the connector is closely linked with the cable. Twisted pair is likely to follow in the future. Table 2-2 gives a brief overview of the supported media.

Table 2-2 A comparison of FCoE physical media

Feature	Fiber optic cable	Direct Attach Copper (DAC) Cable Assembly
Connector	SFP+ modular transceiver with FOCIS-10 connector (LC)	SFP+ module
Reach at 10 Gbps	Depends on fiber type and transceiver <sup>a</sup>	Passive: 7 m Active: 15 m
Port latency	~ 0.2 $\mu$ s	~ 0.2 $\mu$ s
Cost (relative)	\$\$	\$

a. Channel reach of FCoE fiber solutions depend on the fiber type/transceiver selected and the vendor-specific limitations of the FCoE CNA and switch(es) deployed.

### Reach limitations imposed by lossless FCoE and Priority Flow Control

Priority Flow Control (PFC) is one of the functions of IEEE 802.1Q (covering Data Center Bridging) introduced to address reliability in enabling FCoE for certain Class of Services (CoS) required to build a “lossless” Layer 2 medium.

A receiver using PFC must be able to predict, in a timely fashion (under certain CoSs) and be able to respond (to the other end of the link) to a PAUSE frame such that the source sending the PAUSE has time to stop sending before its receive buffer overflows. The length of the channel therefore affects how early the receiving end must send back the PAUSE frame (longer channels must send PAUSE quicker). Buffer on the receive end must be present in sufficient quantity to store information present in the channel and delivered to the receiver in the amount of time it takes for the side transmitting to execute the PAUSE (all packets drained from the channel into the buffer).

Distance limitations placed on the channel are a function of packet size (Maximum Transmission Unit), wire speed (copper versus fiber), latencies of the selected PMDs, and response time of the logic in the send function (maximum of 3840 bytes in any implementation).

This describes known limitations of PFC within the ecosystem of IBM and other FCoE vendors for the purposes of FCoE, or applications that might use PFC functions in the future. PFC is affected by the length of the Ethernet segment between a transmitter and the receiver, and understanding the restrictions placed on both the ends of the link by respective equipment manufacturers is paramount to achieving expected channel reach goals. Consult equipment manufacturers’ specifications to find stated reach while using PFC function.

For a more detailed technical analysis refer to the Cisco white paper *Priority Flow Control: Build Reliable Layer 2 Infrastructure* available from:

[http://ciscosystems.com/en/US/prod/collateral/switches/ps9441/ps9670/white\\_paper\\_c11-542809.html](http://ciscosystems.com/en/US/prod/collateral/switches/ps9441/ps9670/white_paper_c11-542809.html)

## 2.2 FCoE products and technologies

The introduction of a converged fabric required a change in the naming scheme. To distinguish between controllers for traditional Ethernet or Fibre Channel and those that connect to a converged fabric, these new controllers are called converged network adapters (CNAs).

A converged fabric also requires FCoE-enabled switching devices. The devices must recognize the new Ethernet extensions, as described in “Introduction to FCoE” on page 8. There are two types of switches. The first type has only FCoE ports and no dedicated Ethernet or Fibre Channel ports. The second type has dedicated Ethernet or Fibre Channel ports, or both. This type can be used as bridging devices between existing LAN/SAN fabric installations and an FCoE-enabled fabric.

The following two sections introduce the IBM portfolio of CNAs and FCoE-enabled switches for System x and BladeCenter. We then describe the solutions available from Panduit.

### 2.2.1 Converged Network Adapters (CNAs)

At the time of this writing, IBM offers two CNAs for System x and one for BladeCenter. Table 2-3 compares the three adapters and highlights major features.

*Table 2-3 Key features of IBM System x and BladeCenter CNAs*

	<b>Brocade CNA for System x</b>	<b>QLogic CNA for System x</b>	<b>QLogic CNA for BladeCenter</b>
Part number	42C1820	42C1800	42C1830
Host interface	PCIe 2.0 x8	PCIe 2.0 x4 (also works in PCIe Gen 1.1 x8 slot)	CFFh, PCIe 2.0 x8
Number of ports	2	2	2
Supported SFP+ modules	Fiber SR or direct-attach copper	Fiber SR or direct-attach copper	n/a (BladeCenter internal)
FCoE hardware offload	Yes	Yes	Yes
PXE boot	Yes	Yes	Yes
SAN boot over CEE	Yes	Yes	Yes
IOPS per port	500,000	250,000	250,000

All CNAs use PCI Express as the host interface and provide two ports to the Operating system. Table 2-4 on page 12 shows the supported servers for the two System x CNAs; Table 2-5 shows the supported BladeCenter servers.

Table 2-4 Supported IBM System x servers

	x3200 M2	x3200 M3	x3250 M2	x3250 M3	x3350	x3400 M2	x3455	x3500 M2	x3550	x3550 M2	x3650	x3650 M2	x3755	x3850 M2	x3950 M2
Brocade CNA, 42C1820	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
QLogic CNA, 42C1800	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Further information about the available CNAs for IBM System x can be gathered from the following Web site:

<http://www.ibm.com/systems/storage/product/systemx/cna/>

Table 2-5 Supported IBM BladeCenter servers

	HS12	HS21	HS21XM	HS22	HS22V	LS21	LS22	LS41	LS42	JS12	JS21	JS22	JS23/JS43	QS22	PN41
QLogic CNA, 42C1830	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	Y	Y	N	N

### Brocade 10 Gb CNA for System x

The Brocade 10 Gb CNA for System x is a PCI Express Gen2 x8 adapter with two SFP+ cages for optical or copper cables. It offers 10 Gbps maximum bidirectional bandwidth on each port and provides full hardware offload for FCoE protocol processing.

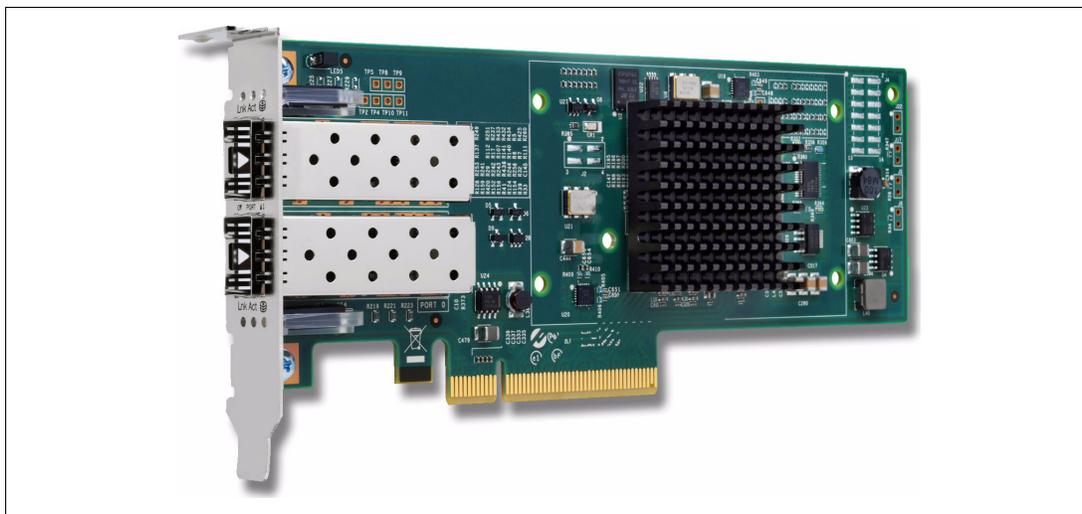


Figure 2-2 Brocade 10 Gb CNA for IBM System x

The Brocade 10 Gb CNA for IBM System x has the following major features:

- ▶ PCI Express x8 Gen2 compliance
- ▶ Two SFP+ cages for either SFP+ Fiber SR or SFP+ Active Copper Cables
- ▶ Standard PCI Express half length card with low profile form factor
- ▶ Support for both standard PCI-E slot and low profile PCI-E slot
- ▶ Support for 10 Gb Converged Enhanced Ethernet (CEE)
- ▶ Support for FC over Converged Enhanced Ethernet (FCoCEE)
- ▶ Full hardware offload for FC protocol processing

- ▶ Support for IPv4 and IPv6
- ▶ Support for SAN boot over CEE, PXE boot, and iSCSI boot
- ▶ Performance: 500,000 IOPS per port
- ▶ Support for 2048 logins and 4096 exchanges

The at-a-glance guide with more detailed technical information about the Brocade 10 Gb CNA is available from:

<http://www.redbooks.ibm.com/abstracts/tips0718.html>

### QLogic 10 Gb CNA for System x

The QLogic 10 Gb CNA for System x is a PCI Express Gen2 x4 adapter with two SFP+ cages for optical or copper cables. It offers 10 Gbps maximum bidirectional bandwidth on each port and provides full hardware offload for FCoE protocol processing.

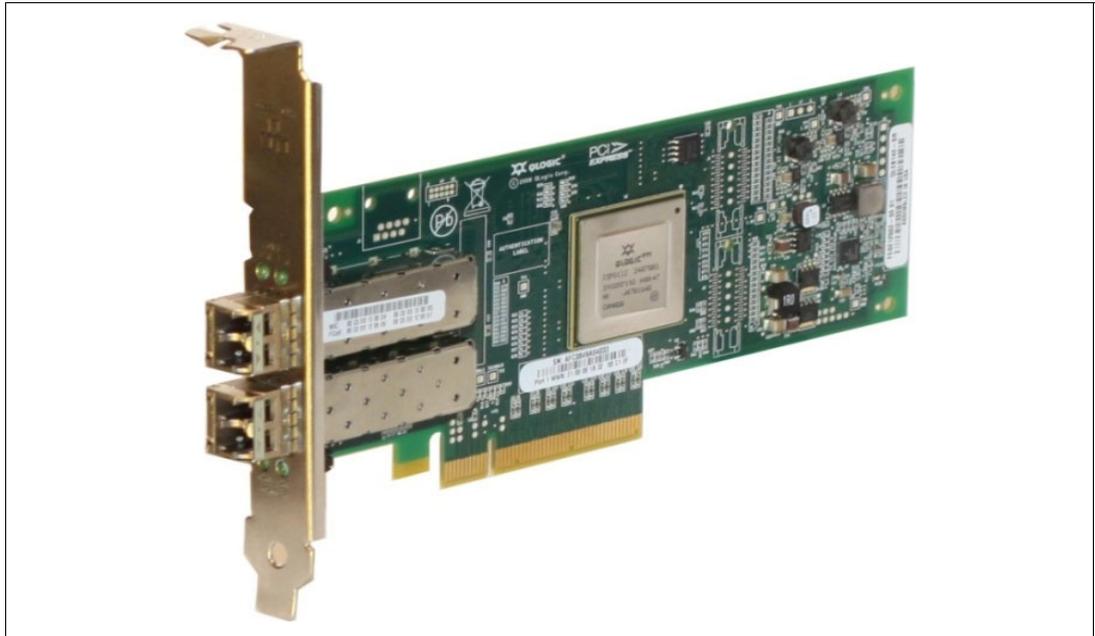


Figure 2-3 QLogic dual-port PCIe CNA Adapter

The QLogic 10 Gb CNA has the following major features:

- ▶ PCI Express 2.0 compliance
- ▶ Operates at PCI Express 2.0 x4 or PCI Express 1.1 x8
- ▶ Two SFP+ cages for either SFP+ Fiber SR or SFP+ Active Copper Cables
- ▶ Standard PCI Express half length card with low profile form factor
- ▶ Support for both standard PCI-E slot and low profile PCI-E slot
- ▶ Support for 10 Gb Converged Enhanced Ethernet (CEE)
- ▶ Support for FC over Converged Enhanced Ethernet (FCoCEE)
- ▶ Full hardware offload for FC protocol processing
- ▶ Support for IPv4 and IPv6
- ▶ Support for SAN boot over CEE and iSCSI boot
- ▶ Performance: 250,000 IOPS per port
- ▶ Support for 2048 logins and 2048 exchanges

The at-a-glance guide with more detailed technical information about the QLogic 10 Gb CNA is available from:

<http://www.redbooks.ibm.com/abstracts/tips0720.html>

## QLogic CNA Adapter for IBM BladeCenter

The QLogic CNA Adapter for IBM BladeCenter is an expansion card for the PCI Express interface on blades. The adapter provides two CEE ports and connects them to the high-speed switch modules (HSSM) on the back of the BladeCenter. Each CEE port offers 8 Gbps Fibre Channel connectivity and 10 Gbps networking.

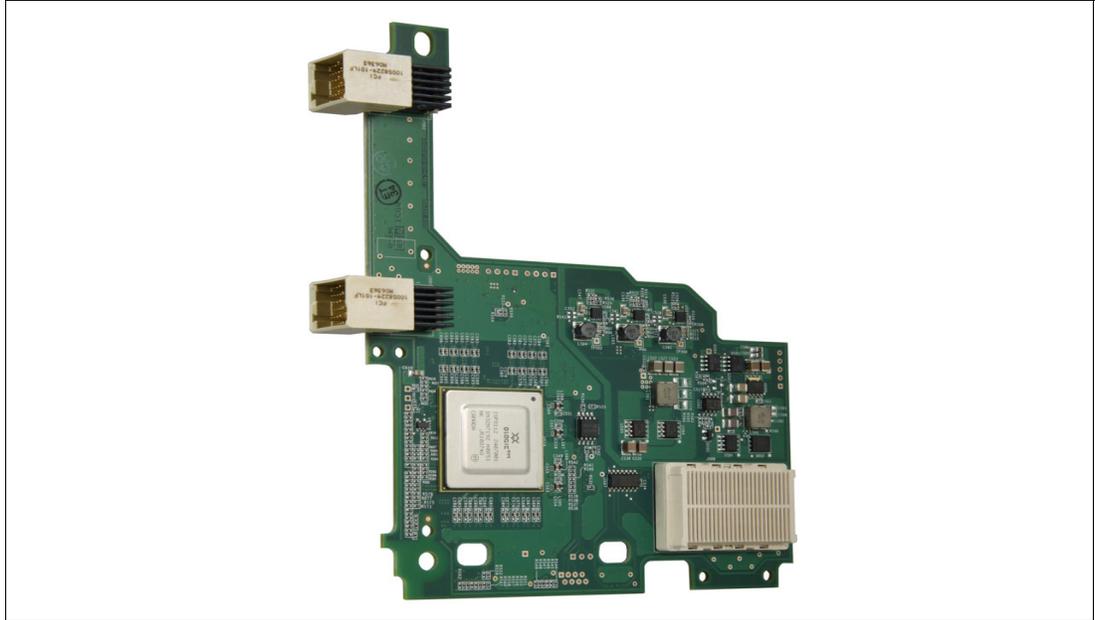


Figure 2-4 QLogic CNA Adapter for IBM BladeCenter

The QLogic CNA for IBM BladeCenter has the following major features:

- ▶ Combo Form Factor (CFFh) PCI Express 2.0 x8 adapter
- ▶ Support for 10 Gb Converged Enhanced Ethernet (CEE)
- ▶ Support for FC over Converged Enhanced Ethernet (FCoCEE)
- ▶ Full hardware offload for FC protocol processing
- ▶ Support for IPv4 and IPv6
- ▶ Support for SAN boot over CEE, PXE boot, and iSCSI boot
- ▶ Support for BladeCenter Open Fabric Manager for BIOS, UEFI, and FCode
- ▶ Performance: 250,000 IOPS per port
- ▶ Support for 2048 logins and 2048 exchanges
- ▶ Compliance with the following IEEE standards:
  - 802.1Qbb rev. 0 (Priority-based Flow Control)
  - 802.1Qaz rev. 0 (Enhanced Transmission Selection)
  - 802.1Qaz rev. 0 (DCBX protocol)

The QLogic CNA adapter is supported in the IBM BladeCenter H and BladeCenter HT with the following I/O modules:

- ▶ 10 Gb Ethernet Pass-Thru Module
- ▶ BNT Virtual Fabric 10 Gb Switch Module
- ▶ Cisco Nexus 4001l Switch Module

The at-a-glance guide with more detailed technical information about the QLogic CNA adapter for IBM BladeCenter is available from:

<http://www.redbooks.ibm.com/abstracts/tips0716.html>

## 2.2.2 FCoE Switches

IBM offers FCoE-enabled switches from several vendors. This allows for a seamless introduction of FCoE into an existing environment that is already customized and uses the management tools from one vendor. Table 2-6 shows the available switches and their major features. For more detailed information, refer to the following switch-specific sections.

Table 2-6 Overview of FCoE-enabled switches for IBM System x and BladeCenter

	IBM Converged Switch B32	QLogic 10 Gb Passthru Module	BNT Virtual Fabric 10 Gb Switch Module	Cisco Nexus 4001I	QLogic Virtual Fabric Extension Module
Form factor	Top-of-rack (TOR)	BladeCenter HSSM	BladeCenter HSSM	BladeCenter HSSM	BladeCenter Standard SM <sup>a</sup>
FCoE Ports (BladeCenter internal)	Not applicable	14	14	14	0
FCoE Ports (external)	24	14	10	6	0
FC Ports	8 ports, each at 1/2/4/8 Gbps FC	0	0	0	6 ports, each at 2/4/8 Gbps FC
Traditional Ethernet ports (not CEE-enabled)	0	0	0	0	0
SFP+ modules	10 Gb: SR, LR <sup>b</sup> , DAC FC ports: 8 Gb SFP+ and 4 Gb SFP modules (SW and LW)	10 Gb: SR, LR, DAC	10 Gb: SR, LR, DAC 1 Gb: 1000BASE-T, SX	10 Gb: SR, LR, DAC 1 Gb: 1000BASE-T, SX, LH	8 Gb SFP+ SW Fibre Channel
Congestion Notification	No	No, passthru only	No	No	Not applicable
Priority-based Flow Control	Yes	No, passthru only	Yes	Yes	Not applicable
Enhanced Transmission Selection	Yes	No, passthru only	Yes	Yes	Not applicable
Data Center Bridging Capabilities Exchange Protocol	Yes	No, passthru only	Yes	Yes	Not applicable

a. The QLogic Virtual Fabric Extension Module requires at least one BNT Virtual Fabric 10 Gb Switch Module to enable CEE traffic from/to the blade servers.

b. When deploying long-reach (LR) SFP+ modules, be aware of the PFC length limitations explained in 2.1.3, "FCoE physical media considerations" on page 9.

## IBM Converged Switch B32

The IBM Converged Switch B32 is designed to support traditional Ethernet, Converged Enhanced Ethernet (CEE), Fibre Channel, and the Fibre Channel over Ethernet (FCoE) protocol within a single switch. Designed as a 1U, top of rack device, its eight Fibre Channel ports support up to 8 Gbps and the 24 Ethernet ports up to 10 Gbps throughput each.



Figure 2-5 IBM Converged Switch B32

Technical features of the QLogic 10 Gb Passthru Module are:

- ▶ 1U, 19 inch, top of rack form factor
- ▶ Eight Fibre Channel ports running at 1/2/4/8 Gbps each, with support for
  - 2/4/8 Gbps SFP+ modules, shortwave and long wave laser
  - 1/2/4 Gbps SFP modules, shortwave and long wave laser
- ▶ 24 10 Gbps Ethernet ports, CEE ready, with support for
  - 10 Gb SFP+ modules, SR or LR
  - 10 Gb DAC cables
- ▶ Support for ISL trunking for Fibre Channel
- ▶ Support for Link Aggregation Control Protocol (LACP) for CEE
- ▶ Support for Jumbo Frames (up to 9048 bytes)

An IBM Redpaper on the IBM Converged Switch B32 is available from:

<http://www.redbooks.ibm.com/abstracts/redp4588.html>

The following table shows the IBM BladeCenter FCoE-enabled switches with the supported expansion cards.

Table 2-7 IBM FCoE-enabled BladeCenter switches and supported expansion cards

	QLogic 10 Gb Passthru Module	BNT Virtual Fabric 10 Gb Switch Module	Cisco Nexus 4001I	QLogic Virtual Fabric Extension Module
2/4 Port Ethernet Expansion Card (CFFh)	No	Yes	Yes	No
NetXen 10 Gb Ethernet Expansion Card (CFFh)	Yes	Yes	Yes	No
Broadcom 10 Gb 2-port Ethernet Exp. Card (CFFh)	Yes	Yes	Yes	No
Broadcom 10 Gb 4-port Ethernet Exp. Card (CFFh)	Yes	Yes	Yes	No
Broadcom 10 Gb Gen 2 2-port Ethernet Exp. Card (CFFh)	No	Yes	Yes	No
Broadcom 10 Gb Gen 2 4-port Ethernet Exp. Card (CFFh)	No	Yes	Yes	No

	QLogic 10 Gb Passthru Module	BNT Virtual Fabric 10 Gb Switch Module	Cisco Nexus 4001I	QLogic Virtual Fabric Extension Module
QLogic 2-port 10 Gb CNA (CFFh)	Yes	Yes	Yes	Yes
Emulex Virtual Fabric Adapter (CFFh)	Yes <sup>a</sup>	Yes	Yes <sup>a</sup>	No

a. The QLogic 10 Gb Passthru Module and the Cisco Nexus 4001I support the Emulex Virtual Fabric Adapter only in the physical NIC (pNIC) mode.

The following sections briefly introduce the available FCoE-enabled switch modules for IBM BladeCenter. Figure 2-6 shows how such an environment might look when using the QLogic 10 Gb Passthru Module.

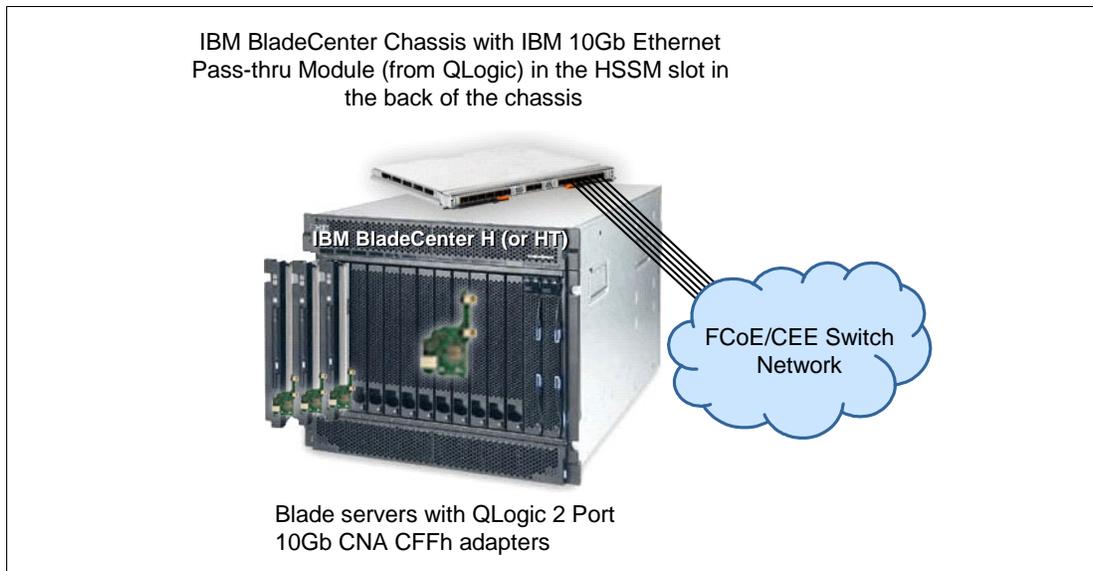


Figure 2-6 QLogic FCoE Networking Solution for IBM BladeCenter

### QLogic 10 Gb Passthru Module for IBM BladeCenter

The QLogic 10 Gb Passthru Module (Figure 2-7) provides a 1-to-1 mapping of the 14 internal ports to 14 external uplink ports (no configuration required). The uplink ports support different SFP+ modules (SR, LR, and DAC).



Figure 2-7 QLogic 10 Gb Passthru Module for IBM BladeCenter

Technical features of the QLogic 10 Gb Passthru Module are:

- ▶ Single-width high speed switch module
- ▶ 14 internal ports at 10 Gbps (no auto-negotiation)
- ▶ 14 external SFP+ ports running at 10 Gbps
- ▶ Support for SR, LR, and DAC cables
- ▶ Direct 1-to-1 mapping of external and internal ports

The at-a-glance guide with more detailed technical information about the QLogic 10 Gb Passthru Module for IBM BladeCenter is available from:

<http://www.redbooks.ibm.com/abstracts/tips0715.html>

### **BNT Virtual Fabric 10 Gb Switch Module for IBM BladeCenter**

The BNT Virtual Fabric 10 Gb Switch Module is a great migration platform for customers who are still using 1 Gbps Ethernet outside the chassis but who plan to introduce Fibre Channel over Ethernet using 10 Gbps CEE. The external switch ports support both 1 and 10 Gbps SFP/SFP+ modules. This allows you to smoothly transition your environment to an FCoE/CEE world. If you have a chassis with multiple servers—some operating at 1 Gbps, some at 10 Gbps, and some transmitting converged packets—this single switch can handle all these workloads and can connect to a 1 Gbps or a 10 Gbps infrastructure, or both.

**Note:** The BNT Virtual Fabric 10 Gb Switch Module replaces the BNT 10 Gb Ethernet Switch Module. BNT firmware release 6.1 introduces features required for Converged Enhanced Ethernet (CEE) and Fibre Channel over Ethernet (FCoE).



*Figure 2-8 BLADE Convergence-ready 10GbE Switch Module for IBM BladeCenter*

Technical features of the BNT Virtual Fabric 10 Gb Switch Module are:

- ▶ Single-width high speed switch module
- ▶ 14 internal auto-negotiating ports at 1 or 10 Gbps
- ▶ Ten external SFP+ cages with support for:
  - 10 Gb SFP+ modules, SR or LR
  - 10 Gb DAC cables
  - 1 Gb SFP modules, 1000BASE-T or SX
- ▶ Very low oversubscription ratio of 14 to 10 (7.2 Gbps per blade server port)
- ▶ Support for up to 32 Kb MAC addresses
- ▶ Support for Jumbo Frames (up to 12288 bytes)

- ▶ Support for up to 1024 VLANs; VLAN numbers from 1 to 4095
- ▶ Support for multiple Layer 2 protocols, among others L2-trunk failover (NIC teaming), VRRP, IGMP snooping
- ▶ Support for Layer 3 protocols, like OSPF and BGP
- ▶ VMready, which means the network is aware of virtual machines on different switch ports and supports VM migration. VMready works with VMware, Xen, KVM, and Hyper-V

With the combination of the BNT switch and the QLogic Virtual Fabric Extension Module, IBM delivers a fully integrated FCoE blade solution without the cost for an additional top of rack switch. For more information, see “QLogic Virtual Fabric Extension Module for IBM BladeCenter” on page 20.

The at-a-glance guide with more detailed technical information about the BNT Virtual Fabric 10 Gb Switch Module for IBM BladeCenter is available from:

<http://www.redbooks.ibm.com/abstracts/tips0708.html>

### Cisco Nexus 4001I Switch Module

The Cisco Nexus 4001I Switch Module 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. It can be used with BladeCenter external FCoE gateways like the Cisco Nexus 5000 Series switches which separate LAN and SAN traffic and provide dedicated connections into the two fabrics.



Figure 2-9 Cisco Nexus Convergence-ready 10GbE Switch Module for IBM BladeCenter

**Note:** The Cisco Nexus 4001I Switch Module is designed to support both 10 Gb Ethernet and Fibre Channel over Ethernet. A software upgrade license (part number 49Y9983) is required for the switch to work in FCoE mode.

Technical features of the Cisco Nexus 4001I are:

- ▶ Single-width high speed switch module
- ▶ 14 internal auto-negotiating ports at 1 or 10 Gbps
- ▶ Six external SFP+ cages with support for:
  - 10 Gbps SFP+ modules, SR or LR
  - 10 Gbps DAC cables
  - 1 Gbps SFP modules, 1000BASE-T, SX, or LX/LH
- ▶ Low, predictable, and consistent latency of 1.5 microseconds regardless of packet size, traffic pattern, or enabled features on 10 Gigabit Ethernet interface

- ▶ Support for up to 8 Kb MAC addresses
- ▶ Support for jumbo frames (up to 9216 bytes)
- ▶ Support for up to 512 VLANs; VLAN numbers from 1 to 4000

The at-a-glance guide with more detailed technical information about the Cisco Nexus 40011 for IBM BladeCenter is available from:

<http://www.redbooks.ibm.com/abstracts/tips0754.html>

### **QLogic Virtual Fabric Extension Module for IBM BladeCenter**

The QLogic Virtual Fabric Extension Module works in combination with the BNT Virtual Fabric 10 Gb Switch Module. If you use a converged network adapter such as the QLogic 2-port 10Gb CNA in the blade server, all packets flow to the BNT high speed switch module where it is split up. LAN traffic exits the BNT switch on its external ports. SAN traffic is internally rerouted to the QLogic Extension Module and exits the module on native Fibre Channel ports. This provides a fully integrated FCoE solution within an IBM BladeCenter chassis, without the cost for an additional top of rack FCoE switch or gateway.



*Figure 2-10 QLogic Virtual Fabric Extension Module for IBM BladeCenter*

Technical features of the QLogic Virtual Fabric Extension Module are:

- ▶ Standard IBM BladeCenter I/O module form factor
- ▶ Six external auto-sensing Fibre Channel ports (8 Gbps, 4 Gbps, or 2 Gbps)
- ▶ Support for transparent mode (NPIV)

The at-a-glance guide with more detailed technical information about the QLogic Virtual Fabric Extension Module for IBM BladeCenter is available from:

<http://www.redbooks.ibm.com/abstracts/tips0717.html>

## **2.2.3 FCoE cabling systems**

The use of a structured cabling system for FCoE is strongly recommended (even for small installations). A structured cabling system provides a protected solution that serves current requirements as well as allowing for easy data center expansion.

A structured fiber cabling system for FCoE provides the following benefits:

- ▶ Reduction in the number of cables under raised floor or overhead, which makes tracking the destination of each cable easier and facilitates tracing of fiber links when problems occur.
- ▶ Disruption under the raised floor caused by moves, adds, and changes is minimized; costs incurred are reduced as well.
- ▶ With individual hardware devices (servers, SAN storage, tape, and so forth) connecting to local patch panels (in a distribution area) with short cables, connectivity is handled in the vicinity of the hardware/patch panels, rather than across the data center.
- ▶ Faster, less disruptive installation and removal of equipment, ease of reconfiguration of equipment, and more efficient use of the under floor space.
- ▶ Potential to improve air movement and reduce air conditioning needs.

For more detailed information about the advantages of structured cabling, refer to Appendix B.3 in the IBM Redbooks publication *IBM SAN Survival Guide*, SG24-6143, available from:

<http://www.redbooks.ibm.com/abstracts/sg246143.html>

### **Cabling system component highlights**

This section describes the cabling offerings from Panduit, which are an integral part of an FCoE solution.

#### ***Panduit 10Gig™ SFP+ Direct Attach Passive Copper Cable Assemblies***

PSF1PXD\*MBU where \*= length in meters to 5 meters

To satisfy the demanding requirements of FCoE (low power server inter-connect and low latency), Panduit offers 10Gig™ SFP+ Direct Attach Passive Copper Cable Assemblies as shown in Figure 2-11.



Figure 2-11 Panduit 10Gig™ SFP+ Direct Attach Passive Copper Cable

By incorporating 10Gig™ SFP+ Direct Attach Passive Copper Cable Assemblies into FCoE server-to-switch deployments, 10 Gigabit port-to-port performance is achieved without additional signal processing or conversion, providing a low power, low latency 10 Gbps option for top of rack switching applications.

These assemblies are 100% factory tested and ensure 10 Gbps performance when mated to any SFF-8431 SFP+ compliant host port. Their design features provide:

- ▶ Robust, low latency performance: High speed 10 GHz twin-axial cable, with two shielded parallel pairs, factory terminated to SFP+ 10 Gbps hot pluggable modular connectors; rated to a minimum 250 mating cycles.
- ▶ Small diameter cable design: Enables proper cable management and improved air flow.

- ▶ Built-in strain relief: Robust, easy-to-use connector latching feature allows easy installation in belly-to-belly high density installations and ensures proper bend radius control for reliable deployment.
- ▶ Variety of lengths: Available 0.5 to 4 meters (in 0.5 meter increments), 5, 6, and 7 meters for design flexibility.

### **Panduit QuickNet SFQ Cassettes**

FQXO-12-10= Optimized SFQ Fiber Cassette 10 GbE 50µm OM3 12-Fiber 6 Duplex LC to (1) MTP



Figure 2-12 Panduit QuickNet SFQ Cassettes

Pre-terminated high density modular fiber optic cassettes comply with IEEE 802.3ae 10 GbE and ANSI T11.2 Fibre Channel requirements. Fiber optic cassettes support network data rates up to 10 Gbps for link lengths up to 300 meters using laser optimized OM3 fiber.

Optimized cassettes provide insertion loss of 0.5 dB maximum, and standard cassettes provide 1.0 dB maximum to meet IEEE 802.3ae maximum channel loss specification of less than 26 dB. Cassettes employ high performance MTP\* connectors on the rear of the units routed to 6 LC adapters on the patch field side. Cassettes interconnect with high density SFF MTP\* ribbon interconnect cable assemblies. Alternatively, they can be supplied hard fixed to distribution-style cabling. High density cassette patch panels hold up to 8 cassettes, allowing up to 96 fiber connections to be deployed in one rack unit (1 RU).

The features of the cassettes include the following:

- ▶ Compatible with QuickNet Patch Panels for up to 96 fiber connections in 1RU with QAPP48HDBL QuickNet 48-port Angled Patch Panel
- ▶ Connect together with MTP cable assemblies as interconnecting network segments (TOR Switch to central patching location)
- ▶ 10 Gb OM3 50/125µm Fiber Cassettes are tested per IEEE 802.3ae 10 GbE to support network transmission speeds up to 10 Gbps for link lengths up to 300 meters at 850 nm; they are backward compatible with all 50/125µm (OM2) MPO or MTP\* system requirements

### **QuickNet HDQ Series High Density Fiber Optic Cassettes**

F1RCXO-9612-10S= QuickNet HDQ Series High Density Fiber Optic Cassette; 1RU 96-fiber; 10 Gb Optimized OM3 12-fiber MTP to LC Pre-terminated



Figure 2-13 Panduit QuickNet HDQ Series High Density Fiber Optic Cassettes

Pre-terminated high density 1U fiber optic cassettes comply with IEEE 802.3ae 10 GbE and ANSI T11.2 Fibre Channel requirements. Fiber optic cassettes support network data rates up to 10 Gbps for link lengths up to 300 meters using laser-optimized OM3 fiber.

Optimized cassettes provide insertion loss of 0.5 dB maximum, and standard cassettes provide 1.0 dB maximum, to meet IEEE 802.3ae maximum channel loss specification of less than 26 dB. Cassettes employ high performance MTP\* connectors on the rear of the units routed to 48 LC adapters on the patch field side. Cassettes interconnect with high density SFF MTP\* ribbon interconnect cable assemblies or can be supplied as hard fixed to distribution-style cabling.

The features of these cassettes include the following:

- ▶ Low insertion loss of 0.5 dB maximum per optimized cassette and 0.75 dB maximum per standard cassette; ensures system meets IEEE 802.3ae maximum channel loss specification of less than 2.6 dB
- ▶ Connect together with high density female MTP\* cable assemblies as interconnection network
- ▶ 10 Gb OM3 50/125µm fiber cassettes are tested per IEEE 802.3ae 10 GbE to support network transmission speeds up to 10 Gbps for link lengths up to 300 meters and up to 550 meters for OM4 at 850 nm; backward compatible with all 50/125µm (OM2) MPO or MTP system requirements

The QuickNet Cassette Rear Cable Manager mounts horizontally to standard rack rails or vertically with Panduit angled cabinet patch panel brackets.

### **QuickNet Patch Panel Components**

QAPP48HDBL, QPP48HDBL - 48-port angled and flat patch panels (accept QuickNet Pre-Terminated Cassettes)



Figure 2-14 Panduit QuickNet Patch Panel

The Panduit QuickNet Cabling System provides a custom, pre-terminated cabling solution that meets unique requirements. Fast and simple to install, the system enables quick network deployment, increased reliability, and lower total cost of ownership as compared to field terminated installations.

Engineered for maximum design flexibility and high rack density utilization (up to 288 fiber ports in one rack space), the system offers 100% factory tested pre-terminated cable assemblies in custom lengths and configurations. QuickNet Angled and Flat Patch Panels accept QuickNet Pre-Terminated Cassettes, Patch Panel Adapters, and Blanks, which snap in and out, with one hand, for quick installation. Features of the panels include the following:

- ▶ Accept QuickNet Copper Cable Assemblies and QuickNet SFQ Series MTP\* Fiber Optic Cassettes, which snap in and out with one hand.
- ▶ Angled patch panels provide bend radius control and minimize the need for horizontal cable managers.
- ▶ Mount to standard EIA 19" racks or 23" racks with optional extender brackets.
- ▶ Optional patch panel adapter (QPPABL) accepts all Mini-Com Modules for UTP, fiber optic, and audio/visual applications.
- ▶ Optional patch panel blank (QPPBBL) promotes proper airflow and cooling.
- ▶ Zero RU brackets accept fiber adapter panels or pre-terminated cassettes by mounting directly to rack or enclosure without utilizing additional rack space.

### ***Harness (Hydra) Assemblies***

FHPX126LM\*\*\*N= Male MTP\* to LC duplex 12-fiber 10 Gb 50/125µm OM3 Multimode Plenum Rated Hydra Cable Assembly; \*\*\*= length in meters

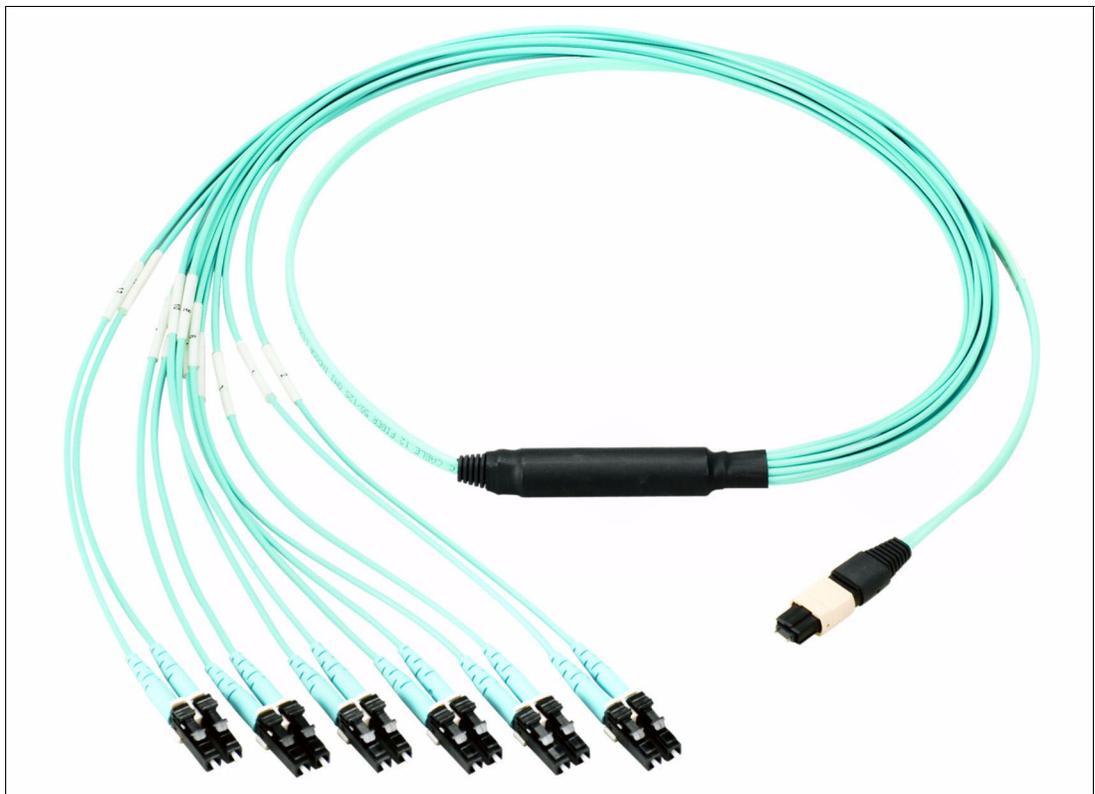


Figure 2-15 Panduit Harness (Hydra) Assembly

Panduit QuickNet Hydra Assemblies allow for rapid deployment of high density, multi-port patch field connectivity for Storage Area Network (SAN) applications. QuickNet Hydra Cable Assemblies optimize SAN patch field organization on the host or switch side, ensuring efficient use of horizontal and vertical rack pathways in the cabinet. QuickNet Hydra Cable Assemblies, built with modular MTP\* connectivity and traditional connectivity (LC), provide compatibility, flexibility, and system performance in high density patch field applications. All Panduit QuickNet Hydra Assemblies are factory terminated and tested to deliver verified optical performance and reliability for improved network integrity. 10 Gb versions provide 10 Gbps network performance up to 300 meters per IEEE 802.3ae 10 GbE standard while maintaining compatibility with legacy systems. Features include:

- ▶ Application-specific designs tailor configuration and breakout construction to application requirements to minimize waste, optimize cable management, speed deployment, and improve flexibility and manageability for lower installation costs.
- ▶ Easily connect network segments and minimize SAN and server patch cord congestion to provide design flexibility, reduce deployment time, improve reliability, and simplify moves, adds, and changes.
- ▶ High density cable uses vertical rack and horizontal patch field pathway space more efficiently to improve manageability and reduce installation costs.
- ▶ Provides low insertion loss and high return loss for improved high performance, minimizing variability for improved channel link loss performance.

***PANDUIT High Density SFQ Small Form Factor QUICKNET Cable Assemblies***

FQPX72V1M\*\*\*N= 72 Fiber 10 Gb plenum rated high density MTP SFQ to MTP connectors cable assembly; \*\*\*= length in meters

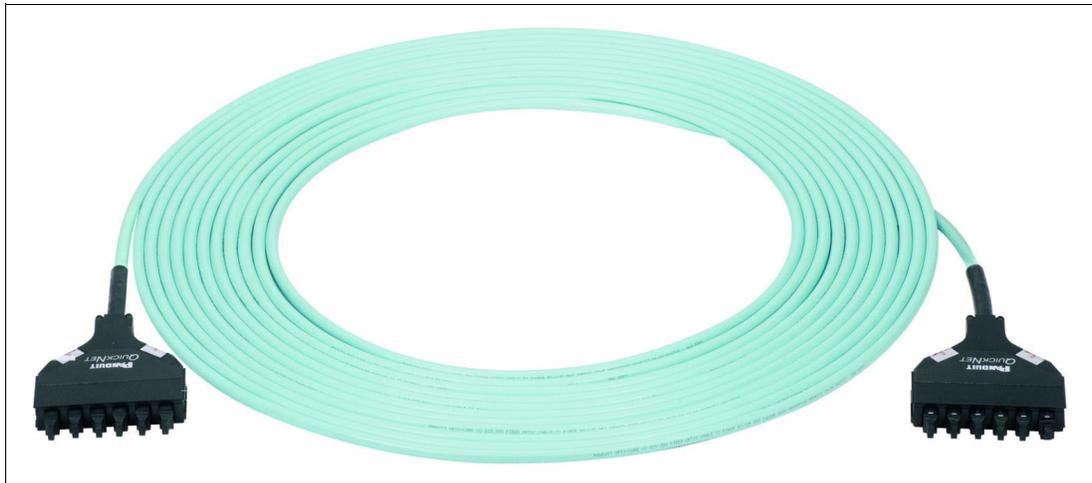


Figure 2-16 Panduit High Density SFQ QuickNet Fiber Optic Cabling

Panduit High Density SFQ QuickNet Fiber Optic Cabling is ideal for SAN EDA applications delivering high reliability and scalability in a cost-effective design. MTP SFQ cassettes achieve up to 72 fiber connections in one cassette, providing up to 576 fiber connections in 1 RU with superior cable access and management capabilities. Panduit QuickNet Fiber and Copper Cassettes feature the same form factor, allowing excellent flexibility within the QuickNet Patch Panels to provide a unified pre-terminated solution.

Panduit QuickNet High Density SFQ Cassette to Cassette Assemblies are 100% tested for link loss and connector loss. Assemblies are built with OM3 laser optimized fiber to provide 10 Gbps Ethernet performance up to 300 meters (per IEEE 802.3ae 10 GbE standard) and 8 Gbps FC performance up to 150m (per ANSI FC-PI-4 8GFC Standard) while maintaining

compatibility with legacy systems. Panduit High Density SFQ Small Form Factor QuickNet Cable Assemblies are part of the full QuickNet Copper and Fiber System. These assemblies provide the following features:

- ▶ High density design maximizes density with up to 72 fibers in one cassette, 576 in 1 RU.
- ▶ Product is compatible with QuickNet patch panels and uses the same cassette form factor as Copper QuickNet patch panels.
- ▶ Identical copper and fiber cassette form factors offer design flexibility and easy maintenance and upgrade, enabling a unified, highly scalable pre-terminated solution.

**PANDUIT QuickNet Pre-Terminated Copper Cable Assemblies**

QCPBCBCBXXxxxN = Plenum, CAT 6A, cassette to cassette assembly, xxx = length in feet

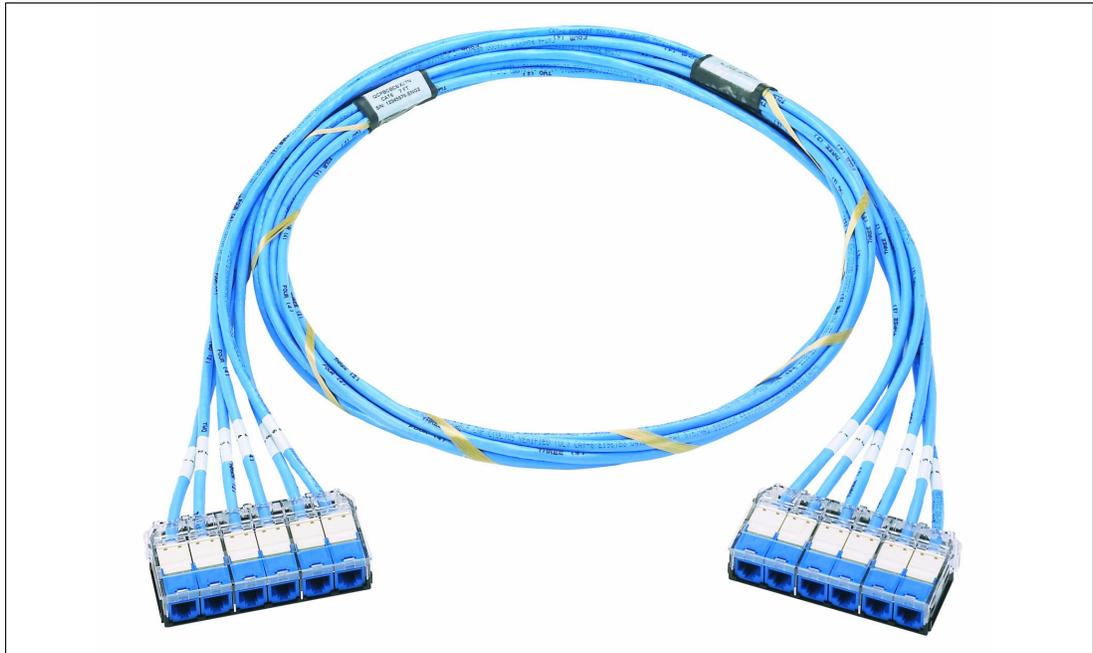


Figure 2-17 Panduit QuickNet Copper Cabling System

The Panduit QuickNet Copper Cabling System provides a custom, pre-terminated cabling solution that can be deployed in the same QuickNet patch panels as the fiber solution (for example, it can be used for management ports on BladeCenter chassis).

Engineered for maximum design flexibility and high rack density utilization (up to 48 ports in one rack space), the system offers 100% factory tested pre-terminated cable assemblies in custom lengths and configurations ideal for intra-rack cabling requirements.

QuickNet Angled and Flat Patch Panels accept QuickNet Pre-Terminated Cassettes, Patch Panel Adapters, and Blanks, which snap in and out, with one hand, for quick installation. Features include the following:

- ▶ QuickNet Cable Assembly is factory tested to electrical permanent link specifications.
- ▶ Copper Cabling System exceeds Category 6A ANSI/TIA/EIA-568-C.2 and ISO 11801 2nd Edition Class Ea standards.
- ▶ Jack modules utilize patented Giga-TX Technology which optimizes performance by maintaining cable pair geometry and eliminating conductor untwist.
- ▶ The plug meets all applicable ANSI/TIA-968-A requirements.

## 2.3 Sample configurations

Fiber connectivity and media systems are agnostic with respect to transmission protocols (FC, FCoE, and Ethernet). Fiber infrastructure solutions based on good practice using a structured cabling approach (as outlined in TIA-942), are managed, with respect to moves, additions, and changes, in a main distribution area (MDA). Port mapping of switching electronics and correspondence to hosts or server pod groups is accomplished through intuitive port configuration and identification as provided by modular fiber cassette solutions in the MDA. The SFP+ converged network solution will seamlessly integrate into an existing fiber cable plant based on the 942 standard, which calls for laser-optimized multimode fiber and a structured and logical approach for all cabling elements.

Fiber connectivity systems on converged electronics are identical to those used in non-converged systems (based on SFF transceivers) and the fiber cable media present in these structured cabling solutions remains unchanged.

The server pod migrations to converged networks, as described in this document, will present minimal disruption or change in design for the fiber cable plant. These migrations are presented in “Day 1” and “Day 2” scenarios, which show the migration from server hosts with “Top-of-Rack” LAN with a separate home run or “collapsed” FC network to several options of converged server pods that still utilize the collapsed FC fiber cabling infrastructure.

- ▶ **“Day 1” architecture** - Intra-row edge LAN switching with collapsed Fibre Channel network

LAN distribution layer switching is deployed either between racks (larger switches servicing servers in multiple racks) or within racks (smaller 1U switches at Top of Rack). FC is delivered to server HBAs and runs back to SAN MDA where SAN edge/core (and cross connect) coexist. FC fiber patch facility either exists within each server rack or is co-located with larger LAN switches that service multiple racks and require a certain number of fibers to connect distribution to access layer (in this case FC HBA connection is direct attach from this HDA).

- ▶ **“Day 2” architectures** - In-rack or in-row converged switching with collapsed Fibre Channel network.

Three scenarios are considered (see server pod elevations that follow):

- a. In-chassis switching - IBM BladeCenter H or HT with converged switches mounted in the Chassis
- b. Discrete server (x3550 M2, x3650 M2, and so forth) with pass through to external converged network switches at the end of each row
- c. Discrete server (x3550 M2, x3650 M2, and so forth) with converged network switching deployed in top-of-rack with end-of-row patching

This section describes four specific configurations that implement an FCoE solution. Two are BladeCenter-based configurations of only one rack, and two are rack-server-based configurations that fill an entire row of rack cabinets.

The configurations are:

- ▶ 2.3.1, “BladeCenter H configuration” on page 28
- ▶ 2.3.2, “BladeCenter HT configuration” on page 30
- ▶ 2.3.3, “Top-of-rack LAN switching and end-of-row SAN switching (LAN aggregation switching)” on page 32
- ▶ 2.3.4, “Top-of-rack LAN/SAN switching and end-of-row patching” on page 35

## 2.3.1 BladeCenter H configuration

Figure 2-18 shows an FCoE deployment in a 19" rack using multiple IBM BladeCenter H chassis and servers with FCoE support hardware. Note that we show the rear of the rack because the BladeCenter H is designed to have all cabling at the rear of the chassis.

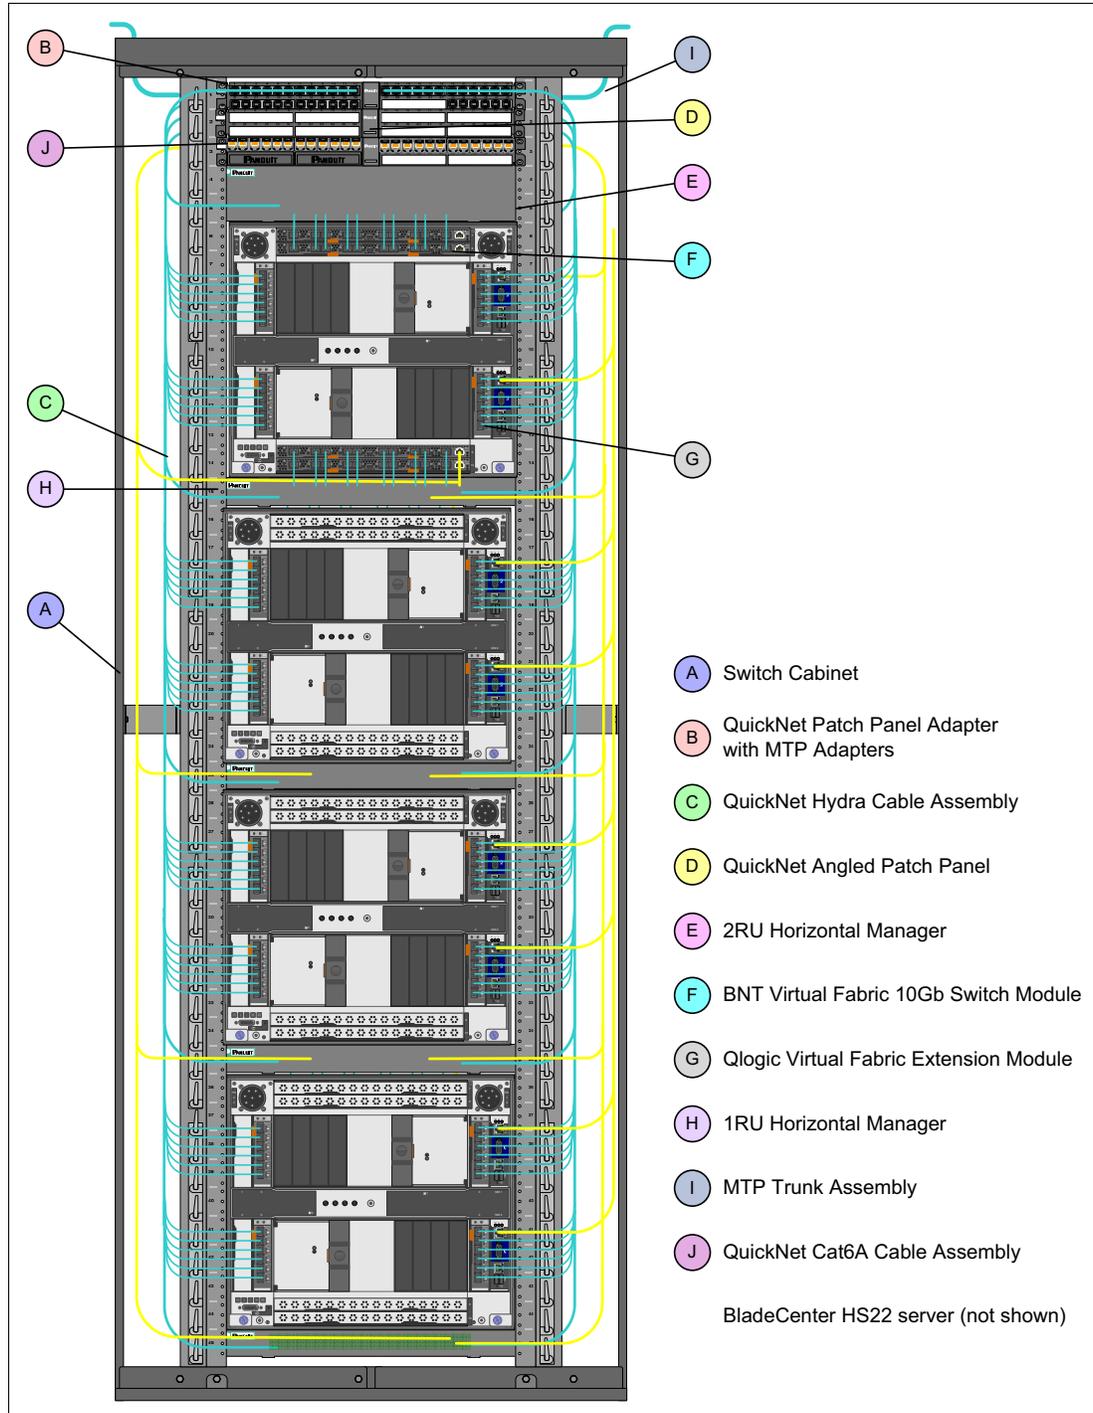


Figure 2-18 IBM BladeCenter H-based configuration (rear view of rack)

Four BladeCenter H chassis can be housed in one 45 U cabinet. This allows for 1 U separation between the switches and 5 U at the top of the cabinet for patching and horizontal

cable management. When the H chassis is fully equipped, it can support up to 14 blade servers and it requires approximately 5800 W. Total power consumption for the cabinet with four BladeCenter H chassis installed is 23,200 W.

Details about the connectivity and accessory cards necessary to outfit the rack shown in Figure 2-18 are presented in Table 2-8.

Table 2-8 Bill of materials for Figure 2-18 on page 28

Callout	Part number	Description	Quantity
None	8852	BladeCenter H chassis.	4
F	46C7191	BNT Virtual Fabric 10 Gb Switch Module.	8
G	42C1830	QLogic 2-port 10 Gb Converged Network Adapter (CFFh).	56
None	7870	BladeCenter HS22 server (not shown due to rear view shown)	56
A	CN1	Cabinet frame with top panel. Dual hinge perforated front door opens to the left or right. Two sets of #12-24 threaded equipment mounting rails. 45 RU cable management on front and rear of front posts. Dimensions: 84.0"H x 31.5"W x 41.1"D (2134 mm x 800 mm x 1044 mm).	1
B	QPPAV612BLMTP/N	QuickNet Patch Panel Adapter loaded with 6 vertically oriented 12-fiber MPO/MTP* Fiber Optic Adapters, (black), non-standard.	16
I	FX12D5-5M**Y	Female MTP* to female MTP* 48-fiber 10 Gb 50/125µm OM3 Multimode Plenum Rated Trunk Cable Assembly (** denotes length in meters).	24
D	QAPP48HDBL	Patch panel that accept up to eight QuickNet SFQ Series MTP* Fiber Optic Cassettes.	2
C	FHPX126LM***N	Male MTP to LC duplex 12F 10 Gb 50/125µm OM3 Hydra Assembly.	32
J	QARBCBCBXX**	QuickNet Cat6A Cable Assembly with pre-terminated jack module cassettes on each end for systems management (** denotes length).	4
None	UTP6A**	Cat6A patch cord for systems management (** denotes length).	24
H	NMF1	1RU Horizontal Cable Manager 1.7"H x 19.0"W x 3.7"D (44 mm x 483 mm x 94 mm).	4
E	NMF2	2RU Horizontal Cable Manager 3.5"H x 20.2"W x 3.7"D (89 mx 573 mm x 94 mm).	1

## 2.3.2 BladeCenter HT configuration

Figure 2-19 shows an FCoE deployment in a 19" rack using multiple IBM BladeCenter HT chassis and servers with FCoE support hardware. Note that we show the front of the rack because the BladeCenter HT is designed to have all cabling at the front of the chassis.

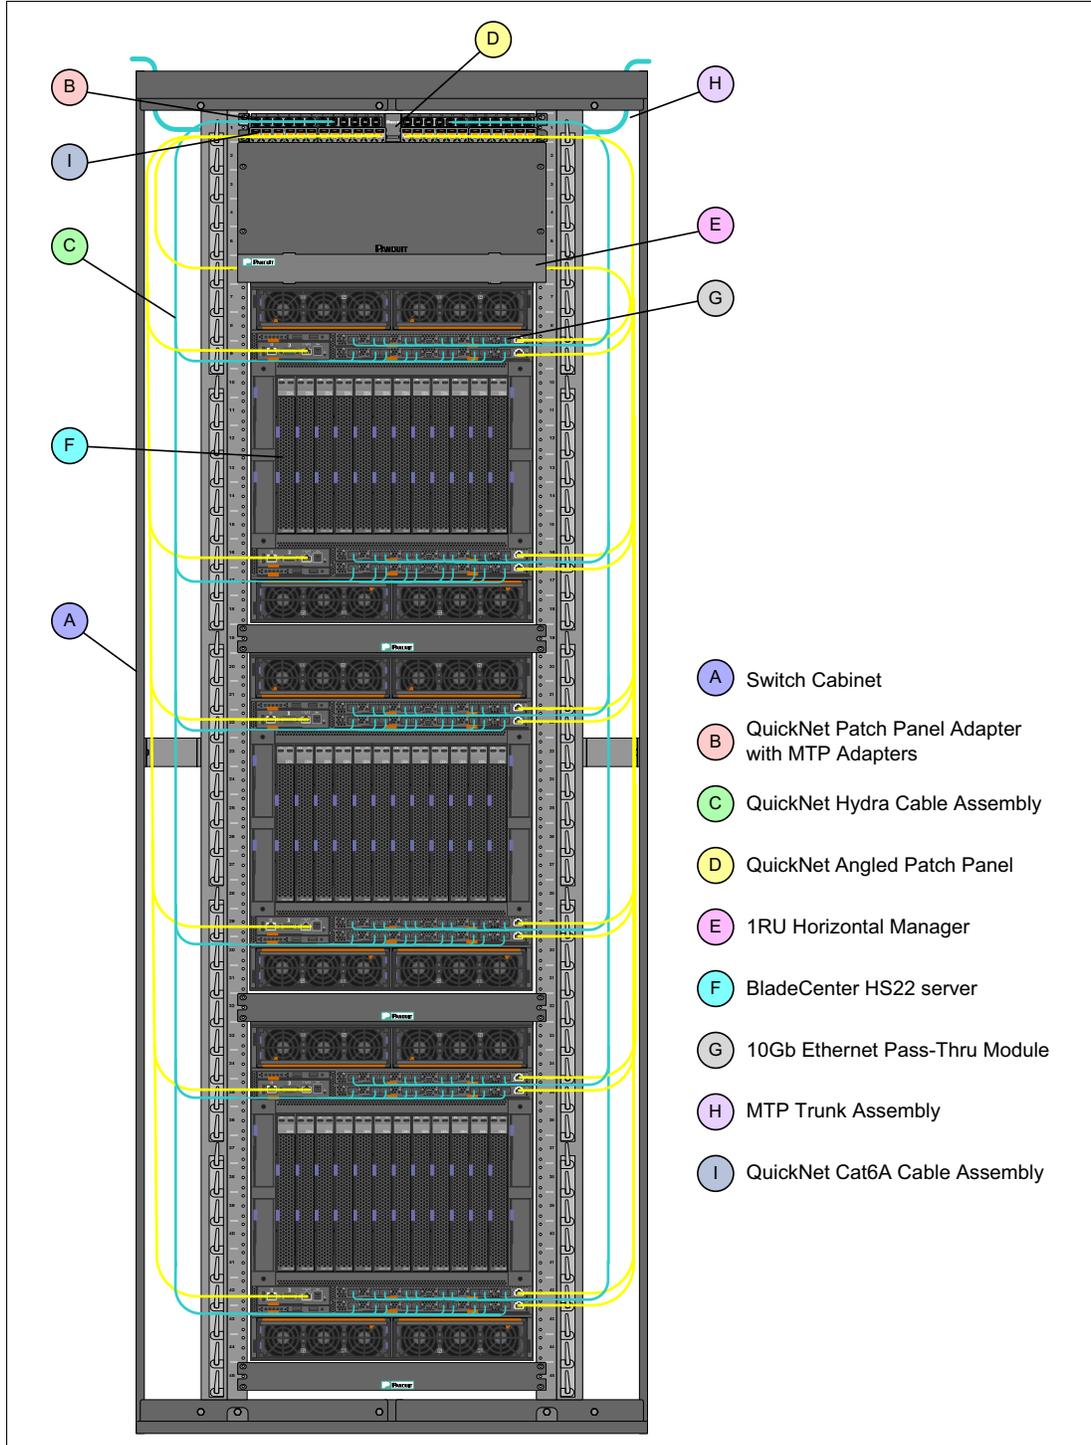


Figure 2-19 IBM BladeCenter HT-based configuration (front view of rack)

Three BladeCenter HT chassis can be housed in one 45 U cabinet. This allows for 1 U separation between the switches and 6 U at the top of the cabinet for patching and horizontal cable management. When the HT chassis is fully equipped, it can support up to 12 blade servers and it requires approximately 3160 W. Total power consumption for the cabinet with three BladeCenter HT chassis installed is 18,960 W.

Details about the connectivity and accessory cards necessary to outfit the rack shown in Figure 2-19 are presented in Table 2-9.

Table 2-9 Bill of materials for Figure 2-19 on page 30

Callout	Part number	Description	Quantity
None	8740/8750	BladeCenter HT	3
None	46M6168	Broadcom 10 Gb Gen 2 4-port Ethernet Expansion Card (CFFh)	12
G	46M6181	10 Gb Ethernet Pass-Thru Module	6
F	7870	IBM BladeCenter HS22 server	36
A	CN1	Cabinet frame with top panel. Dual hinge perforated front door opens to the left or right. Two sets of #12-24 threaded equipment mounting rails. 45 RU cable management on front and rear of front posts. Dimensions: 84.0"H x 31.5"W x 41.1"D (2134 mm x 800 mm x 1044 mm)	1
B	QPPAV612BLMTP/N	QuickNet Patch Panel Adapter loaded with 6 vertically oriented 12-fiber MPO/MTP* Fiber Optic Adapters, (black), non-standard.	4
H	FX12D5-5M**Y	Female MTP* to female MTP* 48-fiber 10 Gb 50/125µm OM3 Multimode Plenum Rated Trunk Cable Assembly (** denotes length in meters)	3
D	QAPP48HDBL	Patch panel that accept up to eight QuickNet SFQ Series MTP* Fiber Optic Cassettes.	2
C	FHPX126LM***N	Male MTP to LC duplex 12F 10 Gb 50/125µm OM3 Hydra Assembly	18
I	QARBCBCBXX**	QuickNet Cat6A Cable Assembly with pre-terminated jack module cassettes on each end for systems management (** denotes length)	2
None	UTP6A**	Cat6A patch cord for systems management (** denotes length)	54
E	NMF1	1U Horizontal Cable Manager 1.7"H x 19.0"W x 3.7"D (44 mm x 483 mm x 94 mm)	4

### 2.3.3 Top-of-rack LAN switching and end-of-row SAN switching (LAN aggregation switching)

The architecture shown in Figure 2-20 is representative of a traditional 1G server layout with top-of-rack LAN switching and end-of-row SAN switching.

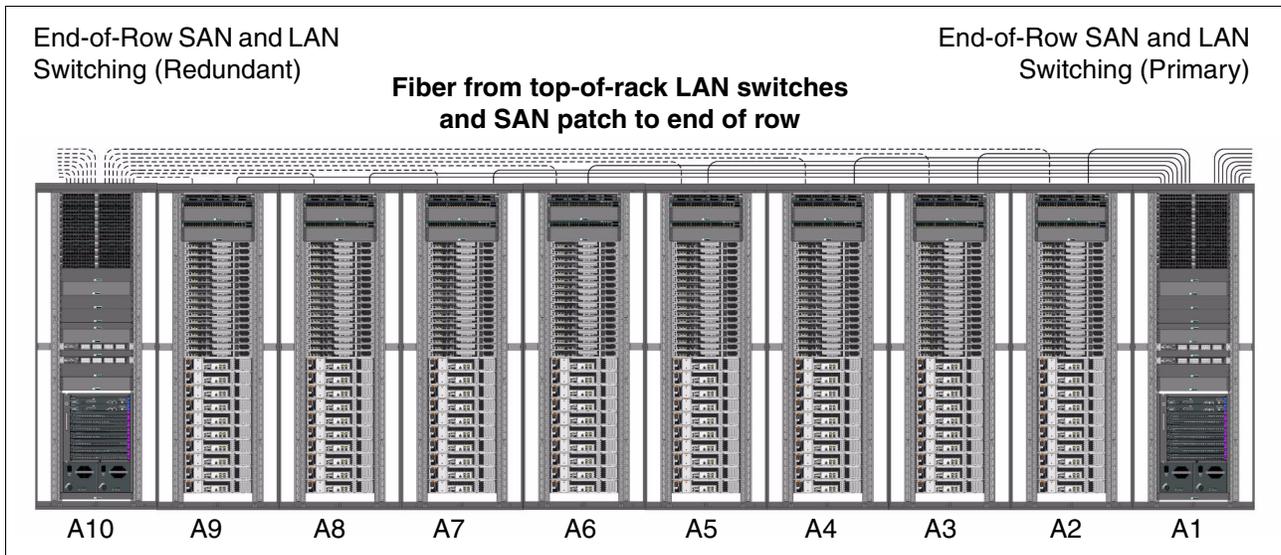


Figure 2-20 Full row of racks: top-of-rack LAN and end-of-row SAN

Figure 2-21 on page 33 shows a detailed view of the racks. It shows:

- ▶ 24 IBM x3550 M2 servers with standard Broadcom BCM5709 Gb Ethernet controllers
- ▶ Eight IBM x3550 M2 servers with both Broadcom BCM5709 Gb Ethernet controllers and QLogic 10 Gb CNA Adapters

The GigE controllers from the servers are patched to a top-of-rack Cisco 4948-10GE switch and the CNA adapters from the servers are patched to a top-of-rack LC patch panel which is then connected to end-of-row IBM converged switch 32B. The fiber uplinks from the LAN switches are run to an end-of-row patch panel via LC-to-LC duplex patch cords and cross-connected to the respective Cisco 6509 aggregation layer switch ports. The end-of-row SAN switch uplinks and the Cisco 6509 uplinks are patched to a top-of-rack patch panel and run back to the LAN and SAN main distribution areas.

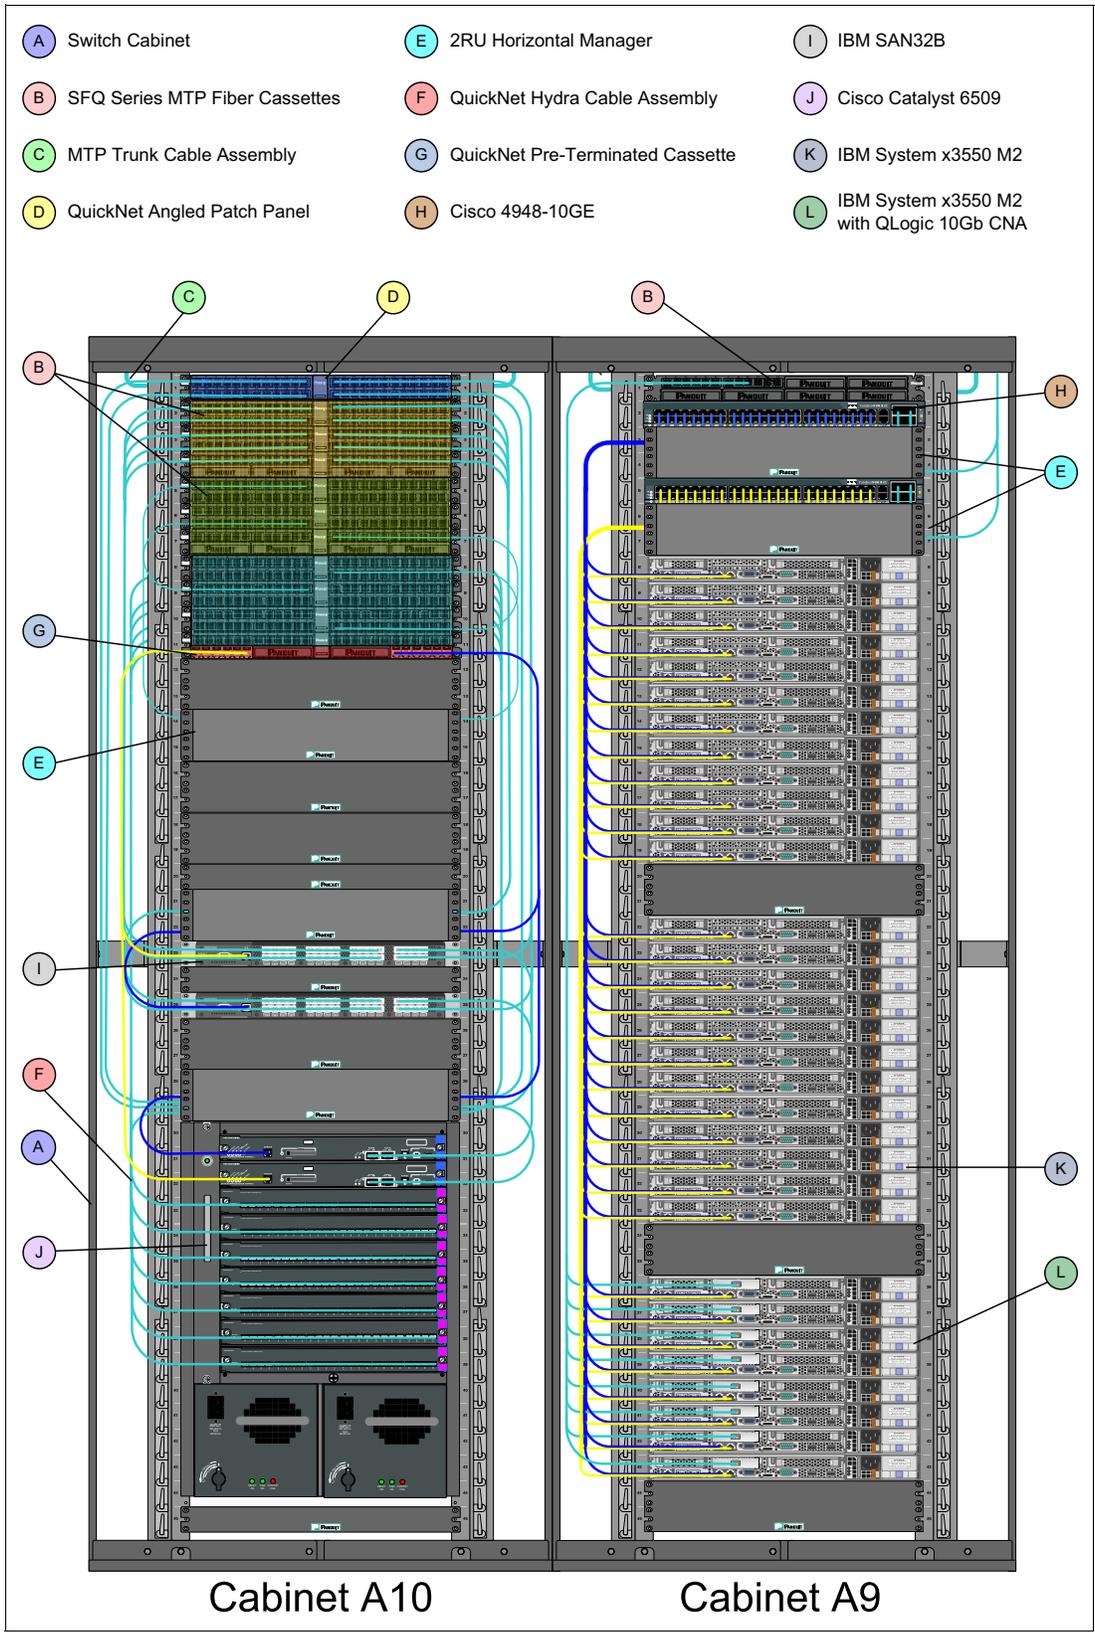


Figure 2-21 Top-of-rack LAN switching and end-of-row SAN switching

The following two tables present detailed information about the configuration shown in Figure 2-21. Table 2-10 shows the details for a typical server cabinet (A9 in the figure); Table 2-11 shows the details for one end-of-row cabinet (A10 in the figure).

Table 2-10 Typical Server Cabinet (A9 in Figure 2-21 on page 33)

Callout	Part number	Description	Quantity
A	CN1	Cabinet frame with top panel. Dual hinge perforated front door opens to the left or right. Two sets of #12-24 threaded equipment mounting rails.45 RU cable management on front and rear of front posts. Dimensions: 84.0"H x 31.5"W x 41.1"D (2134 mm x 800 mm x 1044 mm)	1
H	4948-10GE	Cisco Catalyst top of rack 1U Ethernet switch that offers 48 ports of wire-speed 10/100/1000BASE-T with four alternative wired ports that can accommodate optional 1000BASE-X Small Form-Factor Pluggable (SFP) optics.	2
K	x3550M2	IBM System x3550 M2 1U server with standard Broadcom BCM5709 GigE controllers	24
L	x3550M2 w/CNA adapter	IBM System x3550 M2 1U server with standard Broadcom BCM5709 GigE controllers and QLogic 10 Gb CNA	8
None	FXE10-10M**Y	10G OM3 LC to LC duplex patch cords (** denotes length in meters)	12
D	QAPP48HDBL	Patch panel that accept up to eight QuickNet SFQ Series MTP* Fiber Optic Cassettes.	1
None	UTP6A**	Cat6A patch cord for systems management (** denotes length in feet)	64
E	NMF2	2U Horizontal Cable Manager 3.5"H x 19.0"W x 3.7"D (89 mm x 483 mm x 94 mm)	1

Table 2-11 Typical End-of-Row Cabinet (A10 in Figure 2-21 on page 33)

Callout	Part number	Description	Quantity
A	CN1	Cabinet frame with top panel. Dual hinge perforated front door opens to the left or right. Two sets of #12-24 threaded equipment mounting rails.45 RU cable management on front and rear of front posts. Dimensions: 84.0"H x 31.5"W x 41.1"D (2134 mm x 800 mm x 1044 mm)	1
J	6509-E	9-slot Cisco Catalyst 6509-E Switch provides high port densities that are ideal for many wiring closet, distribution, and core network as well as data center deployments.	2
I	SAN32B	IBM SAN32-B SAN Fabric 1U Switch that provides 24 or 32 active ports and is designed for high performance with 8 Gbps link speeds and backward compatibility to support links running at 4, 2 and 1 Gbps link speeds	2
B	FQXO-12-10	10 Gb 50/125µm (OM3) MM SFQ Series MTP* Cassette	76
None	FXE10-10M**Y	10G OM3 LC to LC duplex patch cords (** denotes length in meters)	***a
F	FHPX126LM***N	Male MTP* to LC Duplex 12F 10 Gb 50/125µm OM3 MM Hydra Cable Assembly	14
D	QAPP48HDBL	Patch panel that accept up to eight QuickNet SFQ Series MTP* Fiber Optic Cassettes.	11
None	UTP6A**	Cat6A patch cord for systems management (** denotes length in feet)	4

Callout	Part number	Description	Quantity
E	NMF2	2U Horizontal Cable Manager 3.5"H x 19.0"W x 3.7"D (89 mm x 483 mm x 94 mm)	3
C	FSPX24-55F***A	Female MTP* to MTP* 10 Gb 50/125µm OM3 MM trunk assembly with pulling eye on one end (***) denotes length in feet)	4
G	QARBCBCBXX**	QuickNet Cat6A Cable Assembly with pre-terminated jack module cassettes on each end for systems management (** denotes length)	2

a. \*\*\* in Quantity of cables depends on number of ports utilized

### 2.3.4 Top-of-rack LAN/SAN switching and end-of-row patching

The architecture shown in Figure 2-22 is representative of a traditional 1G server layout with 10G uplinks connecting the LAN/SAN switches and directors.

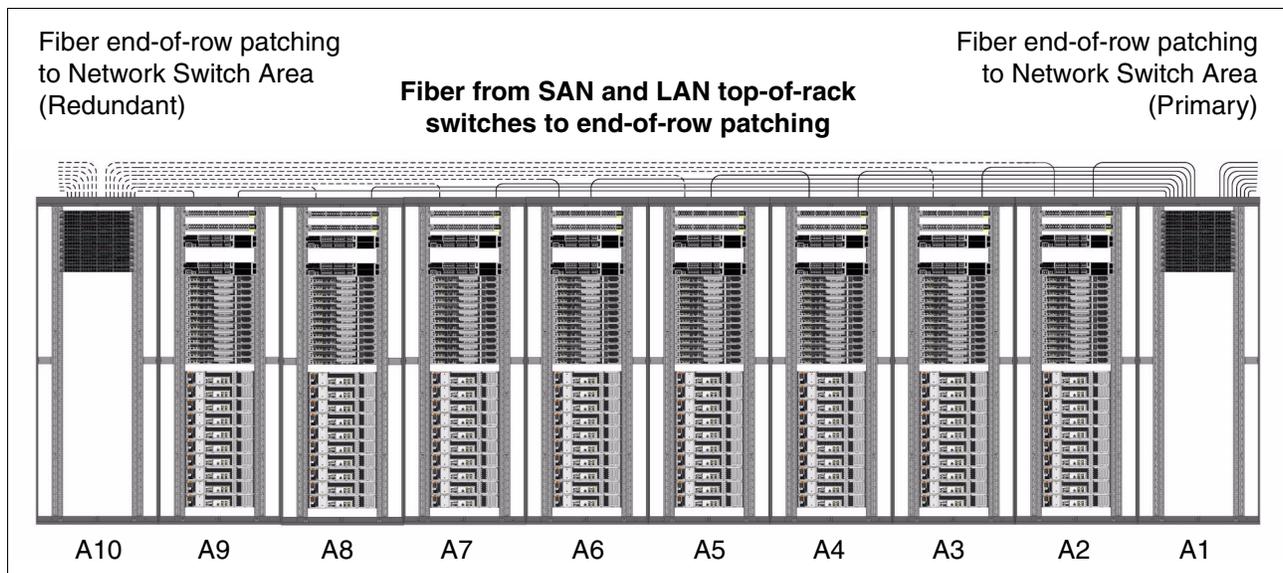


Figure 2-22 Full row: Top-of-rack LAN/SAN Switching and end-of-row patching

Figure 2-23 on page 36 shows detailed views of the racks. It shows:

- ▶ 20 IBM x3550 M2 servers with standard Broadcom BCM5709 GigE controllers
- ▶ Eight IBM x3550 M2 servers with both Broadcom BCM5709 GigE controllers and QLogic 10 Gb CNA Adapters.

The GigE controllers from the servers are patched to a top-of-rack Cisco 4948-10GE switch and the CNA adapters from the servers are patched to a top-of-rack IBM System Storage™ SAN40B-4. The fiber uplinks from the respective switches are run to an end-of-row patch panel via LC-to-LC duplex patch cords. The end-of-row patch then connects to the LAN and SAN equipment distribution areas.



The following two tables present detailed information about the configuration shown in Figure 2-23. Table 2-12 shows the details for a typical server cabinet (A9 in the figure); Table 2-13 shows the details for one end-of-row cabinet (A10 in the figure).

Table 2-12 Typical server cabinet (A9 in Figure 2-23 on page 36)

Callout	Part number	Description	Quantity
A	CN1	Cabinet frame with top panel. Dual hinge perforated front door opens to the left or right. Two sets of #12-24 threaded equipment mounting rails.45 RU cable management on front and rear of front posts. Dimensions: 84.0"H x 31.5"W x 41.1"D (2134 mm x 800 mm x 1044 mm)	1
H	4948-10GE	Cisco Catalyst top of rack 1U Ethernet switch that offers 48 ports of wire-speed 10/100/1000BASE-T with four alternative wired ports that can accommodate optional 1000BASE-X Small Form-Factor Pluggable (SFP) optics	2
G	IBM SAN40B-4	IBM SAN40-B SAN Fabric 1U Switch that provides 24, 32, or 40 active ports and is designed for high performance with 8 Gbps link speeds and backward compatibility to support links running at 4, 2 and 1 Gbps link speeds	2
I	x3550M2	IBM System x3550 M2 1U server with standard Broadcom BCM5709 GigE controllers	20
J	x3550M2 w/CNA adapter	IBM System x3550 M2 1U server with standard Broadcom BCM5709 GigE controllers and QLogic 10Gb CNA	8
None	FXE10-10M**Y	10G OM3 LC to LC duplex patch cords (** denotes length in meters)	****a
None	UTP6A**	Cat6A patch cord for systems management (** denotes length in feet)	****a
C	QAPP48HDBL	Patch panel that accept up to eight QuickNet SFQ Series MTP* Fiber Optic Cassettes	1
E	NMF1	2U Horizontal Cable Manager 3.5"H x 19.0"W x 3.7"D (89 mm x 483 mm x 94 mm)	1
E	NMF2	1U Horizontal Cable Manager 1.7"H x 19.0"W x 3.7"D (44 mm x 483 mm x 94 mm)	2

a. \*\*\* in Quantity of cables depends on number of ports utilized

Table 2-13 Typical end-of-row cabinet (A10 in Figure 2-23 on page 36)

Callout	Part number	Description	Quantity
A	CN1	Cabinet frame with top panel. Dual hinge perforated front door opens to the left or right. Two sets of #12-24 threaded equipment mounting rails.45 RU cable management on front and rear of front posts. Dimensions: 84.0"H x 31.5"W x 41.1"D (2134 mm x 800 mm x 1044 mm)	1
D	FQXO-12-10	10 Gb 50/125µm (OM3) MM SFQ Series MTP* Cassette	40
None	FXE10-10M**Y	10G OM3 LC to LC duplex patch cords (** denotes length in meters)	****a
C	QAPP48HDBL	Patch panel that accept up to eight QuickNet SFQ Series MTP* Fiber Optic Cassettes.	8
None	UTP6A**	Cat6A patch cord (** denotes length in feet)	****a

Callout	Part number	Description	Quantity
B	FSPX24-55F***A	Female MTP* to MTP* 10 Gb 50/125µm OM3 MM trunk assembly with pulling eye on one end (***) denotes footage in feet)	20
F	QARBCBCBXX**	QuickNet Cat6A Cable Assembly with pre-terminated jack module cassettes on each end (** denotes length)	20

a. \*\*\* in Quantity of cables depends on number of ports utilized

## 2.4 FCoE deployment scenarios

Fibre Channel Industry Association (FCIA), which is the standards committee for Fibre Channel and related technologies, defines five scenarios for deployment of FCoE technology in data centers<sup>2</sup> and lays out whether FCoE will be a likely candidate for it.

The five scenarios are:

1. Existing data centers with large investment in FC storage
2. New data center installations
3. Existing data centers with mix of storage (DAS, FC, NAS) requiring near-term data center consolidation
4. Small to medium data centers
5. Remote office installations

We describe solutions to each of the first four of these FCIA scenarios in the rest of this section, using the configurations described as 2.3, “Sample configurations” on page 27 as the basis of our discussion.

Remote office installations are considered to be the least likely targets of FCoE migrations because iSCSI offers the benefit of native IP routing, which allows for longer distances than does FCoE. We therefore do not cover scenario 5 in this section.

Scenario 5 requires cable lengths of several kilometers to reach remote offices. This channel length is currently not being supported while using PFC functionality (see 2.1.3, “FCoE physical media considerations” on page 9). Again, it is important to understand the limitations of FCoE equipment on both ends of the channel and to consult equipment manufacturers’ specifications to find stated reach while using the PFC function.

### 2.4.1 Scenario 1: Existing data centers with large investment in FC storage

Existing data centers with large FC environments are likely candidates for rapid FCoE deployments to leverage and extend investments in their FC infrastructure. Initial FCoE deployments probably will start at the edge with new investments being made into FCoE-enabled switches and CNAs (instead of FC-only switches and HBAs).

### 2.4.2 Scenario 2: New data center installation

For new data centers, the choice between FC, FCoE, and iSCSI is less clear. The two deciding factors are existing employee skill sets and time of data center installation. Companies with employees who are more familiar with Ethernet than Fibre Channel are more

<sup>2</sup> See “FCIA Fibre Channel Solutions Guide Fall 2009”,  
[http://www.fibrechannel.org/documents/doc\\_download/17-fcia-solution-guide-fall-09](http://www.fibrechannel.org/documents/doc_download/17-fcia-solution-guide-fall-09)

likely to stick with iSCSI and not go with FC/FCoE. The same is true for companies that already have employees with a great deal of Fibre Channel expertise. They are more likely to opt for FCoE.

The time of data center installation becomes relevant because FCoE hardware is, as of today, still considered leading edge. The likelihood that a new data center installation goes with iSCSI or FC will significantly decrease during and after 2010, when more companies transition to FCoE and the technology becomes a more well known commodity.

In such greenfield data centers, IT architects will opt for the flexibility of Top-of-Rack converged architectures. This type of architecture (Unified Fabric access layer network) offers low cost copper within the server racks with a lower number of server adapters. Installation time for cabling infrastructure is reduced for such network segments that require less cabling.

FCoE interconnect is only available on equipment and network adapters with SFP+ interfaces. So, BladeCenter server or discrete server applications (like IBM Systems x3650 servers) will use SFP+ SR or LR optical transceivers (between racks and between rows) or SFP+ direct attach copper media for applications up to 7 m (within racks and between racks).

Server pod elevations are considered as design candidates for such greenfield applications.

### **2.4.3 Scenario 3: Existing data centers with mix of storage (DAS, FC, NAS) requiring near-term data center consolidation**

Today, many data centers try to consolidate their storage for any of the following reasons:

- ▶ To reduce cost through server virtualization
- ▶ To reduce environmental costs (power and cooling)
- ▶ To improve manageability (reduced labor expenses)
- ▶ To improve service levels for customers (data accessibility, test/dev)
- ▶ Because of data center space constraints

For this consolidation scenario, the FCIA sees the existing storage environments as the deciding factor for deciding which new technology to invest in. This allows the company to leverage existing storage investments as well as to continue to use the same tools and available skill sets. Companies with lots of network attached storage (NAS) will most likely opt for iSCSI as the preferred remote storage access method. If the company infrastructure is dominated by Fibre Channel, then FCoE is likely to be deployed.

### **2.4.4 Scenario 4: Small to medium data centers**

For small to medium data centers, FCoE and iSCSI are considered to be the best solution. Deciding factors between them are budget, IT personnel skill set, and application requirements. For cost-sensitive, lower end installations iSCSI-based storage networks are sufficient. For high-availability or performance-oriented databases, FCoE becomes a viable option.





## Benefits and business impact

As discussed in Chapter 1, “Introduction” on page 1, a unified fabric consolidates different traffic types onto a single network. This single unified network simplifies the network infrastructure, reduces costs, and increases the efficiency and utilization of the data center. In addition to reducing complexity and TOC, a unified fabric also supports the broader move towards virtualization and consolidation.

FCoE deployed with 10 Gigabit Ethernet allows Fibre Channel to connect every server within an organization to any pool of storage, regardless of their locations, which is the essence of virtualization and a unified networking fabric. FCoE is an evolutionary technology that allows organizations to adopt it at the pace that best serves the organization’s business needs because it works with the existing IT infrastructure. FCoE simplifies the network topology while reducing networking equipment cost, cabling cost, and by eliminating approximately half of the I/O adapter cards. This reduces the power and cooling requirements and increases staffing efficiencies – all while improving bandwidth by taking advantage of 10 Gbps Ethernet.

This chapter describes the advantages of a converged fabric in terms of business benefits. The four key area we describe are:

- ▶ Energy savings
- ▶ A simpler datacenter
- ▶ Supporting future needs
- ▶ Reduced total cost of ownership

### Energy savings

One of the most important benefits of FCoE is the reduction in electricity costs that an organization can realize. Using fewer switches and NICs reduces the growth in power densities, and that results in a lighter load on the cooling infrastructure and an increase in the life of the equipment.

For example, using the top-of-rack LAN switching and end-of-row SAN switching scenario shown in 2.3.3, “Top-of-rack LAN switching and end-of-row SAN switching (LAN aggregation switching)” on page 32 and assuming 12.5 W as the average power consumption for a Fibre Channel HBA, you can save approximately 400 watts by implementing FCoE. For the top-of-rack LAN/SAN switching and end-of-row patching configuration depicted in 2.3.4,

“Top-of-rack LAN/SAN switching and end-of-row patching” on page 35, the energy saved is approximately 350 watts.

As mentioned previously, a direct effect of reducing power consumption is easing the load on the cooling infrastructure. Every watt that is used by equipment results in heat that needs to be removed, requiring increased cooling, resulting in yet increased power consumption to run the cooling infrastructure. Depending on the efficiency of the data center, every watt consumed by the electronics requires an additional 0.8 to 1.5 watts<sup>1</sup> for additional cooling. This implies that an additional 630 W to 1,000 W can be saved per rack — a combined power savings of 980 W to 1.4 KW.

Also note that the power savings are even higher if you take into account the savings in power consumption directly due to fewer switches and the power required to cool those switches.

## A simpler data center

Along with the reduction in the number of switches and HBAs comes a reduction and simplification of the cabling infrastructure. By implementing FCoE and using the same network to carry both IP and Fibre Channel, organizations can reduce the number of network cables.

Using fewer network connections has several advantages, including:

- ▶ Simplifying the connectivity within a data center means that new server racks can be configured and deployed more quickly.
- ▶ Easier and more efficient changes brings with it lower deployment costs because of shorter deployment schedules with less of an impact on the data center’s staff and a reduced chance of any possible interruption of service. The efficiency that is gained allows the data center to respond more quickly to ever-changing business demands.
- ▶ Using fewer cables within a rack results in more efficient airflow, thereby increasing the efficiency of the data center’s cooling infrastructure.
- ▶ Using fewer connections means an increase in the overall reliability of the data center.

Simpler connectivity and reducing the amount of equipment can result in reducing the amount of physical space that is needed for the data center. Less space means even further reductions in the cooling requirements. The reductions in needed space and cooling requirements are especially important to those organizations that lease Co-Location (COLO) space.

As with all consolidated I/O techniques, FCoE increases the demands on network reliability. That is because each network connection is carrying the traffic for several virtual machines; if the connection is lost, the links to several virtual machines are lost. With FCoE, the connection to the storage array is also lost. Standard installation best practices and using top performing connectivity and cabling minimize any potential issues.

Implementing FCoE allows the data center’s managers and administrators to use a common set of tools and software platforms to manage their network, servers, and storage, resulting in further efficiencies. Since FCoE combines the LAN and SAN, each of the data center’s staff members can address a wider array of issues and problems should they arise. This results in quicker response to system outages and meeting of Service Level Agreement requirements.

---

<sup>1</sup> “Doing the Math”, Computer World, April 3, 2006,  
[http://www.computerworld.com/s/article/print/110073/Sidebar\\_Doing\\_the\\_Math](http://www.computerworld.com/s/article/print/110073/Sidebar_Doing_the_Math)

## **Supporting future needs**

Implementing FCoE also provides the data center staff with flexibility regarding migration and expansion plans and provides them with more options. With FCoE, every host or server that is a part of the network can now mount any storage target. This results in higher attach rates and increased efficiencies of the storage arrays.

Virtualization allows enterprises to consolidate applications onto fewer servers in order to lower operating and hardware costs. In addition, it provides the organization's IT administrators the ability to move applications from one server to another to facilitate load balancing. FCoE also provides "storage virtualization," allowing IT administrators to more easily manage their infrastructures and consolidate storage systems from different locations and vendors into one virtual pool of storage.

Another advantage with FCoE is that any network upgrade benefits both the LAN and SAN because they share the same media.

## **Reduces total cost of ownership**

FCoE reduces the capital budgets of upgrades and expansions by requiring fewer network switches, HBAs within the servers, and cabling. Depending on the size of the data center and the extent to which FCoE is deployed, further savings might be achieved due to reduced cooling requirements that allow the use of smaller CRAC units.

Savings in operational expenses can also be realized with FCoE. The cost of operating the equipment is less than a traditional implementation due to the reduction in the amount of hardware (both switches and HBAs), lower power usage, and lower cooling requirements. Reducing the amount of hardware also reduces the cost of annual service contracts. Other operational savings can be found in the area of training and more efficient utilization of data center staff.



# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

## IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 46. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *An Introduction to Fibre Channel over Ethernet, and Fibre Channel over Convergence Enhanced Ethernet*, REDP-4493
- ▶ *IBM Converged Switch B32*, REDP-4588
- ▶ *IBM SAN Survival Guide*, SG24-6143
- ▶ *QLogic 10Gb CNA for IBM System x and IBM Power Systems*, TIPS0720
- ▶ *QLogic 2-port 10Gb Converged Network Adapter (CFFh) for IBM BladeCenter*, TIPS0716
- ▶ *10Gb Ethernet Pass-Thru Module for IBM BladeCenter*, TIPS0715
- ▶ *BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter*, TIPS0708
- ▶ *Cisco Nexus 4001I Switch Module for IBM BladeCenter*, TIPS0754
- ▶ *QLogic Virtual Fabric Extension Module for IBM BladeCenter*, TIPS0717

## Other publications

These publications are also relevant as further information sources:

- ▶ “Doing the Math”, Computer World, April, 3 2006  
[http://www.computerworld.com/s/article/print/110073/Sidebar\\_Doing\\_the\\_Math](http://www.computerworld.com/s/article/print/110073/Sidebar_Doing_the_Math)
- ▶ Priority Flow Control: Build Reliable Layer 2 Infrastructure  
[http://ciscosystems.com/en/US/prod/collateral/switches/ps9441/ps9670/white\\_paper\\_c11-542809.html](http://ciscosystems.com/en/US/prod/collateral/switches/ps9441/ps9670/white_paper_c11-542809.html)
- ▶ FCIA FCoE Recommendations - Minimum Criteria Recommended for Enterprise Data Center IO Unification  
[http://www.fibrechannel.org/documents/doc\\_download/6-fcia-fcoe-guidelines](http://www.fibrechannel.org/documents/doc_download/6-fcia-fcoe-guidelines)
- ▶ FCIA Fibre Channel Solutions Guide  
[http://www.fibrechannel.org/documents/doc\\_download/17-fcia-solution-guide-fall-09](http://www.fibrechannel.org/documents/doc_download/17-fcia-solution-guide-fall-09)

## Online resources

These Web sites are also relevant as further information sources:

- ▶ Product page for IBM System x Converged Network Adapters  
<http://www.ibm.com/systems/storage/product/systemx/cna/>
- ▶ Product page for IBM BladeCenter FCoE products  
<http://www.ibm.com/systems/bladecenter/hardware/openfabric/fcoe.html>
- ▶ Fibre Channel over Ethernet  
<http://www.fcoe.com/>
- ▶ FCoE: the evolutionary approach to SAN expansion  
<http://www.networkworld.com/news/tech/2008/042208-tech-update.html>

## How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this Web site:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Help from IBM

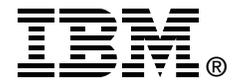
IBM Support and downloads

[ibm.com/support](http://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)





# Planning for Converged Fabrics

## The Next Step in Data Center Evolution



### Best practices when migrating to Fibre Channel Over Ethernet

### Deployment planning for Layer 1 physical infrastructure

### Exploiting products and technology from IBM and Panduit

Over the past decade, almost every industry has witnessed increasing demands for compute capacity and an exponential growth of data, which has resulted in complex and expensive IT infrastructures that are not only cumbersome to deploy, but also to administer and maintain. To counter these increasing demands and complexity in data center infrastructures, as well as to improve utilization, efficiency, and security, there has been greater emphasis on server virtualization and infrastructure consolidation.

The goal of this IBM Redpaper document is to highlight some of the important challenges currently faced in the IT industry with infrastructure growth, and to outline the need for *converged fabrics*. We introduce the emerging IT standards for convergence such as Convergence Enhanced Ethernet (CEE) and Fibre Channel over Ethernet (FCoE). We describe many of the available products and solutions enabling converged fabrics, such as Converged Network Adapters and Converged Fabrics switches that IBM along with business partners such as Panduit are bringing to market.

We continue the discussion by describing the elements of physical deployment planning of convergence technologies, and discuss various scenarios that demonstrate transitioning from traditional to converged infrastructures while following the best practices.

This paper is aimed at data center managers and others who are involved with the planning, implementation, and maintenance of infrastructure in the IT environment.

### INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

### BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:  
[ibm.com/redbooks](http://ibm.com/redbooks)