Connecting an IBM PureFlex System to the Network

IBM Redbooks Solution Guide

Heterogeneous, multi-vendor networks are commonplace and IT managers who deploy a variety of switches in their data centers need to be assured of interoperability. There is always the option of adopting a single vendor strategy, automatically eliminating any interoperability barriers. However, clients choosing a single vendor risk being locked out of new solutions from alternative vendors and unable to exploit cutting-edge technology. IBM® favors a standards-based approach to networking and works closely with its partners, the IEEE, and other leading standards organizations. This approach complements the Open Choice design of the IBM Flex System™ and PureFlex™ System, which gives customers the ability to configure their system with a wide variety of elements from IBM and its partners.

An IBM PureFlex System is shown in Figure 1.

Figure 1. IBM PureFlex System

The Open Data Center Interoperable Network (ODIN) initiative is a work in progress that is designed to help you interpret the open standards that are vital to efficient, interoperable data center networks. IBM System Networking offers a compelling argument based on advanced technology and open standards.

Did you know?

IBM PureSystems™ combines the flexibility of general-purpose computer systems, the elasticity of cloud computing, and the simplicity of an appliance that is tuned to the workload at hand. Expert integrated systems are essentially the building blocks of capability in an enterprise data center. This category of systems represents the collective knowledge of thousands of deployments, established best practice guidelines, innovative thinking, IT leadership, and decades of experience and expertise.

Business value

The offerings in IBM PureSystems are designed to deliver value in the following ways:

- Built-in expertise helps you to address complex business and operational tasks automatically.
- Integration by design helps you to tune systems for optimal performance and efficiency.
- A simplified user experience, from design to purchase to maintenance, creates efficiencies quickly.

The IBM PureSystems offerings are optimized for performance and virtualized for efficiency. These systems offer a no-compromise design with system-level upgradeability. IBM PureSystems are built for the next-generation data center, containing built-in flexibility and simplicity.

Solution overview

IBM expert integrated systems come in two types:

IBM PureFlex System: Pre-configured, pre-integrated

IBM PureFlex System, a part of the IBM PureSystems family of expert integrated systems, combines advanced IBM hardware and software along with patterns of expertise and integrates them into three optimized configurations that are simple to acquire and deploy so you can achieve faster time to value.

If you want a pre-configured, pre-integrated infrastructure with integrated management and cloud capabilities, factory tuned from IBM with x86 and Power hybrid solution, IBM PureFlex System is the answer.

IBM Flex System: Custom built to your requirements

IBM Flex System, the elements that make up IBM PureFlex System, allows you to build your own system to meet your unique IT requirements with a set of no-compromise components including compute, storage, networking, and systems management.

If you want to custom build and tune a configuration to meet your specific requirements utilizing the most advanced blade server in the market with maximum x86 EP compute and memory performance, IBM Flex System is the right fit.

Why choose IBM PureFlex System or IBM Flex System

Extraordinary compute power, storage capability, and networking flexibility, along with the latest processors and a no-compromise design offering:

- Twice the application density in the same floor space
- Energy cost savings of up to 40%
- Reduced software licensing costs with up to 70% fewer cores needed than previous generation systems
- Simplified administration of your whole infrastructure from a single interface
- Management integration of resource pools across compute, storage, and networking reduces management costs up to 50%
- Integrated, automated storage provisioning in 98% less time
- 66% faster set up time

Solution architecture

The IBM Flex System Enterprise Chassis offers intelligent workload deployment and management for maximum business agility. This chassis delivers high-speed performance complete with integrated servers, storage, and networking for multiple chassis management in data center compute environments. Furthermore, its flexible design can meet the needs of varying workloads with independently scalable IT resource pools for higher utilization and lower cost per workload. Increased security and resiliency protect vital information and promote maximum up time, and the integrated, easy-to-use management system reduces setup effort.

The IBM Flex System Enterprise Chassis (machine type 8721) is a 10U next-generation server platform with integrated chassis management. It is a compact, high-density, high-performance, rack-mount, scalable server platform system. It supports up to 14 1-bay compute nodes that can share common resources, such as power, cooling, management, and I/O resources within a single Enterprise Chassis. The chassis can also support up to seven 2-bay compute nodes or three 4-bay compute nodes when the shelves are removed. You can mix and match 1-bay, 2-bay, and 4-bay compute nodes to meet your specific hardware needs.

Figure 2 shows a front view of the IBM Flex System Enterprise Chassis.

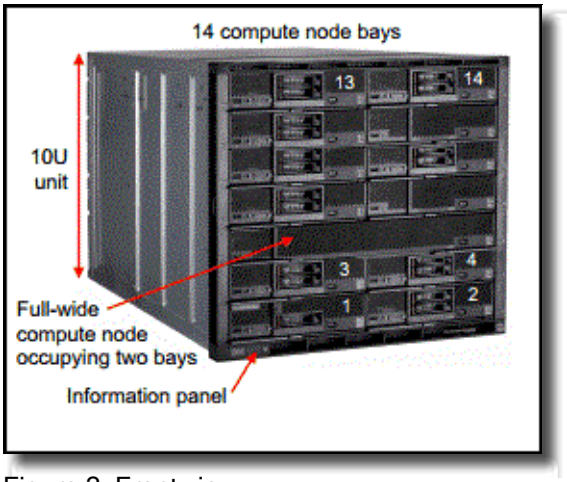

Figure 2. Front view

Figure 3 shows a rear view of the IBM Flex System Enterprise Chassis.

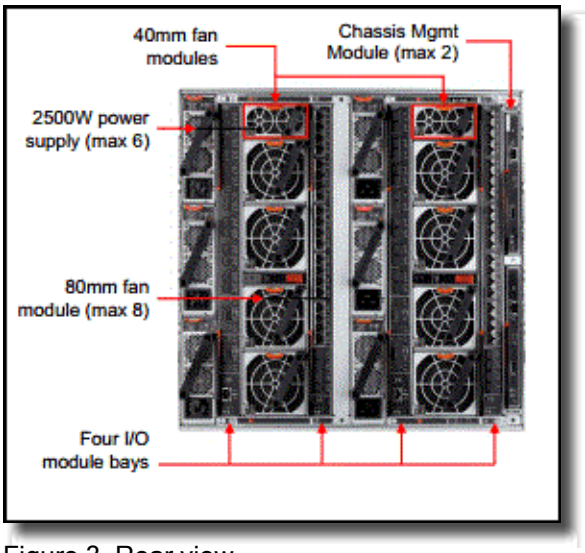

Figure 3. Rear view

Selected Ethernet and Converged I/OM module products

The IBM Flex System Enterprise Chassis features a number of Ethernet and Converged I/O module solutions providing a combination of 1 Gb and 10 Gb ports to the servers and 1 Gb, 10 Gb, and 40 Gb for connectivity to additional chassis or to the outside upstream infrastructure. The IBM Flex System Enterprise Chassis ensures that a suitable selection is available to meet the needs of the server nodes.

There are four Ethernet I/O modules and a Converged I/O module available for deployment with the Enterprise Chassis:

- IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch
- IBM Flex System EN4093R 10Gb Ethernet Scalable switch
- IBM Flex System EN4091 1/10Gb Ethernet pass-thru
- IBM Flex System EN2092 1Gb Ethernet Scalable switch

IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch

The IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch, which has up to 12 external programmable IBM Omni Ports with the choice of SFP+ based 10Gb Ethernet connectivity or 4/8 Gb Fibre Channel connectivity, offers Layer 2/3 switching as well as FCoE Full Fabric and Fibre Channel NPV Gateway operations designed to fit within the I/O module bays of the Enterprise Chassis. The switch has:

- Up to 42 internal 10 Gb ports
- Up to 2 external 10 Gb uplink ports (SFP+ connectors or 1 Gb with SFP transceivers)
- Up to 2 external 40 Gb uplink ports (QSFP+ connectors, or four 10 Gb connections per port)
- Up to 12 external Omni Ports (SFP + connectors)

The switches are particularly suited for clients who:

- Want to implement a converged infrastructure with iSCSI, NAS, or FCoE. The CN4093 includes an integrated Fibre Channel Forwarder (FCF). This allows it to run in full-fabric FC mode or NPV mode. The CN4093 can break out FC traffic within the chassis for native Fibre Channel SAN connectivity.
- Will be building a 10 Gb infrastructure.
- Are implementing a virtualized environment.
- Require investment protection for 40 Gb uplinks.
- Want to reduce TCO, improve performance, and maintain high levels of availability and security.
- Want to avoid or minimize oversubscription (that is, traffic from multiple internal ports attempting to pass through a lower quantity of external ports, leading to congestion and performance impact).

For detailed product and option ordering information, refer to *IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch*, TIPS0910:

http://www.redbooks.ibm.com/abstracts/tips0910.html?Open

IBM Flex System Fabric EN4093R 10Gb Scalable Switch

The IBM Flex System Fabric EN4093R 10Gb Scalable Switch is a 10 Gb 64-port upgradable midrange to high-end switch module, offering Layer 2/3 switching and designed to fit within the I/O module bays of the Enterprise Chassis. The switch has:

- Up to 42 internal 10 Gb ports
- Up to 14 external 10 Gb uplink ports (SFP+ connectors or 1Gb with SFP transceivers)
- Up to 2 external 40 Gb uplink ports (QSFP+ connectors or four 10Gb connections per port)

The switch is particularly suited for clients who:

- Will be building a 10 Gb infrastructure.
- Want to implement a converged infrastructure with iSCSI, NAS, or FCoE. The EN4093R acts as a transit switch for FCoE, passing traffic upstream to a FCF/Gateway, which can be located at the top-of-the-rack, where FC traffic is broken out. Currently, Brocade and Cisco FCF-capable products are supported at the top-of-the-rack.
- Will be implementing a virtualized environment.
- Require investment protection for 40 Gb uplinks.
- Want to reduce TCO, improve performance, and maintain high levels of availability and security.
- Want to avoid oversubscription (that is, traffic from multiple internal ports attempting to pass through a lower quantity of external ports, leading to congestion and performance impact).

For detailed product and option ordering information, refer to *IBM Flex System Fabric EN4093R 10Gb Scalable Switch*, TIPS0864:

http://www.redbooks.ibm.com/abstracts/tips0864.html?Open

IBM Flex System EN4091 10Gb Ethernet Pass-thru

The EN4091 10Gb Ethernet Pass-thru module offers a one for one connection between a single node bay and an I/O module uplink. Unlike the switching modules, the EN4091 does not have the capacity to provide multiple connections to a compute node. It has no management interface. The EN4091 can

support both 1 Gb and 10 Gb dual-port adapters installed in the compute nodes. If quad-port adapters are installed in the compute nodes, only the first two ports have access to the pass-through module's ports; the remaining two ports cannot be used.

The necessary 1 GbE or 10 GbE module (SFP, SFP+ or DAC) must also be installed in the external ports of the pass-through. This configuration must match the speed (1 Gb or 10 Gb) of the adapter ports on the compute nodes.

For detailed product and option ordering information, refer to *IBM Flex System EN4091 10Gb Ethernet Pass-thru Module*, TIPS0865:

http://www.redbooks.ibm.com/abstracts/tips0865.html?Open

IBM Flex System EN2092 1Gb Ethernet Scalable Switch

The EN2092 1Gb Ethernet Switch provides support for L2/L3 switching and routing. The switch has:

- Up to 28 internal 1 Gb ports
- Up to 20 external 1 Gb ports (RJ45 connectors)
- Up to 4 external 10 Gb uplink ports (SFP+ connectors)

This switch is ideal for clients who:

- Are using 1 Gb as their networking infrastructure
- Plan to deploy virtualization and require multiple 1 Gb ports
- Want investment protection for 10 Gb uplinks
- Want to reduce TCO, improve performance, and maintain high levels of availability and security
- Want to avoid oversubscription (multiple internal ports attempting to pass through a lower quantity of external ports, leading to congestion or performance impact).

Which ports are enabled and available depends on the features activated on the I/O module.

For detailed product and option ordering information, refer to *IBM Flex System EN2092 1Gb Ethernet Scalable Switch*, TIPS0861:
http://www.redbooks.ibm.com/abstracts/tips0861.html?Open

For a complete description of the IBM PureFlex and Flex Systems technology refer to *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984:
http://www.redbooks.ibm.com/abstracts/sg247984.html?Open

Usage scenarios

In traditional data centers, which are prevalent today, network infrastructure is optimized for traffic between end user workstations and application servers (north-south traffic). These datacenters are typically not optimized for high-bandwidth server-to-server (east-west) traffic.

Data center requirements are changing. Because of increased use of virtualization, VM mobility and composite (multi-tier) applications, server-to-server traffic is increasing significantly. It is estimated that over 80% of traffic in datacenters soon will be between servers.

VM mobility causes server-to-server data transfers, which require significant amounts of bandwidth for prolonged periods of time. This can negatively affect performance of other network traffic flows, and can consequently hinder application performance.

Another important reason for increased server-to-server traffic demand is the increased use of composite or multi-tier applications. These applications combine multiple functions that exist on different servers and that need to communicate and exchange data quickly and efficiently. Examples of such applications are those that are built with IBM WebShere®, which can include web, application, and database servers.

Next generation data center networks must be designed with these requirements in mind. They must be able to provide for meshed topologies and flatter networks. There is also a need for network traffic management and prioritization, and application management.

The following three scenarios of PureFlex network connectivity are considered here:

- IBM Flex System EN4093R switch modules connected to IBM System Networking RackSwitch™ G8264 switches (which are connected to a Cisco or Juniper infrastructure)
- IBM Flex System EN4093R switch modules connected to a Cisco or Juniper infrastructure
- IBM Flex System EN4091 pass-thru modules connected to network infrastructure

The scenarios are based on IBM EN4093R switch modules and EN4091 pass-thru modules. The new CN4093 switch modules support the same implementation scenarios as EN4093R, and offer the following additional functionality:

- FC external connectivity to Brocade and Cisco SAN fabrics.
- FCF (Fibre Channel Forwarder) functionality implemented in the CN4093, allowing this function to reside within the chassis.
- Omni Ports for ultimate flexibility. These ports are able to operate as either 10 Gb Ethernet or 4/8 Gb Fibre Channel ports, depending on the inserted transceiver.

IBM Flex System EN4093R switch modules connected to IBM System Networking RackSwitch G 8264 switches

This scenario uses EN4093R switch modules in the IBM Flex System chassis, connected to IBM G8264 ToR switches. This topology is appropriate for large deployments with multiple IBM Flex System chassis. EN4093R switch modules handle east-west traffic flows between compute nodes in the same chassis. IBM G8264 switches handle east-west traffic flows between compute nodes in different Flex System chassis, so these traffic flows do not need to go through an upstream Cisco or Juniper network infrastructure. This is the suggested topology for next-generation data center implementations.

Figure 4 shows this network topology. The diagram shows that:

- vLAG is implemented on the pair of EN4093R switch modules inside IBM Flex System chassis.
- vLAG is implemented on the pair of RackSwitch G8264 switches.
- The pair of core switches is configured in a vendor-specific virtual group.

Figure 4 shows only one Flex System chassis, but multiple chassis can be connected to the same pair of IBM G8264 switches in this manner.

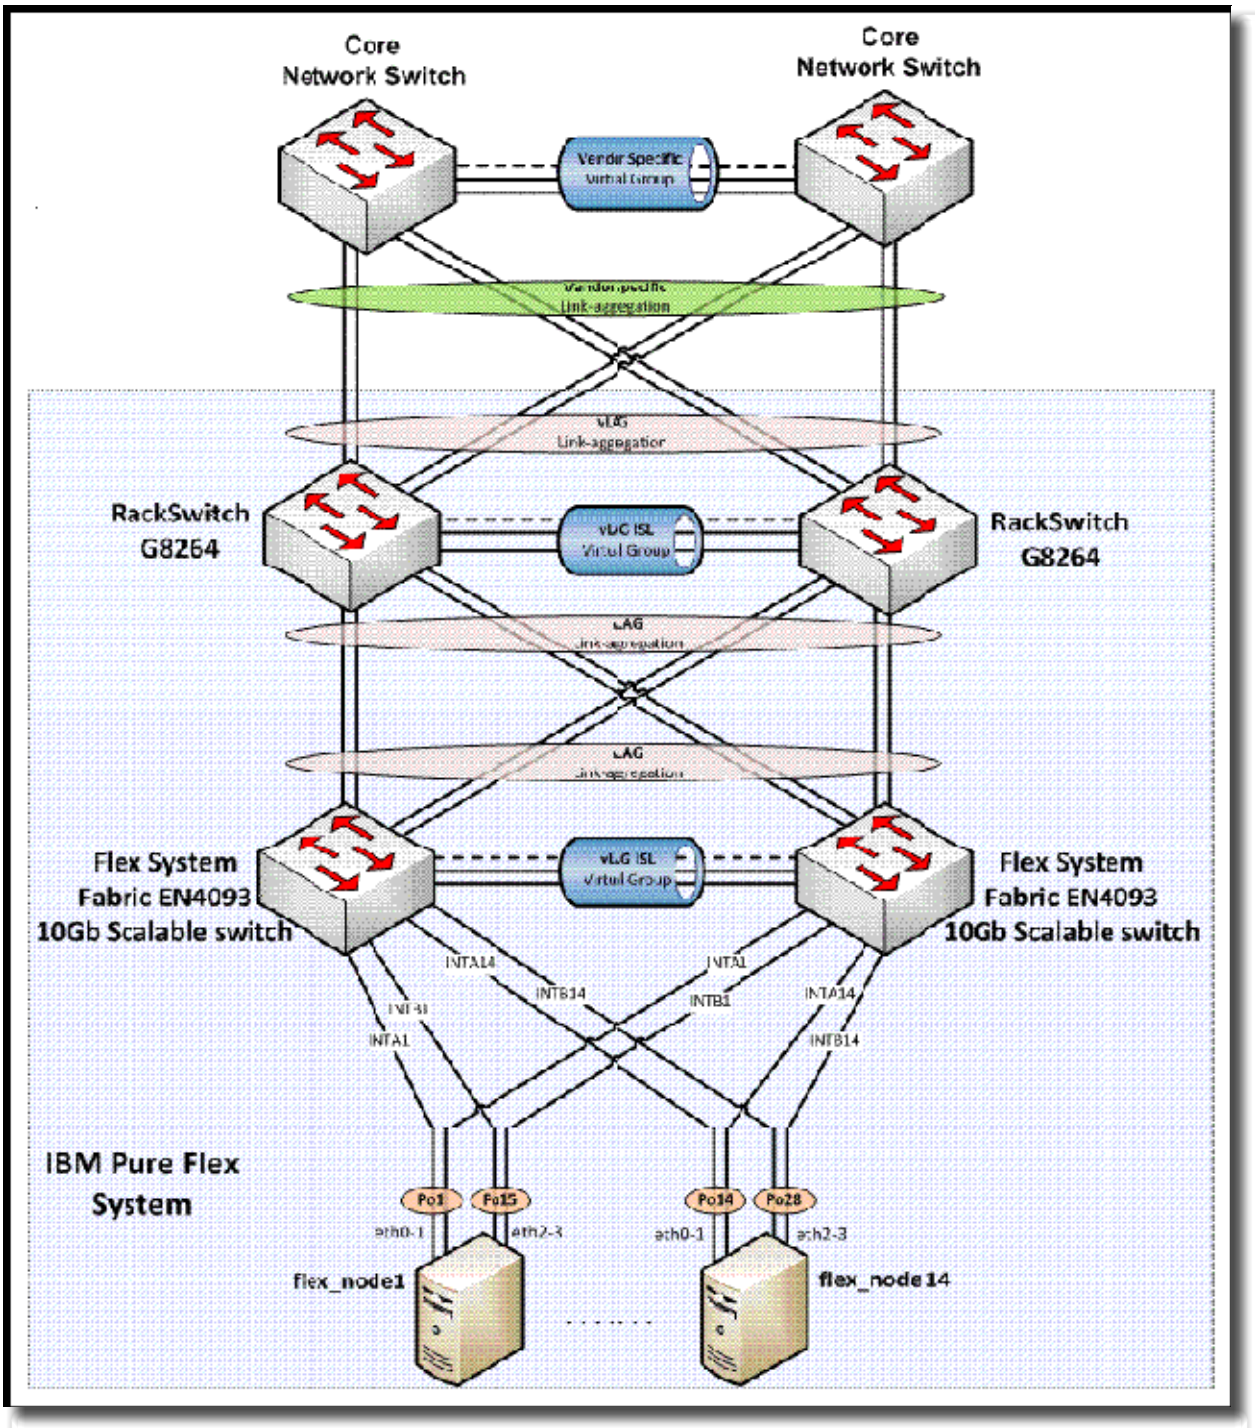

Figure 4. EN4093R connected to G8264

IBM Flex System EN4093 connected to Cisco or Juniper infrastructure

In this scenario, EN4093 switch modules are directly connected to upstream Cisco or Juniper infrastructure. This topology works well in single chassis deployments, where east-west traffic flows are limited to compute nodes within the single Flex System chassis. This topology is not suitable for large environments with multiple Flex System chassis because east-west traffic between compute nodes in different chassis would have to cross the core network switches in the upstream infrastructure.

This network topology is illustrated in Figure 5.

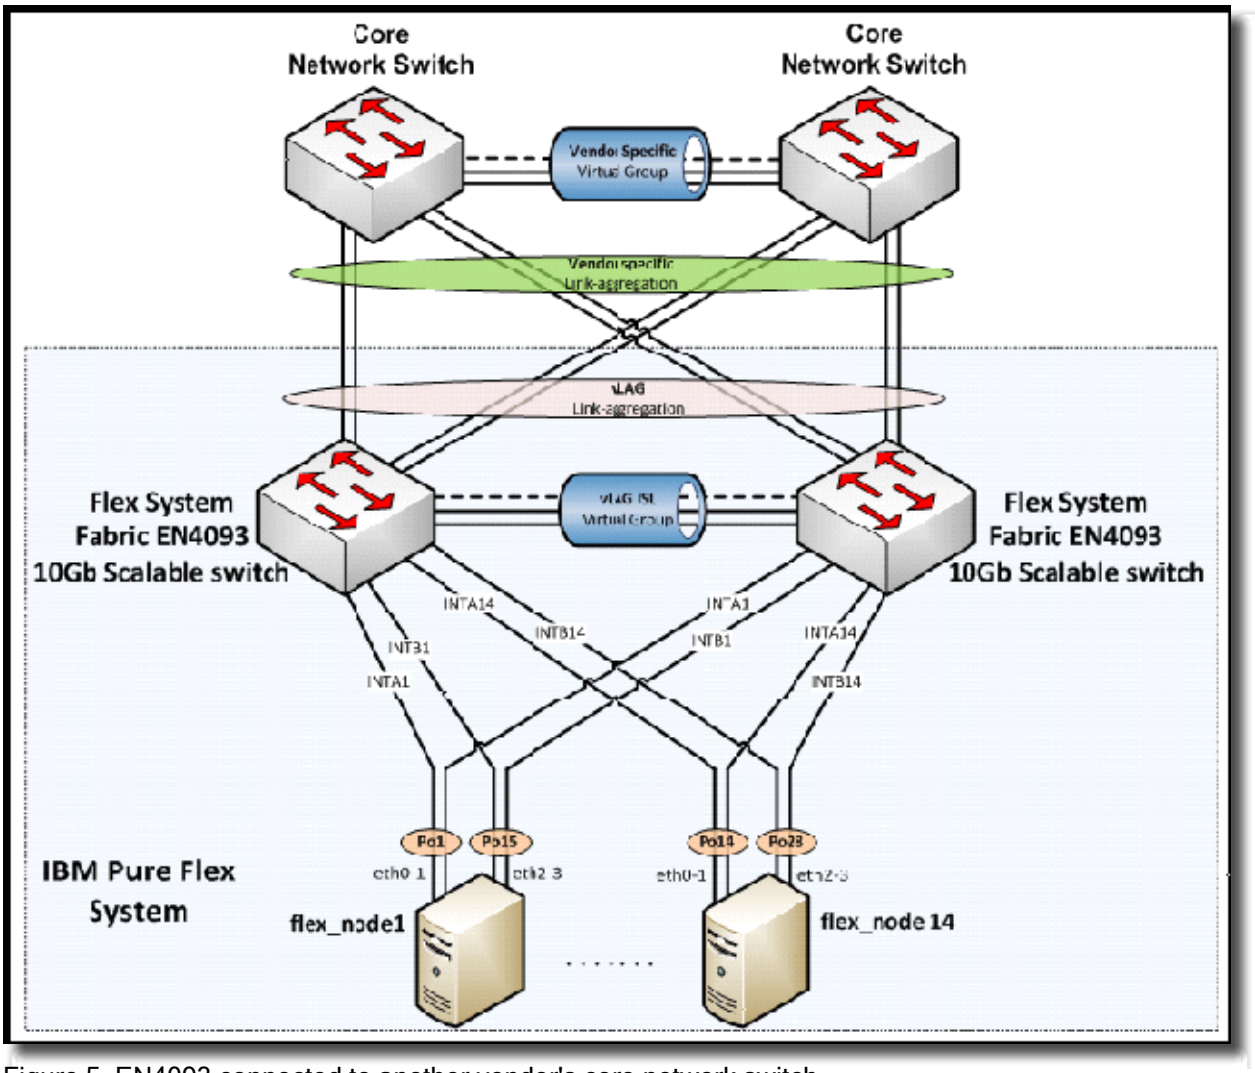

Figure 5. EN4093 connected to another vendor's core network switch

IBM Flex System EN4091 pass-thru modules

Figure 6 shows an example of a network topology utilizing the EN4091 pass-thru modules. EN4091 pass-thru modules in the Flex System chassis only pass the traffic between compute nodes and upstream network switches. They do not provide any switching functionality, and all east-west traffic between compute nodes in the chassis must be handled by upstream network switches. Because of this, pass-through solutions are not a preferred topology in next generation data centers, especially those built with multiple IBM Flex System chassis that will require very high server-to-server bandwidth to support VM mobility and composite applications.

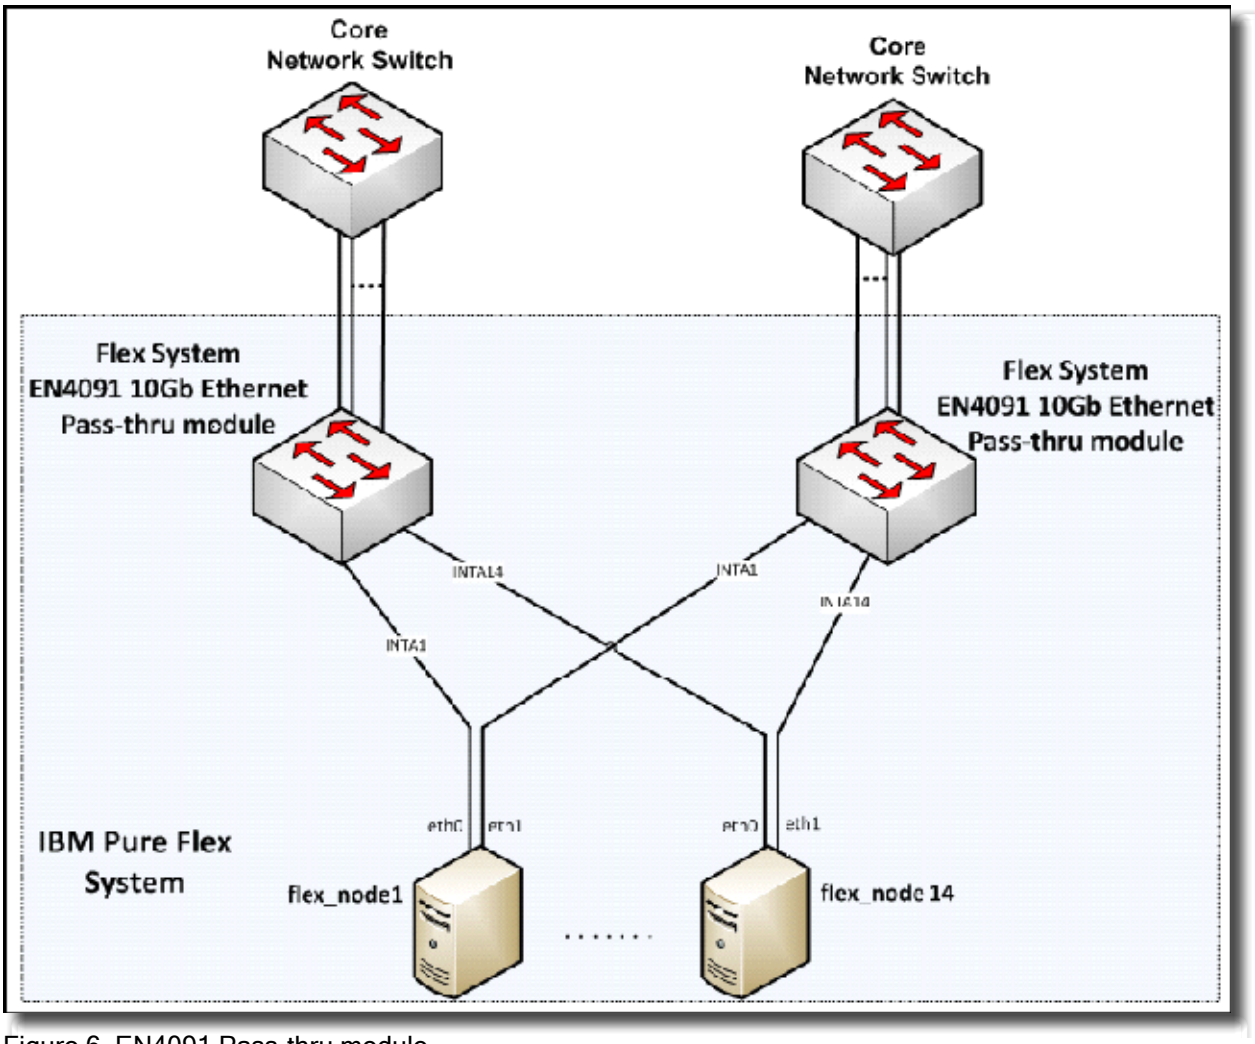

Figure 6. EN4091 Pass-thru module

Integration and interoperability

The Tolly Group (founded by Kevin Tolly in 1989) is an independent test lab and provider of third-party validation services. The Tolly Group has worked with virtually every major vendor, evaluating some of the most important products and technologies to appear over the past decade. The Tolly Group is positioned to certify vendor solutions and provide evidence that their products meet or exceed marketing claims. This proof-of-performance and features or functions lets customers know they can buy with confidence.

A recent Tolly report demonstrated that the PureFlex System and Flex System can seamlessly interoperate between mainstream vendors, such as Cisco and Juniper, using a wide range of networking protocols.

The IBM Flex System EN4093R 10Gb and Flex System EN2092 1Gb Ethernet Scalable Switch, together with the RackSwitches G8052 and G8264, demonstrated interoperability with the Cisco Nexus 5548UP, Nexus 5596UP, Nexus 7000 and Catalyst 6509-E. The switches, features, and versions tested by the Tolly group are shown in Figure 7.

Switch Under Test		Interfaces	Version
IBM	EN4093 Ethernet Scalable Switch	10 / 40 Gigabit Ethernet	7.2 FCoE tests used version 7.5
	EN2092 Ethernet Scalable Switch	1 / 10 Gigabit Ethernet	
	RackSwitch G8052	1 / 10 Gigabit Ethernet	
	RackSwitch G8264	10 / 40 Gigabit Ethernet	
Cisco Systems	Nexus 5548UP	1 / 10 Gigabit Ethernet	NX-OS 5.1(3)N1(1a)
	Nexus 5596P	1 / 10 Gigabit Ethernet	NX-OS 5.1(3)N1(1a)
	Nexus C7009	1 / 10 Gigabit Ethernet	NX-OS 6.0(1)
	Catalyst 6509-E	1 / 10 Gigabit Ethernet	IOS 12.2(18)SXF8

Figure 7. Switches, features, and versions under test

Tolly engineers evaluated Layer 2 features typically used in the data center, including Auto-negotiation, Jumbo frame transmission, 802.1pQ VLAN tag propagation, 10 Gb and 40 GbE LAN PHY support, LACP support, LLDP, STP (including RSTP, PVRST, Cisco PVRST+ and MSTP) and vLAG to Cisco vPC interoperability. Layer 3 protocols were evaluated:RIPv2, BGPv4, OSPFv2, OSPFv3 and PIM-sparse mode. Successful FIP snooping and DCBX negotiation ensured lossless FCoE transmission using Priority Flow Control (PFC). Effective bandwidth allocation among traffic classes using Enhanced Transmission Selection (ETS) was also demonstrated.

Figure 8 shows the certified interoperable features and functions.

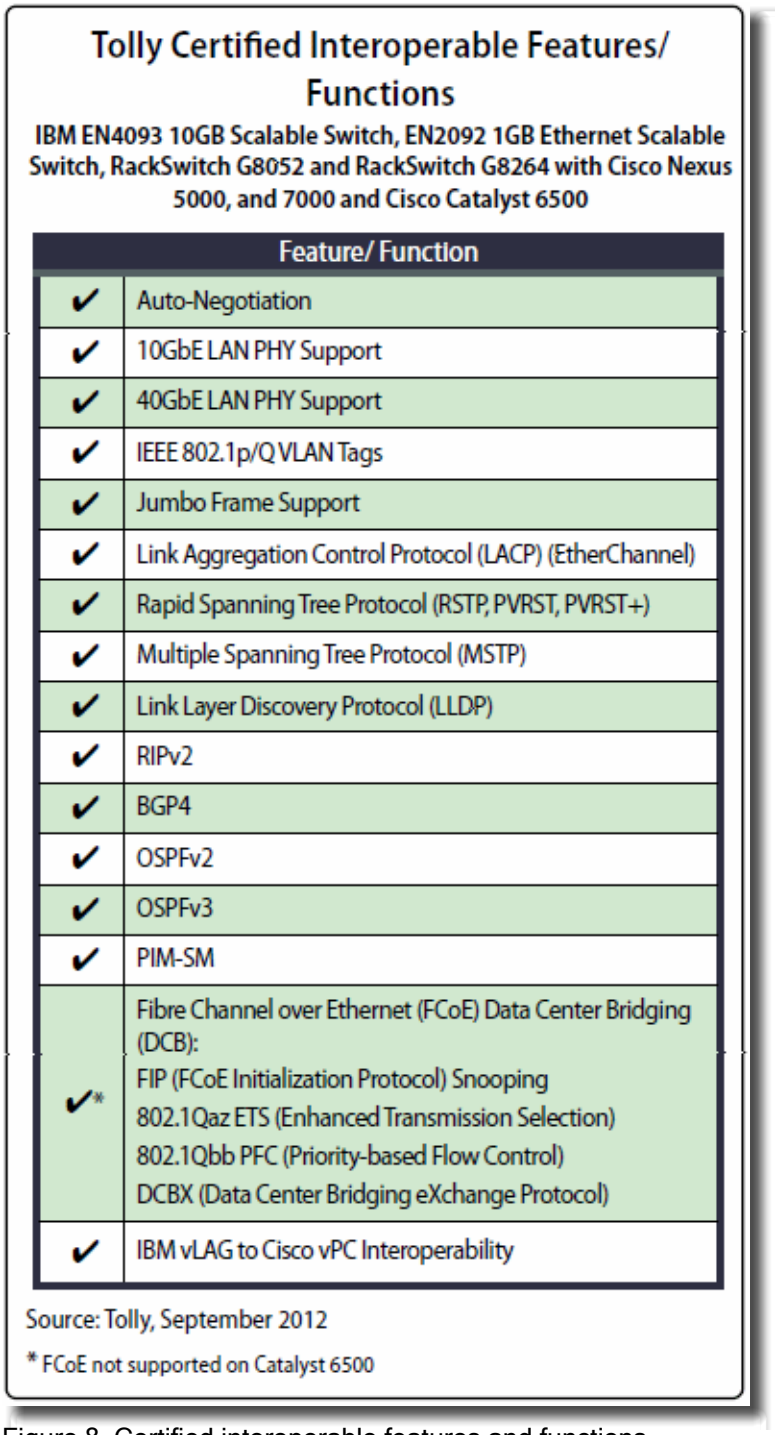

Tolly Certified Interoperable Features/ Functions

IBM EN4093 10GB Scalable Switch, EN2092 1GB Ethernet Scalable Switch, RackSwitch G8052 and RackSwitch G8264 with Cisco Nexus 5000, and 7000 and Cisco Catalyst 6500

	Feature/ Function
✔	Auto-Negotiation
✔	10GbE LAN PHY Support
✔	40GbE LAN PHY Support
✔	IEEE 802.1p/Q VLAN Tags
✔	Jumbo Frame Support
✔	Link Aggregation Control Protocol (LACP) (EtherChannel)
✔	Rapid Spanning Tree Protocol (RSTP, PVRST, PVRST+)
✔	Multiple Spanning Tree Protocol (MSTP)
✔	Link Layer Discovery Protocol (LLDP)
✔	RIPv2
✔	BGP4
✔	OSPFv2
✔	OSPFv3
✔	PIM-SM
✔*	Fibre Channel over Ethernet (FCoE) Data Center Bridging (DCB): FIP (FCoE Initialization Protocol) Snooping 802.1Qaz ETS (Enhanced Transmission Selection) 802.1Qbb PFC (Priority-based Flow Control) DCBX (Data Center Bridging eXchange Protocol)
✔	IBM vLAG to Cisco vPC Interoperability

Source: Tolly, September 2012

* FCoE not supported on Catalyst 6500

Figure 8. Certified interoperable features and functions

Read the full report to understand how IBM Flex System Fabric EN4093 10Gb, Flex System EN2092 1Gb Ethernet Scalable Switch and IBM RackSwitch G8052 and G8264 demonstrate interoperability with Cisco Nexus 5548UP, 5596UP, C7009, and Catalyst 6509-E switches (starts PDF download):

https://ibm.biz/Bdxqaj

Open data center with an interoperable network

Over the past several years, progressive data centers have undergone fundamental and profound architectural changes. Nowhere is this more apparent than in the data center network infrastructure.

An open data center with an interoperable network (ODIN) has a flat, converged, virtualized data center network that is based on open industry standards, as shown in Figure 9. Instead of under-utilized devices, multi-tier networks, and complex management environments, the modern data center is characterized by highly utilized servers running multiple VMs; flattened, lower latency, highly utilized networks; and automated, integrated management tools. New software-defined network approaches (including overlay networks and OpenFlow standards) greatly simplify the implementation of features such as dynamic workload provisioning, load balancing, and redundant paths for high availability and network reconfiguration. Further, high-bandwidth links between virtualized data center resources can extend across multiple data center locations to provide business continuity and backup/recovery of mission-critical data.

The ODIN reference architecture is a work in progress that is designed to address best practices and help you interpret the networking standards that are vital to efficient data center operations. The information about these methods and standards help you make the transition from discrete, special-purpose networks that each have their own management tools, to a converged, flattened network that uses a common set of management tools.

This proven approach has been implemented by IBM within its own data centers using existing products, and also by industry-leading clients worldwide. IBM Flex Systems and the PureFlex System families are key players on this moving strategy as shown in Figure 9.

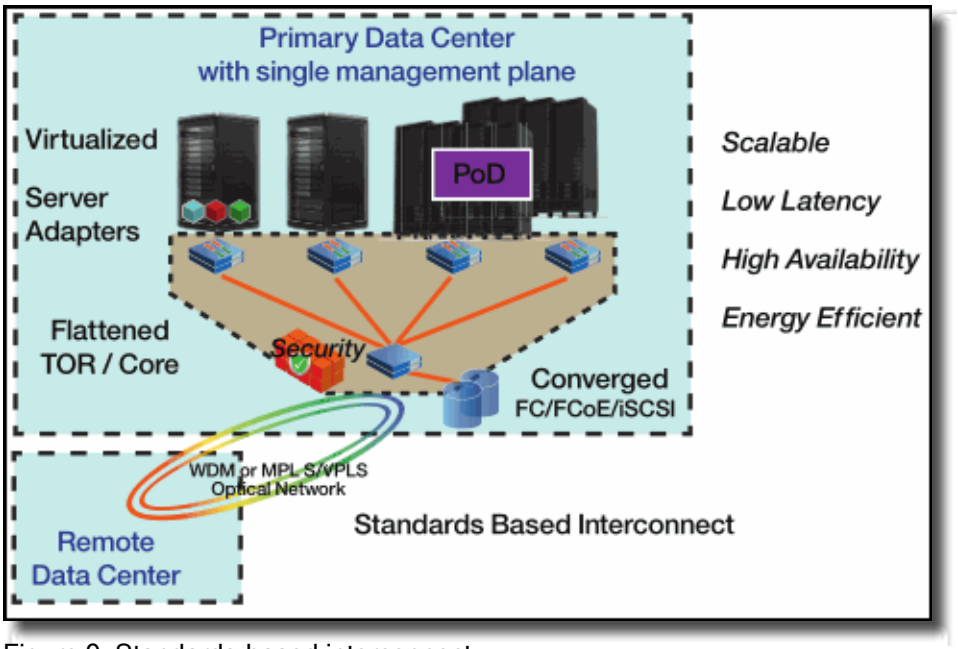

Figure 9. Standards-based interconnect

IBM believes that the practical, cost-effective evolution of data center networks should be based on open industry standards. This can be a challenging and often confusing proposition because there are so many different emerging network architectures, both standard and proprietary. The ODIN materials created by IBM currently address issues such as virtualization and Virtual Machine (VM) migration across flat Layer 2 networks, lossless Ethernet, Software-Defined Networking (SDN) including OpenFlow, and extended distance WAN connectivity (including low latency).

SDN is used to simplify network control and management, automate network virtualization services, and provide a platform from which to build agile network services. By abstracting the control and management aspects of a network into a logical software program, SDN allows real-time programmability and manageability of networks, and it can leverage a centralized logical network view that is easily manipulated via software to implement complex networking rules. This allows networks to achieve unprecedented levels of scalability and flexibility, as well as dynamic behaviors matching cloud service-oriented dynamics.

To learn more about SDN, see:
http://www.ibm.com/systems/networking/solutions/sdn.html

The benefits of ODIN include:

- Customer choice and future-proof designs for data center networks, including vendor-neutral RFQs
- Lower total cost of ownership by enabling a multi-vendor network
- Avoiding confusion in the marketplace between proprietary and vendor-neutral solutions
- Providing best practices, relative maturity, and interpretation of networking standards

By taking advantage of the ODIN reference architecture, you can design a cost-effective and manageable data center that fully utilizes the potential of virtualization. It also gives you the flexibility to migrate to federated data centers, in which computing, storage, and network resources can be treated as dynamically provisioned resource pools that can be rapidly partitioned into any desired configuration.

Over time, emerging standards will help to improve dynamic resource provisioning, offer management flexibility, and deploy new features more rapidly. Similar technologies in software-defined networking will provide for an easily reconfigurable logical network that is aware of the requirements for workload mobility.

To learn more about ODIN, see:

http://www.ibm.com/systems/networking/solutions/odin.html

Conclusion

IBM System Networking's Ethernet product deployments support more than 15 million Ethernet switch ports installed in enterprise and public sector data centers worldwide. Clients recognize the benefits of system networking solutions that enable the combination of best-of-breed components pre-integrated in a rack or integrated system, yet with an open approach that allows these systems to easily connect into existing core network infrastructures, without the need to "rip and replace" or lose investments in existing equipment. IBM System Networking's Ethernet switch solutions have been adopted by more than 350 of the Fortune 500 customers around the world.

IBM System Networking's Ethernet products have been tested in independent industry labs operated by The Tolly Group, which has verified that these products deliver seamless interoperability between mainstream vendors such as Cisco and Juniper. They also found that System Networking products demonstrate significantly lower latency, higher energy efficiency, and better price/performance than comparable products currently in the marketplace.

IBM Smarter Computing solutions equipped with System Networking products bring intelligence, flexibility, and speed to the essential access, distribution, and aggregation layers where server and storage systems are connected to the data center network. Running the robust, reliable, and mature IBM Networking OS, RackSwitch top-of-rack Ethernet switches and embedded switches for Flex System and BladeCenter® can help deliver high-speed, low-latency performance for IBM servers and storage. Clients who deploy System Networking solutions can realize improved economics, better network performance, less complexity, greater energy efficiency, and streamlined management.

Ordering information

For the complete IBM PureFlex and Flex System story, for a quote, or to call IBM, refer to:
http://www.ibm.com/systems/pureflex/index.html

Related information

For more information, see the following documents:

- IBM Offering Information page (to search on announcement letters, sales manuals, or both)
  http://www.ibm.com/common/ssi/index.wss?request_locale=en
  On this page, enter `PureFlex System` or `Flex System` , select the information type, and then click **Search**. On the next page, narrow your search results by geography and language.
- BM Flex System and RackSwitch Networking Solutions flyer
  https://ibm.biz/Bdxqge
- Open Datacenter Interoperable Networks (ODIN)
  http://www.ibm.com/systems/networking/solutions/odin.html
- Software Defined Networking (SDN)
  http://www.ibm.com/systems/networking/solutions/sdn.html
- *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984
  http://www.redbooks.ibm.com/abstracts/sg247984.html?Open
- IBM PureFlex System and IBM Flex System overview
  http://www.ibm.com/systems/pureflex/index.html
- *Implementing IBM System Networking 10Gb Ethernet Switches*, SG247960
  http://www.redbooks.ibm.com/abstracts/sg247960.html?Open
- *IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch*, TIPS0910
  http://www.redbooks.ibm.com/abstracts/tips0910.html?Open
- *IBM Flex System Fabric EN4093R 10Gb Scalable Switch*, TIPS0864
  http://www.redbooks.ibm.com/abstracts/tips0864.html?Open
- *IBM Flex System EN4091 10Gb Ethernet Pass-thru Module*, TIPS0865
  http://www.redbooks.ibm.com/abstracts/tips0865.html?Open
- *IBM Flex System EN2092 1Gb Ethernet Scalable Switch*, TIPS0861
  http://www.redbooks.ibm.com/abstracts/tips0861.html?Open
- *IBM System Networking RackSwitch G8264/G8264T*, TIPS0815
  http://www.redbooks.ibm.com/abstracts/tips0815.html?Open
- *Implementing IBM System Networking 10Gb Ethernet Switches*, SG24-7960
  http://www.redbooks.ibm.com/abstracts/sg247960.html?Open

- IBM System Networking Redbooks
  http://www.redbooks.ibm.com/portals/networking
- IBM PureFlex System and IBM Flex System overview
  http://www.ibm.com/systems/pureflex/index.html

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.

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

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.

This document was created or updated on January 30, 2013.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
  ibm.com/redbooks
- Send your comments in an e-mail to:
  redbook@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at http://www.ibm.com/redbooks/abstracts/tips0941.html .

Trademarks