


IBM Flex System and PureFlex System Network Implementation with Cisco Systems

Connect IBM PureFlex Systems to a Cisco Network

Troubleshoot and maintain the IBM EN4093

Understand Cisco IOS and IBM N/OS differences



Jon Tate
Jure Arzensek
David Cain
William King

Gaston Sancassano Rodriguez
Tiago Nunes dos Santos

Redbooks



International Technical Support Organization

**IBM Flex System and PureFlex System Network
Implementation with Cisco Systems**

August 2013

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

First Edition (August 2013)

This edition applies to the IBM PureFlex System and Cisco Nexus 5000 software and hardware available in September 2012. This may, or may not, include pre-GA code.

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

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

Contents

Notices	vii
Trademarks	viii
Preface	ix
Authors	ix
Now you can become a published author, too!	xi
Comments welcome	xi
Stay connected to IBM Redbooks	xii
Chapter 1. Introduction	1
1.1 Networking requirements	2
1.2 Data center architecture	3
1.2.1 The IBM PureFlex System and IBM Flex System family	3
1.3 The goal of this book	4
1.4 Networking equipment	4
1.4.1 IBM System Networking	4
1.4.2 Cisco Switches	6
Chapter 2. Layer 1 Overview	7
2.1 Layer 1 networking concepts and terminology	8
2.1.1 Ethernet cabling	8
2.1.2 Twisted-pair copper cabling	8
2.1.3 Fiber optic cabling	10
2.1.4 Physical configuration parameters	15
2.2 Physical layer on IBM Flex System Enterprise Chassis	16
2.3 IBM Flex System Ethernet I/O modules	19
2.3.1 IBM Flex System EN2092 1Gb Ethernet Scalable Switch	19
2.3.2 IBM Flex System Fabric EN4093/EN4093R 10 Gb Scalable Switch	22
2.3.3 IBM Flex System EN4091 10Gb Ethernet Pass-thru module	26
2.3.4 Cables and transceivers for I/O modules	27
2.4 IBM Flex System Ethernet adapters	28
2.4.1 IBM Flex System CN4054 10Gb Virtual Fabric Adapter	28
2.4.2 IBM Flex System EN2024 4-port 1Gb Ethernet Adapter	30
2.4.3 IBM Flex System EN4132 2-port 10Gb Ethernet Adapter	32
Chapter 3. Layer 2 Overview	35
3.1 Basic Frame Forwarding Concept	36
3.2 Virtual local area network (VLAN) and tagging	37
3.2.1 Tagged frames	38
3.3 Spanning tree	38
3.3.1 Spanning Tree Protocol (STP) IEEE802.1D	38
3.3.2 Rapid Spanning Tree (RSTP) IEEE802.1w	39
3.3.3 Multiple Spanning Tree (MSTP) IEEE802.1s	39
3.3.4 Per VLAN Rapid Spanning Tree (PVRST)	40
3.4 Dynamic Link Aggregation Control Protocol (LACP)	40
3.5 Virtual Link Aggregation Groups (VLAG)	41
3.6 Cisco Virtual Port Channel (vPC)	41
3.7 Link Layer Discovery Protocol (LLDP)	42
3.8 Layer 2 Failover	42

Chapter 4. Layer 3 Overview	45
4.1 Overview of Layer 3	46
4.2 Static routes	46
4.3 Default gateways	46
4.4 Equal-cost multi-path (ECMP) static routes	46
4.5 Routing Information Protocol v2 (RIPv2)	46
4.6 Enhanced Interior Gateway Routing Protocol (EIGRP)	47
4.7 Open Shortest Path First (OSPF) for IPv4	47
4.7.1 OSPF area types	47
4.7.2 Neighbors and adjacencies	48
4.7.3 Link State Database (LSDB)	48
4.7.4 OSPF router types	49
4.7.5 Shortest path first	49
4.8 Border Gateway Protocol (BGP)	50
4.9 IPv6	51
4.9.1 Address size	51
4.9.2 Address usage	51
4.9.3 Address hierarchy	51
4.9.4 Address autoconfiguration/plug-and-play	51
4.10 Open Shortest Path First for IPv6 (OSPFv3)	52
4.11 FHRP (First Hop Redundancy Protocols) VRRP and HSRP	52
4.11.1 Active-active redundancy	54
4.11.2 VRRP high availability with VLAGs	54
Chapter 5. Interoperability Use Cases: Connecting to a Cisco Network	55
5.1 Introduction	56
5.2 High availability overview	56
5.2.1 Looped and blocking design	57
5.2.2 Non-looped, single upstream device design	58
5.2.3 Non-looped, multiple upstream devices design	59
5.3 Fully redundant with virtualized chassis technology (VSS/vPC/vLAG)	60
5.3.1 Components used	60
5.3.2 Network topology and physical setup	61
5.3.3 EN4093flex_1 configuration	62
5.3.4 G8264tor_1 configuration	67
5.3.5 Nexus5548core_1 vPC primary switch configuration	69
5.3.6 Cisco Nexus 5548core_2 vPC secondary switch configuration	71
5.3.7 Verification and show command output	73
5.3.8 Full configuration files	91
5.4 Fully redundant with traditional spanning-tree	122
5.4.1 Topology and requirements	122
5.4.2 Components used	123
5.4.3 Network diagram and physical setup	123
5.4.4 EN4093flex_1 configuration	124
5.4.5 G8264tor_1 configuration	128
5.4.6 Nexus5548core_1 STP primary switch configuration	132
5.4.7 Nexus5548core_2 STP secondary switch configuration	133
5.4.8 Verification and show command output	135
5.4.9 Full configuration files	158
5.5 Fully redundant with Open Shortest Path First (OSPF)	187
5.5.1 Topology and requirements	188
5.5.2 Network diagram and physical setup	189
5.5.3 EN4093flex_1 configuration	190

5.5.4	G8264tor_1 configuration	194
5.5.5	G8264tor_2 configuration	199
5.5.6	Nexus5548core_1 switch configuration	203
5.5.7	Nexus5548core_2 configuration	205
5.5.8	Verification and show command output	208
5.5.9	Full configuration files	233
Chapter 6.	Troubleshooting and maintenance	265
6.1	Troubleshooting	266
6.1.1	Basic troubleshooting procedures	266
6.1.2	Connectivity troubleshooting	270
6.1.3	Port mirroring	271
6.1.4	Serial cable troubleshooting procedures	273
6.2	Configuration management	274
6.2.1	Configuration files	274
6.2.2	Configuration blocks	274
6.2.3	Managing configuration files	274
6.2.4	Resetting to factory defaults	276
6.2.5	Password recovery	282
6.3	Firmware management	282
6.3.1	Firmware images	282
6.3.2	Upgrading the firmware with ISCLI	283
6.3.3	Recovering from a failed firmware upgrade	287
6.4	Logging and reporting	290
6.4.1	System logs	290
6.4.2	SNMP	292
6.4.3	Remote Monitoring (RMON)	296
6.4.4	Using sFlow to monitor traffic	298
Appendix A.	Cisco IOS to IBM isCLI Command Comparison	301
	General configuration	302
	Authentication	302
	Local authentication	302
	Remote authentication	303
	BPDU Guard	304
	Cisco IOS	304
	IBM isCLI	304
	DHCP snooping	304
	Cisco IOS	304
	IBM isCLI	305
	Hostname and DNS server configuration	305
	Cisco IOS	305
	IBM isCLI	305
	Banner configuration	305
	Cisco IOS	305
	IBM isCLI	306
	Interface speed and duplex	306
	Cisco IOS	306
	IBM isCLI	306
	LLDP	306
	Cisco IOS	307
	IBM isCLI	307
	Management network configuration	307

NTP	307
Cisco IOS	307
IBM isCLI	308
OSPF configuration	308
Cisco IOS	308
IBM isCLI	308
Port mirroring	309
Cisco IOS	309
IBM isCLI	309
SNMP	309
Cisco IOS	309
IBM isCLI	310
Spanning Tree Protocol (STP)	311
Cisco IOS	311
IBM isCLI	311
SSH and Telnet	311
Cisco IOS	312
IBM isCLI	312
Syslog	312
Cisco IOS	312
IBM isCLI	312
Port aggregation (static)	313
Cisco IOS	313
IBM isCLI	313
Port aggregation (LACP)	313
Cisco IOS	313
IBM isCLI	314
VLAN tagging (802.1q)	314
Cisco IOS	314
IBM isCLI	314
Appendix B. Easy Connect	315
Introduction to IBM Easy Connect	316
Single Mode	316
Storage Mode	318
Easy Connect Multi-Chassis Mode	320
Implementation with CN/EN4093/R	321
Implementation with G8264	323
Customer examples with diagrams	324
Telecommunications customer	325
State government customer	325
Medical center customer	326
Easy Connect limitations	327
Related publications	329
IBM Redbooks	329
Online resources	329
Help from IBM	329

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 grant you any license to these patents. You can send license inquiries, in writing, to:

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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


Trademarks

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

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

AIX®
BladeCenter®
developerWorks®
DS4000®
Extreme Blue®
IBM®
IBM Flex System™

Micromuse®
Netcool®
POWER®
PureFlex™
PureSystems™
RackSwitch™
Redbooks®

Redbooks (logo) ®
System Storage®
System x®
Tivoli®
VMready®
zEnterprise®

The following terms are trademarks of other companies:

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

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

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

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

Preface

To meet today's complex and ever-changing business demands, you need a solid foundation of server, storage, networking, and software resources. It must be simple to deploy and able to quickly and automatically adapt to changing conditions. You also need access to, and the ability to take advantage of, broad expertise and proven best practices in systems management, applications, hardware maintenance, and more.

IBM® PureFlex™ System is part of the IBM PureSystems™ family of expert integrated systems. It 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. With the PureFlex System, 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.

In this IBM Redbooks® publication, the examples use a Cisco Nexus 5000 Series Switch, although any configurations should also apply to the Cisco Nexus 7000 Series Switch too. However, it is wise to check as there might be minor differences.

This book also covers the different variations for the implementation of these use cases when you use Cisco Catalyst Series Switches.

Authors

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



Jon Tate is a Project Manager for IBM System Storage® SAN Solutions at the International Technical Support Organization, San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing Level 2 support for IBM storage products. Jon has 26 years of experience in storage software and management, services, and support, and is both an IBM Certified IT Specialist and an IBM SAN Certified Specialist. He is also the UK Chairman of the Storage Networking Industry Association.



Jure Arzenšek is an Advisory IT Specialist for IBM Slovenia, and works for the EMEA level 2 team supporting PureFlex and IBM BladeCenter® products. He has been with IBM since 1995, and has worked in various technical support and technical education roles. Jure holds a degree in Computer Science from the University of Ljubljana. His other areas of expertise include IBM System x® servers; SAN; System Storage DS3000, DS4000®, and DS5000 products; and network operating systems for the Intel platform. He has co-authored eleven other IBM Redbooks publications.



David Cain is a network and systems engineer for the IBM Software Group in Research Triangle Park, North Carolina. He has nine years of experience in the data center, with expertise in Ethernet switching, storage, SAN, security, virtualization, System x, and Linux server infrastructure. Dave holds a Bachelor of Science degree in Computer Science from North Carolina State University, and has co-authored two patents and invention disclosures in the networking field. He joined IBM full-time in 2006 after gaining valuable experience on various internships with IBM while a student, including an Extreme Blue® internship in 2005.



William King works for IBM Software Group, Tivoli® Division, IBM UK, as part of the Network Management team. His role is as a network architect developing scenarios on the test network that is used by the ITNM and ITNCM development teams. As a former Micromuse® employee, he has been working on the Tivoli Netcool® suite of products for over 10 years. He is familiar with a wide range of different network equipment from optical and MPLS WAN topologies to data center Fibre Channel and iSCSI storage. He has worked with Cisco, Juniper, Huawei, Nortel, IBM System Networking, Brocade, Foundry, and Extreme equipment. He has a PhD in Immunology from Birmingham University



Gaston Sancassano Rodriguez is a Network Specialist for IBM Uruguay. He has almost seven years of experience working in the design and implementation of networking and security projects. His main specialties include routing, switching, and wireless. He holds an Engineering degree in Telecommunications from Universidad ORT, and several Cisco and Juniper certifications in routing and switching.



Tiago Nunes dos Santos is a Gold Redbooks author and the Infrastructure Strategy leader for the IBM Linux Technology Center, IBM Brazil. He is a Staff Software Engineer and specialized System Administrator, and an expert on the Operating Systems/Application stack, network architecture, and IT User Support processes. Tiago has been working on both Enterprise and Open Source community for over seven years, accumulating expertise in innovation, IT architecture, and strategy leadership. His knowledge on IT Infrastructure architecture helped him become an IBM Inventor, and he is also a member of the Brazilian developerWorks® technical reviewing board.

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

Sangam Racherla
International Technical Support Organization, San Jose Center

Syed Jawad Ameer
Pushkar Patil
Tim Shaughnessy
Yousuff Shaik
IBM San Jose

Scott Lorditch
IBM Denver

Scott Irwin
IBM Dallas

Shawn Raess
Dilminder (Tony) Virk
IBM Raleigh

Casimer DeCusatis
IBM Dallas

The authors would also like to express their thanks to the following Cisco Systems people for their support of this project:

Santiago Freitas
Cisco Systems

Now you can become a published author, too!

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

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

- Send your comments in an email to:

redbooks@us.ibm.com

- Mail your comments to:

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

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/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

This chapter addresses some of the key requirements and trends in the data center that influence the purchase decision of networking hardware.

This chapter includes the following sections:

- ▶ Networking requirements
- ▶ Data center architecture
- ▶ The goal of this book
- ▶ Networking equipment

1.1 Networking requirements

Computer networking is now pervasive. The massive network traffic that is generated by multimedia data (audio, video, streaming), cloud oriented storage and big data is driving research and development into increasing bandwidth and lowering latency.

Combined with CFO demands for minimizing infrastructure costs while keeping (or improving) the services offered, data center administrators are using virtualization techniques to maximize resource usage.

Data center consolidation means that network architects must consider constantly changing business needs and respond to them. They must often do with minimal disruption to the business in a way that is transparent to the users.

Additionally, 10 Gb Ethernet is beginning to replace 1 Gb Ethernet technology as the base on which the new data center environment is built.

As always, network security is becoming more critical as data can be accessed from many different devices at any time of the day. This is one of the main reasons that cloud computing is gaining the foothold that it has. Local data is becoming less and less in demand, with data being stored in the cloud. This means that the network, both virtual and physical, must be able to support the high-bandwidth demands of its users.

Putting all this together, it comes down to two main choices:

- ▶ Implement an industry standards-based data center to ensure the highest level of interoperability between vendors.
- ▶ Implement products from only one vendor.

The IT industry has worked for a long time towards developing industry-based standards that ensure that their products work with other vendors in certain areas. But this does not mean that each company loses the ability to design and develop products and features that give them a competitive advantage.

However, adopting a single vendor-only strategy automatically implies that interoperability barriers are eliminated. It can also mean that if a vendor's competitors introduce new and exclusive solutions, the client might be locked into the single vendor and miss the latest enhancements.

The IBM strategy is to work with the standards bodies, and is active on a number of committees that are driving industry-wide standards.

1.2 Data center architecture

For a long time, the prevailing data center network design paradigm was to have each application attached to its exclusive system server. This approach is well-described and documented, and there is much network design expertise built around it.

According to this paradigm, any application that is associated to its particular system physically resides on that system, and is uniquely identified by the properties of the system in the network. These properties can be elements such as the physical network address of the adapter for that server, or a port number on a switch to which the server is connected, among others.

Virtualized servers solutions are the industry's answer to fit the growing requirements for power and cooling efficiency, optimizing resource utilization. Server virtualization technologies support the effective increase of resource utilization, while also lowering operational and management costs. With this approach, each physical server hosts multiple virtual machines (VMs), and applications that are run in these VMs. There is usually one application per VM. Physical NICs and HBAs are shared between VMs to provide network and storage connectivity.

This virtualized environment approach demands new thinking during network design.

The natural path to the next paradigm is to adopt converged network solutions, which are gaining popularity because of their reduction of network complexity, simplification of network management, and overall improvement on data center operation tasks. This network approach is inherent in the IBM Flex System™ family.

Each business approaches and overcomes these issues in different ways, depending on the company culture and its history. Because of this, no network infrastructure is identical. For more information, see *IBM Flex System Networking in an Enterprise Data Center*, REDP-4834.

1.2.1 The IBM PureFlex System and IBM Flex System family

The IBM PureFlex System and the IBM Flex System products are the next generation of Smarter Computing. They offer intelligent workload deployment and management for maximum business agility. This chassis delivers high-speed performance with integrated servers, storage, and networking for multi-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. Although increased security and resiliency protect vital information and promote maximum uptime, the integrated, easy-to-use management system reduces setup time and complexity, thus providing a quicker path to return on investment (ROI).

With the release of IBM Flex systems, IBM launched a second hybrid computing platform to the market. While IBM zEnterprise® with zBX is focused on mainframe affine applications with a simplified workload-oriented management approach, PureSystems offers a large variety of implementation possibilities focused on a cloud-oriented customer strategy. This new platform adopts a building blocks approach and they are: Management, Compute Nodes, and Storage Networking.

The IBM PureSystems and IBM Flex System family provides a large choice of adapters and switches. All components are standard-based and integrated into the management of the

chassis. This variety provides a combination of features that fits into the existing infrastructure. The modular concept offers the possibility to adapt to future requirements.

A connection to an existing network is required to use the capabilities of PureSystems, in most cases. However, modern data centers rely on a complex network infrastructure. The introduction of active networking components within an existing infrastructure can affect all components and introduce risks. Therefore, many customers are reluctant to introduce such solutions.

1.3 The goal of this book

The goal of this book is to demonstrate that the new IBM PureFlex Systems family can interoperate with Cisco switches.

It does so by implementing practical use case scenarios that involve typical setups that are used by industry customers, and detail the steps needed to configure them.

1.4 Networking equipment

This section describes the IBM System Networking, Cisco Nexus Series Switches, and Cisco Catalyst Series Switches product families.

1.4.1 IBM System Networking

In today's infrastructure, it is common to build networks that are based on 10 Gb Ethernet technology. The IBM portfolio of 10 Gb system networking products includes Top-of-Rack switches, and the embedded switches in the IBM PureFlex System and IBM Flex System families. In 2010, IBM formed the IBM System Networking business (by acquiring BLADE Network Technologies), which is now focused on driving data center networking by using the latest in Ethernet technologies.

The physical layout of most corporate networks has evolved over time. Classic hub/router topologies have given way to faster switched topologies, particularly now that switches are increasingly intelligent. IBM System Networking switches are intelligent and fast enough to run routing functions on par with wire-speed Layer 2 switching. The combination of faster routing and switching in a single device provides another service: You can build versatile topologies that account for earlier configurations.

IBM System Networking switches support up to 1024 VLANs per switch. Even though the maximum number of VLANs supported at any time is 1024, each can be identified by a number from 1 - 4095.

In a routed environment, routers communicate with one another to track available routes. Routers can learn about available routes dynamically by using the Routing Information Protocol (RIP). IBM Networking OS supports RIP version 1 (RIPv1) and RIP version 2 (RIPv2) for exchanging TCP/IPv4 route information with other routers.

IBM System Networking also currently supports these standards and technologies:

- ▶ VLAN: Virtual Local Area Network
 - PVID: PortVLAN IDs
 - VLAN tagging
 - PVLAN: protocol-based VLANs
- ▶ STP: Spanning Tree Protocol
 - RSTP: Rapid Spanning Tree Protocol
 - PVRST: Per-VLAN Rapid Spanning Tree Protocol
 - MSTP: Multiple Spanning Tree Protocol
- ▶ IP routing: Internet Protocol routing
 - Static routes
 - ECMP: Equal-Cost Multi-Path static routes
 - RIP: Routing Information Protocol (RIPv1, RIPv2)
- ▶ OSPF: Open Shortest Path First
- ▶ BGP: Border Gateway Protocol
 - eBGP
 - iBGP
- ▶ IP multicast
- ▶ IGMP: Internet Group Management Protocol
- ▶ PIM: Protocol Independent Multicast
 - PIM Sparse Mode
 - PIM Dense Mode
- ▶ IPv6: Internet Protocol version 6
- ▶ ND: Neighbor Discovery protocol
- ▶ Port mirroring
- ▶ ACL-based mirroring
- ▶ sFlow monitoring
- ▶ RMON: Remote Monitoring
- ▶ Trunking
 - Static trunk groups (portchannel)
 - Dynamic LACP trunk groups
- ▶ LACP: Link Aggregation Control Protocol (IEEE 802.3ad)
- ▶ VLAG: Virtual Link Aggregation Groups
- ▶ Fast Uplink Convergence
- ▶ NIC teaming and Layer 2 failover
- ▶ VRRP: Virtual Router Redundancy Protocol
- ▶ AMP: Active Multipath Protocol

Also, the stacking capability provides the ability to implement a group of up to eight IBM System Networking switches that work together as a unified system. Stacking is supported only on Virtual Fabric 10Gb Switch Module devices.

For more information about these features, see *Implementing IBM System Networking 10Gb Ethernet Switches*, SG24-7960.

The examples in this book use the IBM Flex System Fabric EN4093 10Gb Scalable Switch. However, it equally applies to the IBM Flex System Fabric EN4093R 10Gb Scalable Switch. For more information about these switches, see:

<http://www.redbooks.ibm.com/abstracts/tips0864.html>

1.4.2 Cisco Switches

Cisco Nexus Series Switches

The Cisco Nexus family of data center-class switches was developed on the Cisco Data Center 3.0 framework. It is designed to help build, implement and operate a virtualized, next-generation data center. For more information, see:

http://www.cisco.com/en/US/products/ps9441/Products_Sub_Category_Home.html

The Cisco Nexus Family of switches is designed to allow you to upgrade to 10-Gigabit Ethernet in a granular, cost-effective manner as part of your data center transformation strategies.

Cisco Catalyst Series Switches

Cisco Catalyst Series Switches are a high performance Top-of-Rack Switch. Because of its flexibility, it is able to handle small/medium businesses. However, it can also be used in large data centers.

Use cases approach

the examples in this book use a Cisco Nexus 5000 Series Switch because of its popularity in customer data centers and because it can handle most of the features customers are currently using.

Any configurations that are detailed here generally apply to the Cisco Nexus 7000 Series Switch as well, but it is wise to check as there might be minor differences.

Different variations for the implementation of these use cases when using Cisco Catalyst Series Switches are also covered.



Layer 1 Overview

This chapter provides details about IBM PureFlex System networking from the physical layer perspective. It explains Layer 1 networking concepts and terminology, and describes IBM PureFlex System networking components (midplane connections, switches, adapters). It also lists transceivers and cables that are used with IBM PureFlex System networking options.

This chapter includes the following sections:

- ▶ Layer 1 networking concepts and terminology
- ▶ Physical layer on IBM Flex System Enterprise Chassis
- ▶ IBM Flex System Ethernet I/O modules
- ▶ IBM Flex System Ethernet adapters

2.1 Layer 1 networking concepts and terminology

Layer 1 of the OSI model is the layer at which the physical transmission of data occurs. This section explains some of the common concepts that are important at the Layer 1 level. These include Ethernet cabling, copper and Fibre Channel media, transceivers and Direct Attached Cables, and physical configuration parameters.

2.1.1 Ethernet cabling

Ethernet cabling typically comes in one of two forms: Copper cabling or fiber optic cabling. Copper is the less expensive choice in terms of materials, components, and installation cost. Copper cabling is the method that is commonly used to connect devices to the access layer switches.

Fiber optic cabling comes at a higher cost than copper cabling. The optical components for devices and switches and the cost of any customer cabling is typically higher. However, the higher costs are often easily justified by the benefits of fiber optic cabling. Fiber optic cabling yields longer cable lengths and is immune to signal distortion that is caused in copper cabling by electromagnetic interference.

2.1.2 Twisted-pair copper cabling

Twisted-pair copper cabling is a common media for Ethernet networking installations. Twisted-pair cabling is available as unshielded twisted pair (UTP) or shielded twisted pair (STP). This shielding helps prevent electromagnetic interference.

Several different categories of twisted-pair cabling are available as listed in Table 2-1. These categories indicate the signaling capabilities of the cabling.

Table 2-1 TIA/EIA cabling categories

TIA/EIA cabling category	Maximum network speeds supported
Cat 1	Telephone or ISDN
Cat 2	4 Mb Token Ring
Cat 3	10 Mb Ethernet
Cat 4	16 Mb Token Ring
Cat 5	100 Mb Ethernet
Cat 5e	1 Gb Ethernet
Cat 6	10 Gb Ethernet Short Distance - 55 m (180 ft.)
Cat 6a	10 Gb Ethernet

The RJ45 connector used for Ethernet twisted-pair cabling is the most recognizable and associated with networking. The RJ45 connector is shown in Figure 2-1.




Figure 2-1 RJ45 Copper Connector

Twisted-pair cabling contains four pairs of wire inside the cable, as illustrated in Figure 2-2.




Figure 2-2 Straight through Ethernet cable

An Ethernet operating in 10/100 Mb mode uses only two pairs, pairs 1-2 and 3-6. An Ethernet operating in 1 Gb mode uses all four pairs: Pairs 1-2, 3-6, 4-5, and 7-8. Distances up to 100 meters are supported.

Twisted-pair crossover requirements

In 10/100 Mbps Ethernet operations, one pair of wire is used for data transmission and one pair is used for receiving data. When a device, such as a PC, is attached to a hub or switch, the ports are designed so that the transmitting and receiving pairs are properly matched. When directly connecting two like devices, such PC-PC, hub-hub, or switch-switch, a crossover in the pairs must be made.

A crossover function can be made internally by the port of one of the devices, or can be achieved by using a crossover cable as illustrated in Figure 2-3.




Figure 2-3 10/100 Mbps crossover cable

Ethernet ports without crossover are known as *Medium Dependent Interface* (MDI). Ports with crossover are known as *Medium Dependent Interface Crossover* (MDIX). The “X” means crossover. To simplify cabling, ports can sense whether crossover is needed and configure the port properly. This function is known as *Auto MDIX*. For Gigabit Ethernet, the auto crossover function is an optional part of the 1000Base-T Ethernet standard.

Today’s 1 Gb and 10 Gb Ethernet switches typically use Auto MDIX to automatically determine the correct port configuration.

2.1.3 Fiber optic cabling

In copper cabling, electric signals are used to transmit data through the network. The copper cabling is the medium for that electrical transmission. In fiber optic cabling, light is used to transmit the data. Fiber optic cabling is the medium for channeling the light signals between devices in the network.

Two modes of fiber optic signaling are single-mode and multimode. The difference between the modes is the wavelength of the light used for the transmission as illustrated in Figure 2-4.




Figure 2-4 Multimode versus single-mode optic signaling

Single-mode fiber

Single-mode optical fiber (SMF) uses long wavelength light to transmit data and requires a cable with a small core for transmission (Figure 2-5). The core diameter for single-mode cabling is 9 microns in diameter.




Figure 2-5 Single-mode fiber cable

Single-mode fiber cabling allows for much longer cable lengths than multimode. For example, when you use 10GBASE-ER transceivers and suitable single mode fiber cable, it is possible to reach distances up to 40 km.

Multimode fiber

Multi-mode optical fiber (MMF) uses short wavelength light to transmit data, and requires a cable with a larger core for transmission (Figure 2-6). The core diameter for multimode cabling can be 50 or 62.5 microns in diameter.




Figure 2-6 Multimode fiber cable

The color of the outer coating is sometimes used to identify if a cable is a multimode or single-mode fiber cable, but the color is not a reliable method. The TIA-598C standard suggests the outer coating to be yellow for single mode fiber and orange for multimode fiber for civilian applications.

This guideline is not always implemented as shown in Figure 2-7, which shows a blue cable.



Figure 2-7 Blue 62.5 micron MMF cable

The reliable method is to look at the specifications of the cable printed on the outer coating of the cabling. Figure 2-8 shows an SMF cable in the standard yellow.




Figure 2-8 Yellow SMF cable

Figure 2-9 shows an MMF cable in the standard orange.



Figure 2-9 Orange 50 micron MMF cable

With multimode cabling, 10 Gbps Ethernet supports cable lengths of up to 550 m, and 40 Gbps Ethernet supports cable lengths of up to 125 m.

Connector types

The most common connector type for fiber optic media that is used in networking today is the LC connector, which is shown in Figure 2-10.




Figure 2-10 LC fiber connector

Other connectors that are commonly encountered in Ethernet networks are the SC connector (Figure 2-11), and the ST connector (Figure 2-12).




Figure 2-11 SC fiber connector

Figure 2-12 shows the ST connector.



Figure 2-12 ST fiber connectors

Transceivers

A *transceiver* or *transmitter/receiver* is the fiber optic port of a device. It is where the fiber optic cables connect. Transceiver performs conversion from electric signals to optical, and vice versa. Some devices might have an integrated transceiver, which limits the flexibility in the type of cabling that can be used. However, most devices provide a slot for a modular transceiver to be inserted, providing flexibility of use for single or multimode implementations.

In today's Ethernet networks, SFP, SFP+, XFP, and QSFP transceivers are typically used.

Figure 2-13 shows SFP, SFP+, and XFP transceivers (MMF and SMF varieties).



Figure 2-13 From left to right: SFP-MMF, SFP-SMF, SFP+-MMF, XFP-MMF, and XFP-SMF

A QSFP transceiver is shown in Figure 2-14.



Figure 2-14 QSFP transceiver

Table 2-2 shows comparison between different transceiver types.

Table 2-2 Comparison of transceivers

Type	Data rate	Supported standards
SFP	Up to 4.25 Gbps	1 Gb Ethernet, up to 4 Gb FC
SFP+	Up to 10 Gbps	10 Gb Ethernet, 8 Gb FC, OTU2
XFP	10 Gbps	10 Gb Ethernet, 10 Gb FC, SONET, OTU2
QSFP	40 Gbps	40 Gb Ethernet, 20 Gb/40 Gb InfiniBand

Direct Attach Cable

Direct Attach Cable (DAC) is a twinaxial (twinax) cable that can be used in 10 Gb Ethernet environments. The DAC has SFP+ housing on each end, which means that you can plug it directly into SFP+ slot on a switch. DAC can be either passive or active:

- ▶ Passive DAC contains no active components, and supports cable lengths up to five meters.
- ▶ Active DAC contains active electronic components in SFP+ housing for enhanced signal quality, and supports cable lengths up to ten meters.

DAC cables are quite cost-effective in comparison with FC cables in cases when short cable lengths are sufficient.

Figure 2-15 shows a DAC example: 3m IBM Passive DAC SFP+ cable, P/N 90Y9430.



Figure 2-15 3m IBM Passive DAC SFP+ cable, P/N 90Y9430

2.1.4 Physical configuration parameters

The physical layer (Layer 1) properties include elements such as line speed and duplex.

Speed

Speed in Ethernet refers to data rates such as 10 Mbps, 100 Mbps, 1 Gbps, 10 Gbps, and 40 Gbps.

Duplex

Duplex modes are either full or half duplex. *Half duplex* is when a device can only send or receive at a time (Figure 2-16).




Figure 2-16 Half-duplex mode

Full duplex devices can send and receive at the same time (Figure 2-17).




Figure 2-17 Full-duplex mode

Autonegotiation

In an Ethernet network, the speed and duplex of a device that is attached to a segment must match. Autonegotiation of the speed and duplex of a device usually works well, but it is not 100% reliable. The problems usually occur with older 10/100 devices. Newer devices rarely have an issue when negotiating with each other.

One step to reduce negotiation problems is to ensure that both devices on a switch segment are configured the same. Either configure both devices for autonegotiation, or “hard code” (manually configure) both the speed and duplex settings of both devices to the same settings.

2.2 Physical layer on IBM Flex System Enterprise Chassis

This section provides IBM Flex System Enterprise Chassis networking physical layer details. It addresses physical connectivity between network adapters installed in compute nodes and switches installed in I/O bays. It also describes the switches and network adapters available for IBM Flex System.

The Ethernet networking I/O architecture for the IBM Flex System Enterprise Chassis includes an array of connectivity options for compute nodes installed in the enclosure. Users can decide to use a local switching model that provides superior performance, cable reduction and a rich feature set. Or they can use pass-through technology and allow all Ethernet networking decisions to be made external to the Enterprise Chassis.

By far, the most versatile option is to use modules that provide local switching capabilities and advanced features that are fully integrated into the operation and management of the Enterprise Chassis. In particular, the EN4093/EN4093R 10Gb Scalable Switch module offers the maximum port density, highest throughput, and most advanced data center-class features to support the most demanding compute environments.

The Enterprise Chassis has four I/O bays in the rear of the chassis. This is where you can install up to four network switch modules. The physical layout of these I/O module bays is shown in Figure 2-18.




Figure 2-18 Rear view of the Enterprise Chassis showing I/O module bays

From a midplane wiring point of view, the Enterprise Chassis provides 16 lanes between each half-wide compute node bay and each I/O bay. Each lane can provide 16 Gbps or higher speeds. How these lanes are used depends on these factors:

- ▶ Network adapters that are installed in a node
- ▶ I/O module that is installed in the I/O bay
- ▶ Port licenses enabled on the I/O module

Figure 2-19 shows how the midplane lanes connect between the compute node bays up front and the I/O bays in the rear. The concept of an I/O module partition is also illustrated in Figure 2-19. From a physical perspective, a partition in this context is a bank of 14 ports that can be implemented on a switch module. By default, all I/O modules include the base partition, and thus have 14 internal ports, each connected to a corresponding compute node bay in the front. By adding an upgrade license to the I/O module, you can add more banks of 14 ports (partitions) to an I/O module (assuming that module supports the partition). If a node is connected to one of the ports on one of the additional partitions, that partition is enabled through an upgrade on the I/O module. The node needs an adapter that has the necessary physical ports to connect to the wanted lanes. Those lanes connect to the ports in the I/O partition that is enabled on the I/O module.

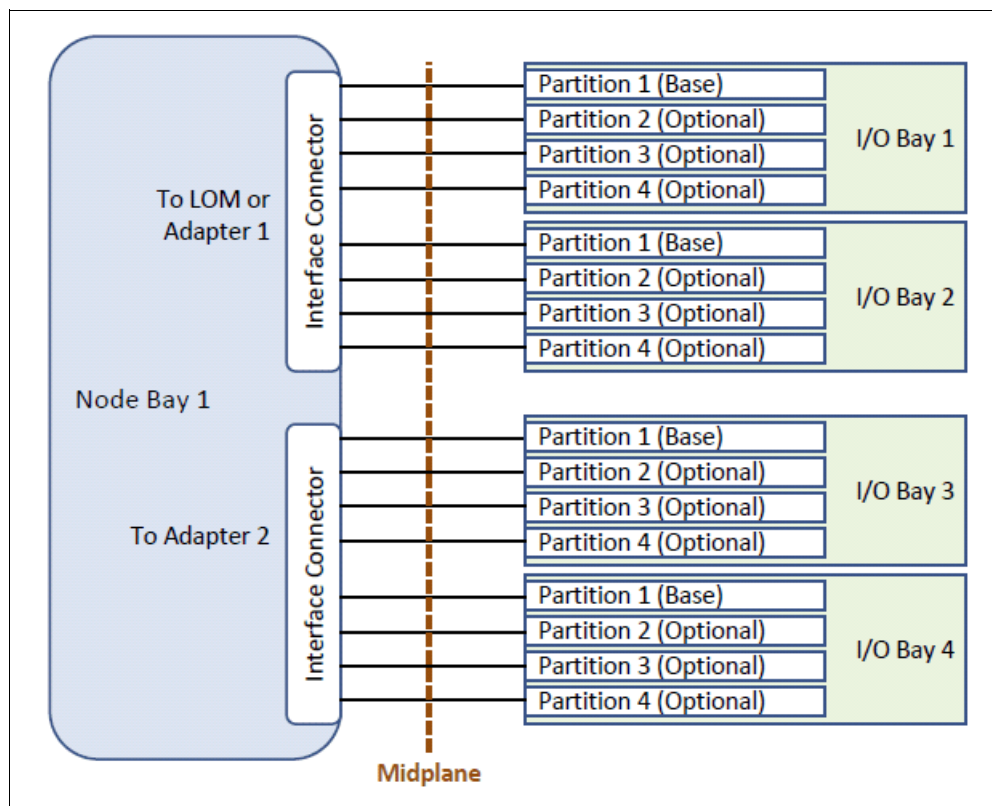


Figure 2-19 Sixteen lanes total of a single half-wide node bay toward the I/O bays

For example, if a dual port LAN on motherboard (LOM) adapter is installed on compute node, only two of the 16 lanes are used (one to I/O bay 1 and one to I/O bay 2), as shown in Figure 2-20 on page 18.

If two quad port network adapters are installed on compute node, eight of the 16 lanes are used (two to each of the four I/O bays).

This installation can provide up to 320 Gbps of full duplex Ethernet bandwidth (16 lanes x 10 Gbps x 2) to a single half-wide node, and up to 640 Gbps of bandwidth to a full-wide node.

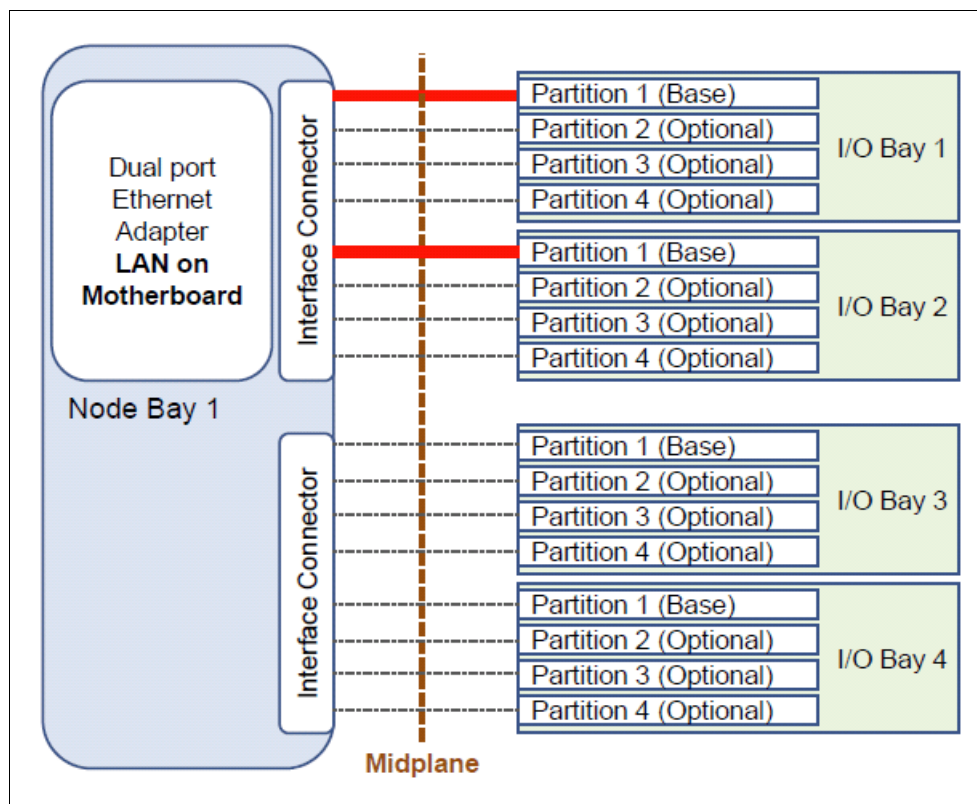


Figure 2-20 Dual port LOM connecting to partition on I/O bays 1 and 2 (all other lanes unused)

There are limits on the port density of the node network adapters and the number of ports available from each switch in the I/O bays that lead to the nodes. However, the Enterprise Chassis can easily scale to high bandwidth to meet demand.

Nodes are currently limited to a maximum of two quad port adapters on a single half-wide node. This limits the connection to eight lanes of 10 Gb Ethernet for a half-wide server.

On the I/O module side, the number of links that connect to the lanes toward the nodes is the gating factor. By default, each I/O module provides a single connection (lane) to each of the 14 half-wide node bays up front. By adding port licenses, a single EN2092 1Gb Ethernet Switch can offer two 1 Gb ports to each half-wide node bay,. The EN4093/EN4093R 10Gb Scalable Switch can provide up to three 10 Gb ports to each of the 14 half-wide node bays. Because it is a one-for-one 14-port pass-through, the EN4091 10Gb Ethernet Pass-thru I/O module can only ever offer a single link to each of the half-wide node bays.

All I/O modules include a base partition of 14 downstream ports, with the pass-through module supporting only the single partition. Both the EN4093/EN4093R 10Gb Scalable Switch and the EN2092 1Gb Ethernet Switch support more than the base partition. Table 2-4 on page 23 and Table 2-3 on page 19 show the available I/O module partition upgrades.

At the time of writing, no I/O modules and node adapter combinations can use all 16 lanes between a compute node bay and the I/O bays. The extra lanes ensure that the Enterprise Chassis can accommodate future capacity demands.

2.3 IBM Flex System Ethernet I/O modules

The IBM Flex System Enterprise Chassis features a number of Ethernet I/O module solutions that provide a combination of 1 Gb and 10 Gb ports to the servers, and 1 Gb, 10 Gb, and 40 Gb for uplink connectivity 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 three Ethernet I/O modules available for deployment with the Enterprise Chassis:

- ▶ “IBM Flex System EN2092 1Gb Ethernet Scalable Switch”
- ▶ “IBM Flex System Fabric EN4093/EN4093R 10 Gb Scalable Switch”
- ▶ “IBM Flex System EN4091 10Gb Ethernet Pass-thru module”

2.3.1 IBM Flex System EN2092 1Gb Ethernet Scalable Switch

The EN2092 1Gb Ethernet Switch is primarily a 1 Gb switch, offering up to 28 x 1 Gb downlinks to the internal nodes. It has a total combination of up to 20 x 1 Gb RJ45 uplinks and four 10 Gb uplinks with “pay-as-you-grow” scalability.

Figure 2-21 shows the EN2092 1Gb Ethernet Switch.



Figure 2-21 The EN2092 1Gb Ethernet Switch

Ports that are enabled and available depend on the features activated on the I/O module. Table 2-3 describes the port configurations for the EN2092 1Gb Ethernet Switch.

Table 2-3 Port counts for EN2092 1Gb Ethernet Switch

Part number	Product name	Switch function	Total ports	
49Y4294	IBM Flex System EN2092 1 Gb Ethernet Switch	14x 1 Gb internal ports and 10x 1 Gb uplinks	14x 1 Gb internal, 10x 1 Gb uplinks	
90Y3562	IBM Flex System EN2092 1 Gb Ethernet Switch (Upgrade 1)	Adds extra 14x 1 Gb internal ports and extra 10x 1 Gb external uplinks	28x 1 Gb internal 20x 1 Gb uplinks	28x 1 Gb internal 20x 1 Gb uplinks 4x 10 Gb uplinks
49Y4298	IBM Flex System EN2092 1 Gb Ethernet Switch (10 Gb Uplinks)	Enables the 4x 10 Gb external uplink ports	14x 1 Gb internal 10x 1 Gb uplinks 4x 10 Gb uplinks	

Upgrade 1 and the 10 Gb Uplinks upgrade do not depend on each other. If only one upgrade is activated, the total number of enabled ports is shown in the respective row in Table 2-3. The table also shows the number of enabled ports when both upgrades are activated.

The EN2092 1 Gb Ethernet Scalable Switch has the following features and specifications:

- ▶ Internal ports:
 - Twenty-eight internal full-duplex Gigabit ports with 14 ports that are enabled by default; an optional Features on Demand (FoD) capability license is required to activate the other 14 ports
 - Two internal full-duplex 1 GbE ports connected to the chassis management module
- ▶ External ports:
 - Four ports for 1 Gb or 10 Gb Ethernet SFP+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10 GBASE-SR, or 10 GBASE-LR) or SFP+ copper DACs. These ports are disabled by default, and an optional FoD license is required to activate them. SFP+ modules are not included and must be purchased separately.
 - A total of 20 external 10/100/1000 1000BASE-T Gigabit Ethernet ports with RJ-45 connectors (10 ports are enabled by default, an optional FoD license is required to activate the other 10 ports).
 - One RS-232 serial port (mini-USB connector) that provides an extra means to configure the switch module.
- ▶ Scalability and performance:
 - Fixed-speed external 10 Gb Ethernet ports for maximum uplink bandwidth
 - Autosensing 10/1000/1000 external Gigabit Ethernet ports for bandwidth optimization
 - Non-blocking architecture with wire-speed forwarding of traffic
 - Media Access Control (MAC) address learning:
 - Automatic update
 - Support of up to 32,000 MAC addresses
 - Up to 128 IP interfaces per switch
 - Static and LACP (IEEE 802.3ad) link aggregation with up to:
 - 60 Gb of total uplink bandwidth per switch
 - 64 trunk groups
 - 16 ports per group
 - Support for jumbo frames (up to 9,216 bytes)
 - Broadcast/multicast storm control
 - Internet Group Management Protocol (IGMP) snooping for limit flooding of IP multicast traffic
 - IGMP filtering to control multicast traffic for hosts that participate in multicast groups
 - Configurable traffic distribution schemes over trunk links that are based on source/destination IP, MAC addresses, or both
 - Fast port forwarding and fast uplink convergence for rapid STP convergence
- ▶ Availability and redundancy:
 - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy
 - IEEE 802.1D STP for providing L2 redundancy
 - IEEE 802.1s Multiple STP (MSTP) for topology optimization, up to 32 STP instances supported by single switch
 - IEEE 802.1w Rapid STP (RSTP) provides rapid STP convergence for critical delay-sensitive traffic, such as voice or video

- Per-VLAN Rapid STP (PVRST) enhancements
- Layer 2 Trunk Failover to support active/standby configurations of network adapter teaming on compute nodes
- Hot Links provides basic link redundancy with fast recovery for network topologies that require Spanning Tree to be turned off
- ▶ VLAN support:
 - Up to 1024 VLANs supported per switch, with VLAN numbers that range from 1 - 4095 (4095 is used for the connection of the management module only)
 - 802.1Q VLAN tagging support on all ports
 - Private VLANs
- ▶ Security:
 - VLAN-based, MAC-based, and IP-based ACLs
 - 802.1x port-based authentication
 - Multiple user IDs and passwords
 - User access control
 - Radius, TACACS+, and LDAP authentication and authorization
- ▶ Quality of service (QoS):
 - Support for IEEE 802.1p, IP ToS/DSCP, and ACL-based (MAC/IP source and destination addresses, VLANs) traffic classification and processing
 - Traffic shaping and remarking based on defined policies
 - Eight weighted round robin (WRR) priority queues per port for processing qualified traffic
- ▶ IP v4 Layer 3 functions:
 - Host management
 - IP forwarding
 - IP filtering with ACLs, up to 896 ACLs supported
 - VRRP for router redundancy
 - Support for up to 128 static routes
 - Routing protocol support (RIP v1, RIP v2, OSPF v2, and BGP-4), up to 2048 entries in a routing table
 - Support for DHCP Relay
 - Support for IGMP snooping and IGMP relay
 - Support for Protocol Independent Multicast (PIM) in Sparse Mode (PIM-SM) and Dense Mode (PIM-DM).
- ▶ IP v6 Layer 3 functions:
 - IPv6 host management (except default switch management IP address)
 - IPv6 forwarding
 - Up to 128 static routes
 - Support for OSPF v3 routing protocol
 - IPv6 filtering with ACLs

- ▶ Virtualization:
 - IBM VMready®
- ▶ Manageability:
 - Simple Network Management Protocol (SNMP V1, V2, and V3)
 - HTTP browser GUI
 - Telnet interface for CLI
 - SSH
 - Serial interface for CLI
 - Scriptable CLI
 - Firmware image update (TFTP and FTP)
 - Network Time Protocol (NTP) for switch clock synchronization
- ▶ Monitoring:
 - Switch LEDs for external port status and switch module status indication
 - Remote Monitoring (RMON) agent to collect statistics and proactively monitor switch performance
 - Port mirroring for analyzing network traffic that passes through the switch
 - Change tracking and remote logging with the **syslog** feature
 - Support for the sFLOW agent for monitoring traffic in data networks (to monitor elsewhere in the network, you need an external sFLOW analyzer)
 - POST diagnostic tests

For more information, see *IBM Flex System EN2092 1Gb Ethernet Scalable Switch*, TIPS0861 at:

<http://www.redbooks.ibm.com/abstracts/tips0861.html>

2.3.2 IBM Flex System Fabric EN4093/EN4093R 10 Gb Scalable Switch

The EN4093/EN4093R 10Gb Scalable Switch is primarily a 10 Gb switch. It can provide up to 42 10 Gb internal node-facing ports, and up to 14 SFP+ 10 Gb and two QSFP+ 40 Gb external upstream facing ports, depending on the applied upgrade licenses.

A view of the face plate of the EN4093/EN4093R 10Gb Scalable Switch is shown in Figure 2-22.



Figure 2-22 The IBM Flex System Fabric EN4093 10Gb Scalable Switch

Information about available upgrade options for this module is provided in Table 2-4.

Table 2-4 EN4093/EN4093R Scalable Switch

Part number	Product name	Switch function	Total ports
49Y4270	IBM Flex System Fabric EN4093 10Gb Scalable Switch	14x 10 Gb internal ports and 10x 10 Gb uplinks	14x 10 Gb internal 10x 10 Gb uplinks
95Y3309	IBM Flex System Fabric EN4093R 10Gb Scalable Switch	14x 10 Gb internal ports and 10x 10 Gb uplinks	14x 10 Gb internal 10x 10 Gb uplinks
49Y4798	IBM Flex System Fabric EN4093 10Gb Scalable Switch (Upgrade 1)	Adds extra 14x 10 Gb internal ports and enables 2x 40 Gb external uplinks	28x 10 Gb internal 10x 10 Gb uplinks 2x 40 Gb uplinks
88Y6037	IBM Flex System Fabric EN4093 10Gb Scalable Switch (Upgrade 2) ^a	Adds extra 14x 10 Gb internal ports and 4x 10 Gb external uplinks	42x 10 Gb internal 14x 10 Gb uplinks 2x 40 Gb uplinks

a. Upgrade 2 requires Upgrade 1, 49Y4798. Internal ports that are enabled with Upgrade 2 require a 6-port adapter, which is unavailable as of this writing.

The IBM Flex System Fabric EN4093/EN4093R 10 Gb Scalable Switch has the following features and specifications:

- Internal ports:
 - A total of 42 internal full-duplex 10 Gigabit ports (14 ports are enabled by default; optional FoD licenses are required to activate the remaining 28 ports)
 - Two internal full-duplex 1 GbE ports that are connected to the chassis management module
- External ports:
 - A total of 14 ports for 1 Gb or 10 Gb Ethernet SFP+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10 GBASE-SR, or 10 GBASE-LR) or SFP+ copper DACs. Ten ports are enabled by default and an optional FoD license is required to activate the remaining four ports. SFP+ modules and DAC cables are not included and must be purchased separately.
 - Two ports for 40 Gb Ethernet QSFP+ transceivers or QSFP+ DACs (these ports are disabled by default. An optional FoD license is required to activate them). QSFP+ modules and DAC cables are not included and must be purchased separately.
 - One RS-232 serial port (mini-USB connector) that provides an extra means to configure the switch module.
- Scalability and performance:
 - 40 Gb Ethernet ports for extreme uplink bandwidth and performance
 - Fixed-speed external 10 Gb Ethernet ports to use 10 Gb core infrastructure
 - Autosensing 10/100/1000 external Gigabit Ethernet ports for bandwidth optimization
 - Non-blocking architecture with wire-speed forwarding of traffic and aggregated throughput of 1.28 Tbps
 - MAC address learning:
 - Automatic update
 - Support of up to 128,000 MAC addresses
 - Up to 128 IP interfaces per switch

- Static and LACP (IEEE 802.3ad) link aggregation with up to:
 - 220 Gb of total uplink bandwidth per switch
 - 64 trunk groups
 - 16 ports per group
- Support for jumbo frames (up to 9,216 bytes)
- Broadcast/multicast storm control
- IGMP snooping to limit flooding of IP multicast traffic
- IGMP filtering to control multicast traffic for hosts that participate in multicast groups
- Configurable traffic distribution schemes over trunk links based on source/destination IP, MAC addresses, or both
- Fast port forwarding and fast uplink convergence for rapid STP convergence
- ▶ Availability and redundancy:
 - VRRP for Layer 3 router redundancy
 - IEEE 802.1D STP for providing L2 redundancy
 - IEEE 802.1s Multiple STP (MSTP) for topology optimization, up to 32 STP instances are supported by single switch
 - IEEE 802.1w Rapid STP (RSTP) provides rapid STP convergence for critical delay-sensitive traffic, such as voice or video
 - Per-VLAN Rapid STP (PVRST) enhancements
 - Layer 2 Trunk Failover to support active/standby configurations of network adapter that team on compute nodes
 - Hot Links provides basic link redundancy with fast recovery for network topologies that require Spanning Tree to be turned off
- ▶ VLAN support:
 - Up to 1024 VLANs supported per switch, with VLAN numbers that range from 1- 4095 (4095 is used for the connection of the management module only)
 - 802.1Q VLAN tagging support on all ports
 - Private VLANs
- ▶ Security:
 - VLAN-based, MAC-based, and IP-based ACLs
 - 802.1x port-based authentication
 - Multiple user IDs and passwords
 - User access control
 - Radius, TACACS+, and LDAP authentication and authorization
- ▶ Quality of service (QoS):
 - Support for IEEE 802.1p, IP ToS/DSCP, and ACL-based (MAC/IP source and destination addresses, VLANs) traffic classification and processing
 - Traffic shaping and remarking based on defined policies
 - Eight Weighted Round Robin (WRR) priority queues per port for processing qualified traffic

- ▶ IP v4 Layer 3 functions:
 - Host management
 - IP forwarding
 - IP filtering with ACLs, up to 896 ACLs supported
 - VRRP for router redundancy
 - Support for up to 128 static routes
 - Routing protocol support (RIP v1, RIP v2, OSPF v2, and BGP-4), up to 2048 entries in a routing table
 - Support for DHCP Relay
 - Support for IGMP snooping and IGMP relay
 - Support for Protocol Independent Multicast (PIM) in sparse mode (PIM-SM) and dense mode (PIM-DM).
- ▶ IP v6 Layer 3 functions:
 - IPv6 host management (except default switch management IP address)
 - IPv6 forwarding
 - Up to 128 static routes
 - Support of OSPF v3 routing protocol
 - IPv6 filtering with ACLs
- ▶ Virtualization:
 - Virtual Fabric with vNIC (virtual NICs)
 - 802.1Qbg Edge Virtual Bridging (EVB)
 - VMready
- ▶ Converged Enhanced Ethernet:
 - Priority-Based Flow Control (PFC) (IEEE 802.1Qbb) extends 802.3x standard flow control to allow the switch to pause traffic based on the 802.1p priority value in the VLAN tag of each packet
 - Enhanced Transmission Selection (ETS) (IEEE 802.1Qaz) provides a method for allocating link bandwidth based on the 802.1p priority value in the VLAN tag of each packet
 - Data Center Bridging Capability Exchange Protocol (DCBX) (IEEE 802.1AB) allows neighboring network devices to exchange information about their capabilities
- ▶ Manageability:
 - Simple Network Management Protocol (SNMP V1, V2, and V3)
 - HTTP browser GUI
 - Telnet interface for CLI
 - SSH
 - Serial interface for CLI
 - Scriptable CLI
 - Firmware image update (TFTP and FTP)
 - NTP for switch clock synchronization

- Monitoring:
 - Switch LEDs for external port status and switch module status indication
 - Remote Monitoring (RMON) agent to collect statistics and proactively monitor switch performance
 - Port mirroring for analyzing network traffic that passes through switch
 - Change tracking and remote logging with syslog feature
 - Support for sFLOW agent for monitoring traffic in data networks (separate sFLOW analyzer required elsewhere)
 - POST diagnostic testing

For more information, see the *IBM Flex System Fabric EN4093 and EN4093R 10Gb Scalable Switches*, TIPS0864, at:

<http://www.redbooks.ibm.com/abstracts/tips0864.html>

2.3.3 IBM Flex System EN4091 10Gb Ethernet Pass-thru module

The EN4091 10Gb Ethernet Pass-thru module offers one-to-one connections between compute node bays and I/O module uplinks. It has 14 internal ports and 14 external ports. Each internal port is wired to its matching external port.

The module has no management interface, and can support 1 Gb and 10 Gb dual port adapters installed on the nodes. If quad port adapters are used in a node, only the first two ports access the pass-through modules. The necessary 1 Gb or 10 Gb modules (SFP, SFP+, or DAC) must also be installed in the external ports of the pass-thru module to support the wanted speed (1 Gb or 10 Gb) and medium (fiber or copper) for adapter ports on the node.

The EN4091 10Gb Ethernet Pass-thru module is shown in Figure 2-23.



Figure 2-23 The IBM Flex System EN4091 10Gb Ethernet Pass-thru

The part number for the EN4091 10Gb Ethernet Pass-thru module is listed in Table 2-5. There are no upgrades available for this I/O module at the time of writing.

Table 2-5 IBM Flex System EN4091 10Gb Ethernet Pass-thru part number

Part number	Description
88Y6043	IBM Flex System EN4091 10Gb Ethernet Pass-thru

The IBM Flex System EN4091 10 Gb Ethernet Pass-thru includes the following features and specifications:

- ▶ **Internal ports**
A total of 14 internal full-duplex Ethernet ports that can operate at 1 Gb or 10 Gb speeds.
- ▶ **External ports**
A total of 14 ports for 1 Gb or 10 Gb Ethernet SFP+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10 GBASE-SR, or 10 GBASE-LR) or SFP+ copper DACs. SFP+ modules and DAC cables are not included and must be purchased separately.
- ▶ This device is unmanaged and has no internal Ethernet management port. However, it provides its vital product data (VPD) to the secure management network in the Chassis Management Module.

For more information, see the *IBM Flex System EN4091 10Gb Ethernet Pass-thru Module*, TIPS0865, at:

<http://www.redbooks.ibm.com/abstracts/tips0865.html>

2.3.4 Cables and transceivers for I/O modules

Table 2-6 lists supported cables and transceivers for IBM PureFlex System Ethernet I/O modules.

Table 2-6 Modules and cables that are supported in Ethernet I/O modules

Part number	Description	EN2092 1 GbE Switch	EN4093 10 GbE Switch	EN4091 10 GbE Pass-thru
44W4408	10 GBase-SR SFP+ (MMFiber)	Yes	Yes	Yes
46C3447	10 GBase-SR SFP+ (MMFiber)	Yes	Yes	Yes
90Y9412	IBM SFP+ LR (SMFiber)	Yes	Yes	Yes
81Y1622	1000Base-SX SFP (MMFiber)	Yes	Yes	Yes
81Y1618	1000Base-T SFP	Yes	Yes	Yes
90Y9424	1000Base-LX SFP	Yes	Yes	Yes
49Y7884	IBM QSFP+ 40 Gbase-SR	No	Yes	No
90Y9427	1m IBM Passive DAC SFP+	Yes	Yes	No
90Y9430	3m IBM Passive DAC SFP+	Yes	Yes	No
90Y9433	5m IBM Passive DAC SFP+	Yes	Yes	No
49Y7886	1m 40 Gb QSFP+ to 4 x 10 Gb SFP+ Cable	No	Yes	No
49Y7887	3m 40 Gb QSFP+ to 4 x 10 Gb SFP+ Cable	No	Yes	No
49Y7888	5m 40 Gb QSFP+ to 4 x 10 Gb SFP+ Cable	No	Yes	No
90Y3519	10m IBM MTP Fiber Optical Cable	No	Yes	No
90Y3521	30m IBM MTP Fiber Optical Cable	No	Yes	No
49Y7890	1m QSFP+ to QSFP+ DAC	No	Yes	No

Part number	Description	EN2092 1 GbE Switch	EN4093 10 GbE Switch	EN4091 10 GbE Pass-thru
49Y7891	3m QSFP+ to QSFP+ DAC	No	Yes	No
95Y0323	IBM 1m 10 GBase Copper SFP+ TwinAx (Active)	No	No	Yes
95Y0326	IBM 3m 10 GBase Copper SFP+ TwinAx (Active)	No	No	Yes
95Y0329	IBM 5m 10 GBase Copper SFP+ TwinAx (Active)	No	No	Yes
81Y8295	1m 10 GE Twinax Act Copper SFP+ DAC (active)	No	No	Yes
81Y8296	3m 10 GE Twinax Act Copper SFP+ DAC (active)	No	No	Yes
81Y8297	5m 10 GE Twinax Act Copper SFP+ DAC (active)	No	No	Yes

All Ethernet /O modules are restricted to using the SFP, SFP+, and QSFP modules that are listed in Table 2-6 on page 27. However, OEM Direct Attached Cables can be used if they meet the MSA standards.

2.4 IBM Flex System Ethernet adapters

The IBM Flex System portfolio contains a number of Ethernet I/O adapters. The cards differ in physical port speeds (1 Gbps versus 10 Gbps) and in functions that they support (base Ethernet connectivity versus converged networks and virtual NIC support).

The following Ethernet I/O adapters are covered:

- IBM Flex System CN4054 10Gb Virtual Fabric Adapter
- IBM Flex System EN2024 4-port 1Gb Ethernet Adapter
- IBM Flex System EN4132 2-port 10Gb Ethernet Adapter

2.4.1 IBM Flex System CN4054 10Gb Virtual Fabric Adapter

The IBM Flex System CN4054 10Gb Virtual Fabric Adapter is a 4-port 10 Gb converged network adapter (CNA) for Intel processor-based compute nodes that can scale up to 16 virtual ports and support Ethernet, iSCSI, and FCoE. The adapter supports up to eight virtual NIC (vNIC) devices, where each physical 10 GbE port can be divided into four virtual ports with flexible bandwidth allocation. The CN4054 Virtual Fabric Adapter Upgrade adds FCoE and iSCSI hardware initiator functions.

The CN4054 adapter is shown in Figure 2-24.




Figure 2-24 IBM Flex System CN4054 10Gb Virtual Fabric Adapter

The ordering information for the CN4054 adapter is listed in Table 2-7.

Table 2-7 CN4054 ordering part numbers and descriptions

Part number	Description
90Y3554	IBM Flex System CN4054 10 Gb Virtual Fabric Adapter
90Y3558	IBM Flex System CN4054 Virtual Fabric Adapter Upgrade

The IBM Flex System CN4054 10 Gb Virtual Fabric Adapter includes the following features:

- ▶ Four-port 10 Gb Ethernet adapter
- ▶ Dual-ASIC Emulex BladeEngine 3 (BE3) controller
- ▶ Connection to either 1 Gb or 10 Gb data center infrastructure (1 Gb and 10 Gb autonegotiation)
- ▶ PCI Express 3.0 x8 host interface
- ▶ Full duplex (FDX) capability
- ▶ Bus-mastering support
- ▶ Direct memory access (DMA) support
- ▶ Preboot Execution Environment (PXE) support
- ▶ IPv4/IPv6 TCP, UDP checksum offload:
 - Large send offload (LSO)
 - Large receive offload

- Receive side scaling (RSS)
- IPv4 TCP Chimney Offload
- TCP Segmentation Offload
- ▶ VLAN insertion and extraction
- ▶ Jumbo frames up to 9000 bytes
- ▶ Load balancing and failover support, including:
 - Adapter fault tolerance (AFT)
 - Switch fault tolerance (SFT)
 - Adaptive load balancing (ALB)
 - Teaming support
 - IEEE 802.3ad
- ▶ Enhanced Ethernet (draft):
 - Enhanced Transmission Selection (ETS) (P802.1Qaz)
 - Priority-based Flow Control (PFC) (P802.1Qbb)
 - Data Center Bridging Capabilities eXchange Protocol, CIN-DCBX, and CEE-DCBX (P802.1Qaz)
- ▶ Operates either as a 4-port 1/10 Gb Ethernet adapter or supports up to 16 vNICs
- ▶ In virtual NIC (vNIC) mode, it supports:
 - Virtual port bandwidth allocation in 100 Mbps increments
 - Up to 16 virtual ports per adapter (four per port)
 - With the CN4054 Virtual Fabric Adapter Upgrade, 90Y3558, four of the 16 vNICs (one per port) support iSCSI or FCoE
- ▶ Supports for two vNIC modes: IBM Virtual Fabric Mode and Switch Independent Mode
- ▶ Wake On LAN support
- ▶ With the CN4054 Virtual Fabric Adapter Upgrade, 90Y3558, the adapter adds FCoE and iSCSI hardware initiator support

iSCSI support is implemented as a full offload and presents an iSCSI adapter to the operating system
- ▶ TCP/IP Offload Engine (TOE) support with Windows Server 2003, 2008, and 2008 R2 (TCP Chimney) and Linux:
 - Connection and its state are passed to the TCP offload engine
 - The data transmit and receive function is handled by adapter
 - Supported by iSCSI

For more information, see the *IBM Flex System CN4054 10Gb Virtual Fabric Adapter and EN4054 4-port 10Gb Ethernet Adapter*, TIPS0868, at:

<http://www.redbooks.ibm.com/abstracts/tips0868.html>

2.4.2 IBM Flex System EN2024 4-port 1Gb Ethernet Adapter

The IBM Flex System EN2024 4-port 1Gb Ethernet Adapter is a quad-port Gigabit Ethernet network adapter. When it is combined with the IBM Flex System EN2092 1Gb Ethernet Switch, clients can use an end-to-end 1 Gb solution on the IBM Flex System Enterprise Chassis. The EN2024 adapter is based on the Broadcom 5718 controller, and offers a PCIe 2.0 x1 host interface with MSI/MSI-X. It also supports I/O virtualization features such as VMware NetQueue and Microsoft VMQ technologies.

The EN2024 adapter is shown in Figure 2-25.




Figure 2-25 IBM Flex System EN2024 4-port 1Gb Ethernet Adapter

The ordering information for the EN2024 adapter is listed in Table 2-8.

Table 2-8 EN2024 ordering part number and description

Part number	Description
49Y7900	IBM Flex System EN2024 4-port 1 Gb Ethernet Adapter

The IBM Flex System EN2024 4-port 1 Gb Ethernet Adapter has the following features:

- ▶ Dual Broadcom BCM5718 ASICs
- ▶ Quad-port Gigabit 1000BASE-X interface
- ▶ Two PCI Express 2.0 x1 host interfaces, one per ASIC
- ▶ Full-duplex (FDX) capability, enabling simultaneous transmission and reception of data on the Ethernet network
- ▶ MSI and MSI-X capabilities, up to 17 MSI-X vectors
- ▶ I/O virtualization support for VMware NetQueue, and Microsoft VMQ
- ▶ A total of 17 receive queues and 16 transmit queues
- ▶ A total of 17 MSI-X vectors supporting per-queue interrupt to host
- ▶ Function Level Reset (FLR)

- ▶ ECC error detection and correction on internal SRAM
- ▶ TCP, IP, and UDP checksum offload
- ▶ Large Send offload, TCP segmentation offload
- ▶ Receive-side scaling
- ▶ Virtual LANs (VLANs): IEEE 802.1q VLAN tagging
- ▶ Jumbo frames (9 KB)
- ▶ IEEE 802.3x flow control
- ▶ Statistic gathering (SNMP MIB II, Ethernet-like MIB [IEEE 802.3x, Clause 30])
- ▶ Comprehensive diagnostic and configuration software suite
- ▶ ACPI 1.1a-compliant; multiple power modes
- ▶ Wake-on-LAN (WOL) support
- ▶ Preboot Execution Environment (PXE) support
- ▶ RoHS-compliant

For more information, see the *IBM Flex System EN2024 4-port 1Gb Ethernet Adapter*, TIPS0845, at:

<http://www.redbooks.ibm.com/abstracts/tips0845.html>

2.4.3 IBM Flex System EN4132 2-port 10Gb Ethernet Adapter

The IBM Flex System EN4132 2-port 10Gb Ethernet Adapter provides the highest-performing and most flexible interconnect solution for servers used in enterprise data centers, high-performance computing, and embedded environments.

The IBM Flex System EN4132 2-port 10Gb Ethernet Adapter is shown in Figure 2-26.




Figure 2-26 The EN4132 2-port 10Gb Ethernet Adapter for IBM Flex System

The ordering part number for the EN4132 adapter is listed in Table 2-9.

Table 2-9 Ordering part number and description

Part number	Description
90Y3466	EN4132 2-port 10Gb Ethernet Adapter

The IBM Flex System EN4132 2-port 10Gb Ethernet Adapter has the following features:

- ▶ Based on Mellanox Connect-X3 technology
- ▶ IEEE Std. 802.3 compliant
- ▶ PCI Express 3.0 (1.1 and 2.0 compatible) through an x8 edge connector up to 8 GT/s
- ▶ 10 Gbps Ethernet
- ▶ Processor offload of transport operations
- ▶ CORE-Direct application offload
- ▶ GPUDirect application offload
- ▶ RDMA over Converged Ethernet (RoCE)
- ▶ End-to-end QoS and congestion control
- ▶ Hardware-based I/O virtualization
- ▶ TCP/UDP/IP stateless offload
- ▶ Ethernet encapsulation (EoIB)
- ▶ RoHS-6 compliant

For more information, see the *IBM Flex System EN4132 2-port 10Gb Ethernet Adapter*, TIPS0873, at:

<http://www.redbooks.ibm.com/abstracts/tips0873.html>



Layer 2 Overview

This chapter explains the Layer 2 fundamental networking protocols and terminology used in the rest of this book.

This chapter includes the following sections:

- ▶ Basic Frame Forwarding Concept
- ▶ Virtual local area network (VLAN) and tagging
- ▶ Spanning tree
- ▶ Dynamic Link Aggregation Control Protocol (LACP)
- ▶ Virtual Link Aggregation Groups (VLAG)
- ▶ Cisco Virtual Port Channel (vPC)
- ▶ Link Layer Discovery Protocol (LLDP)
- ▶ Layer 2 Failover

3.1 Basic Frame Forwarding Concept

Each frame contains a source and a destination MAC address. A network Bridge or Switch, also called Layer 2 device, is responsible to transport the Ethernet frame based on the destination MAC address.

Figure 3-1 shows the simplified principle of frame forwarding.

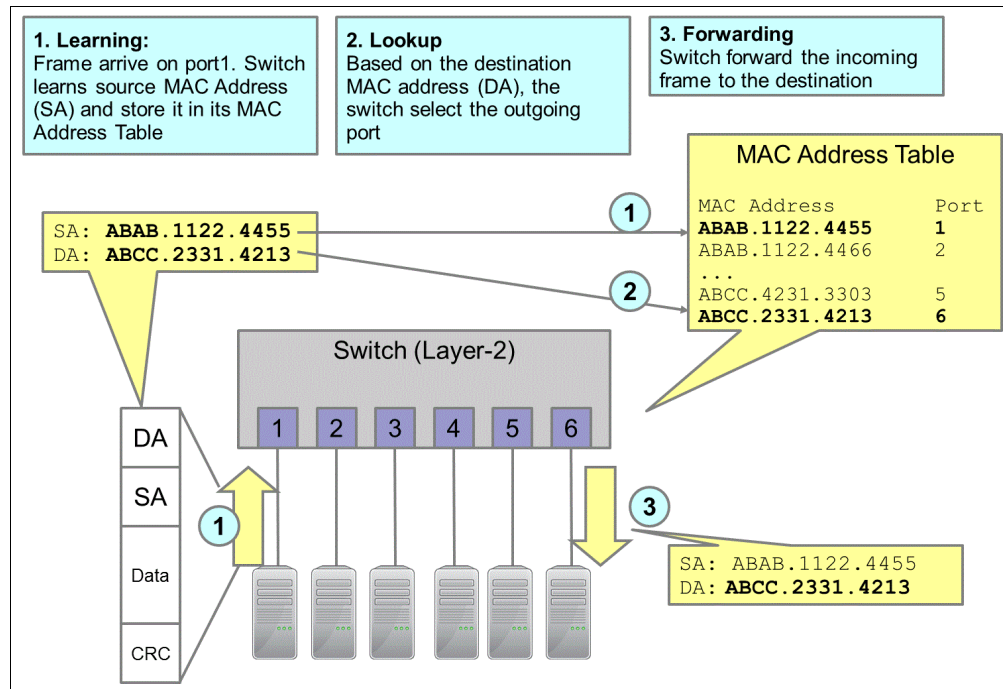


Figure 3-1 Frame forwarding principle

The forwarding of an incoming frame (on port 1 in this case) is divided into these phases:

- Learning** Ethernet Frame arrives on port1. Switch learns source MAC Address (SA) and store it in its MAC Address Table that this address belongs to port 1.
- Lookup** Based on the destination MAC address (DA), the switch performs a lookup in its MAC address table and selects the outgoing port (port 6).
- Forwarding** The switch forwards the Ethernet frame to the destination MAC address through port 6.

If the switch does not know the destination address, it forwards the packet on all ports except the port it was received from.

3.2 Virtual local area network (VLAN) and tagging

A VLAN is a networking concept in which a network is logically divided into smaller virtual LANs so that distinct broadcast domains are created. The Layer 2 traffic in one VLAN is logically isolated from other VLANs as illustrated in Figure 3-2.

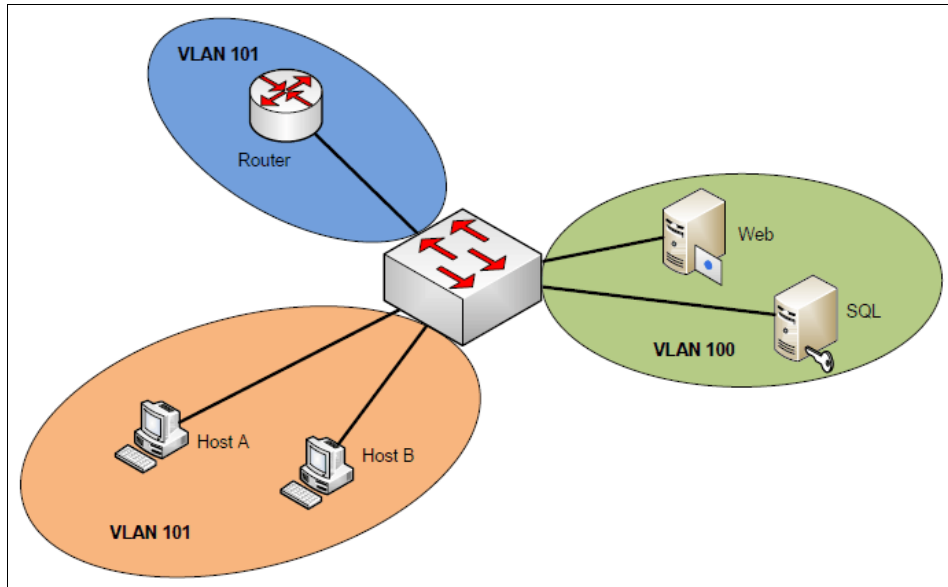


Figure 3-2 Virtual local area network

As shown in Figure 3-3, there are two methods for sharing VLANs across devices:

- ▶ Using dedicated cabling for each VLAN to keep them isolated
- ▶ Marking packets through tagging so that a single interconnect can be used to transport data for multiple VLANs.




Figure 3-3 VLAN tagging

The first method does not scale well because it uses many ports in networks with multiple VLANs and multiple switches. Also, this method does not use link capacity efficiently when traffic in the LANs is not uniform.

The second method is highly scalable because only a single link is required to provide connectivity to many VLANs. This configuration provides for better use of the link capacity when VLAN traffic is not uniform.

3.2.1 Tagged frames

The protocol for VLAN tagging of frames in a LAN environment is defined by the IEEE 802.1P/Q standard. The standard provides an extra 4 bytes of information to be added to each Ethernet frame. A frame that includes this extra information is known as a *tagged frame*.

The 4-byte tag has four component fields:

- ▶ A type field that is 2 bytes long with the hexadecimal value of x8100 to identify the frame as an 802.1P/Q tagged frame.
- ▶ A priority field of 3 bits long to allow a priority value of eight different values to be included in the tag. It is the “P” portion of the 802.1P/Q standard.
- ▶ A Canonical Format Indicator field that is 1 bit long to identify when the contents of the payload field are in canonical format.
- ▶ A VLAN ID field that is 12 bits long to identify which VLAN the frame is a member of, with 4096 different VLANs possible.

3.3 Spanning tree

Because of the history of LANs and Ethernet, there are some shortcomings in the protocol, particularly Ethernet, which was not designed to use frame forwarding. Therefore, the frame format does not include a hop count field (or time to live, TTL) that allows it to detect and discard a looping packet. Packets sent in a loop between multiple switches are forwarded without reaching their destination, which can cause significant load.

The most simple approach to prevent looping packets is to create a network topology where frames with a certain target can take only one path on each individual switch element. For Ethernet the tree topology was chosen, which is the most simple topology that ensures this requirement. Bridges and Switches were enhanced to support a topology configuration protocol called the *Spanning Tree Protocol*.

The Spanning Tree Protocol (STP) provides Layer 2 loop prevention by deactivating redundant routes between network elements. Over the years it has been further enhanced into the following different forms:

- ▶ Spanning Tree Protocol (STP)
- ▶ Rapid STP (RSTP)
- ▶ Multiple STP (MSTP)
- ▶ Per VLAN STP (PVST) and Per VLAN Rapid STP (PVRST)

3.3.1 Spanning Tree Protocol (STP) IEEE802.1D

STP uses Bridge Protocol Data Unit (BPDU) packets to exchange information with other switches. BPDUs send out hello packets at regular intervals to exchange information across bridges and detect loops in a network topology.

Three types of BPDUs are available:

- ▶ Configuration BPDUs: These BPDUs contain configuration information about the transmitting switch and its ports, including switch and port MAC addresses, switch priority, port priority, and port cost.

- ▶ Topology Change Notification (TCN) BPDUs: When a bridge must signal a topology change, it starts to send TCNs on its root port. The designated bridge receives the TCN, acknowledges it, and generates another one for its own root port. The process continues until the TCN reaches the root bridge.
- ▶ Topology Change Notification Acknowledgement (TCA) BPDUs: These frames are sent by the root bridge to acknowledge the receipt of a TCN BPDU.

STP uses the information that is provided by the BPDUs to perform these tasks:

- ▶ Select a root bridge
- ▶ Identify root ports for each switch
- ▶ Identify designated ports for each physical LAN segment
- ▶ Prune specific redundant links to create a loop-free tree topology

All leaf devices calculate the best path to the root device and place their ports in blocking or forwarding states based on the best path to the root. The resulting tree topology provides a single active Layer 2 data path between any two end stations.

Requirement: The root bridge election is an extremely important point in a network design. To avoid suboptimal Layer 2 paths, it is always necessary to manually adjust the bridge priority on each switch in a Layer 2 network.

3.3.2 Rapid Spanning Tree (RSTP) IEEE802.1w

RSTP provides better reconvergence time than the original STP. RSTP identifies certain links as point to point. When a point-to-point link fails, the alternate link can make the transition to the forwarding state.

RSTP adds new bridge port roles to speed convergence after a link failure. The RSTP bridge ports can have these roles:

Root port	The “best path” to the root device.
Designated port	Indicates that the switch is the designated bridge for the other switch that connects to this port.
Alternate port	Provides an alternate root port.
Backup port	Provides an alternate designated port.

RSTP was originally defined in the IEEE 802.1w draft specification, and later incorporated into the IEEE 802.1D-2004 specification.

3.3.3 Multiple Spanning Tree (MSTP) IEEE802.1s

Although RSTP provides faster convergence time than STP, it still does not solve a problem inherent in STP. All VLANs within a LAN must share this spanning topology, while many links in the network can be unused. To solve this problem, the existing STP concepts are no longer applied to physical ports, but to the connectivity of multiple individual groups of VLANs, called spanning-tree regions.

In an MSTP region, a group of bridges can be modeled as a single bridge. An MSTP region contains multiple spanning tree instances (MSTIs). MSTIs provide different paths for different VLANs. This function facilitates better load sharing across redundant links.

An MSTP region can support up to 64 MSTIs, and each instance can support anywhere from 1 - 4094 VLANs.

MSTP was originally defined in the IEEE 802.1s draft specification, and later incorporated into the IEEE 802.1Q-2005 specification.

3.3.4 Per VLAN Rapid Spanning Tree (PVRST)

PVRST is a nonstandard spanning tree extension and based on RSTP introduced by Cisco Systems. In PVRST mode, each VLAN is assigned to an own spanning-tree group.

Like RSTP, PVRST mode provides rapid Spanning Tree convergence. Each VLAN has its own Spanning-Tree instance and tree which allows to utilize different paths.

PVRST use 802.1Q tagged frames to differentiate STP BPDUs for each VLAN. The IBM System Networking implementation of PVRST is fully compatible to Cisco RSTP/PVRST+ protocol. A maximum of 127 spanning-tree groups are currently allowed in IBM System Networking switches.

3.4 Dynamic Link Aggregation Control Protocol (LACP)

LACP is a vendor independent standard for dynamically building aggregated links between switches, and was first defined in 802.3ad. The standard was later included in the mainline 802.3 standard, but then was pulled out into the current standard 802.1AX-2008. LACP is a dynamic way of determining whether both sides of the link might be aggregating.

As shown in Figure 3-4, link aggregation combines multiple physical links to operate as a single larger logical link. The member links no longer function as independent physical connections, but as members of the larger logical link.




Figure 3-4 Link aggregation

Link aggregation provides greater bandwidth between the devices at each end of the aggregated link. Another advantage of link aggregation is increased availability because the aggregated link is composed of multiple member links. If one member link fails, the aggregated link continues to carry traffic over the remaining member links.

Each of devices that are interconnected by the aggregated link uses a hashing algorithm to determine on which of the member links frames to transmit on. The hashing algorithm might use varying information in the frame to make the decision. This algorithm might include a

source MAC, destination MAC, source IP, destination IP, and more. It might also include a combination of these values.

3.5 Virtual Link Aggregation Groups (VLAG)

VLAG is an extension to link aggregation to allow more redundancy. For a standard LAG (static or dynamic) all ports that are building an aggregated link must be on the same switch. VLAG allows two switches to appear as a single virtual entity to build an aggregated link that is distributed to both switches. From the perspective of the target device, the ports that are connected to the VLAG peers appear to be a single VLAN tagging link connecting to a single logical device.

As shown in Figure 3-5, a switch in the access layer can be connected to more than one switch in the aggregation layer to provide for network redundancy. Typically, STP is used to prevent broadcast loops, blocking redundant uplink paths. This configuration has the unwanted consequence of reducing the available bandwidth between the layers. In addition, STP might be slow to resolve topology changes that occur during a link failure, and can result in considerable MAC address flooding.

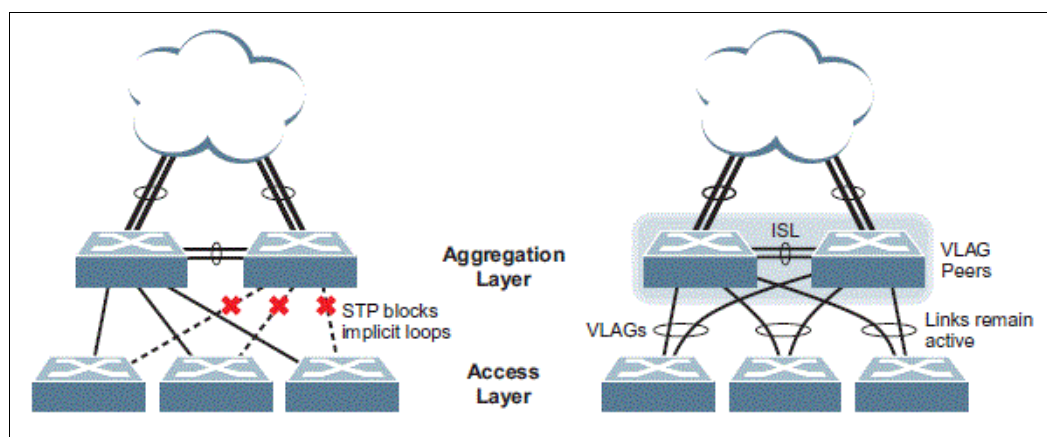


Figure 3-5 Spanning-tree versus VLAG

The VLAG-capable switches synchronize their logical view of the access layer port structure and internally prevent implicit loops. The VLAG topology also responds more quickly to link failure, and does not result in unnecessary MAC flooding.

3.6 Cisco Virtual Port Channel (vPC)

On the Nexus Platform, Cisco implemented the vLAG concept as a version of a Multichassis Etherchannel (MEC), called the vPC. The vPC combines the advantages of hardware redundancy and the loop management of an aggregated link. The pair of switches that form the vPC appear to any Portchannel-attached device as a single switch from Layer 2 perspective, although they operate as two independent devices with independent switch control and management, as shown in Figure 3-6 on page 42

If using a vPC, the STP is no longer needed to manage the loops, it can be disabled on these links and all disadvantages of it can be eliminated. The major advantages are the usability of all bandwidth of the installed links and the fast handling of link failures within the vPC.




Figure 3-6 Schematic drawing of vPC

The pair of switches that form the vPC are seen as a single switch from the device connected to the Port channel. This device can be either a server, a switch, or any other network device

3.7 Link Layer Discovery Protocol (LLDP)

LLDP is a vendor independent protocol for network devices to advertise information about their identity and capabilities. It is referred to as *Station and Media Access Control Connectivity Discovery*, which is specified in the 802.1ab standard.

LLDP performs functions similar to several proprietary protocols, such as the Cisco Discovery Protocol (CDP) and others.

LLDP-capable devices transmit information in Type Length Values (TLV) messages to neighbor devices at fixed intervals. Device information can include specifics such as chassis and port identification, system name, and system capabilities.

Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) is an enhancement of LLDP. Network devices can learn and distribute other information such as auto-discovery of LAN policies, inventory management, and so on.

With this information, the switch can quickly identify devices, resulting in a LAN that interoperates smoothly and efficiently.

3.8 Layer 2 Failover

The primary application for Layer 2 Failover is to support Network Adapter Teaming. With Network Adapter Teaming, two or more NICs on each server share an IP address, and are configured into a team. One NIC is the Active link, and the other is a Standby link. They can also be configured in an Active-Active pairing, in which both NICs are in a forwarding state. For more information, see the documentation for your Ethernet adapter.

Layer 2 Failover can be enabled on any link aggregation group in IBM System Networking switches, including LACP. Aggregated links can be added to failover trigger groups. Then, if

some specified number of monitor links fail, the switch disables all the control ports in the switch. When the control ports are disabled, it causes the NIC team on the affected servers to fail over from the primary to the backup NIC. This process is called a failover event.

When the appropriate number of links in a monitor group return to service, the switch enables the control ports. This configuration causes the NIC team on the affected servers to fail back to the primary switch (unless Auto-Fallback is disabled on the NIC team). The backup switch processes traffic until the primary switch's control links come up, which can take up to 5 seconds.

Figure 3-7 is a simple example of Layer 2 Failover. One switch is the primary, and the other is used as a backup. In this example, all ports on the primary switch belong to a single LAG, with Layer 2 Failover enabled, and the Failover Limit set to 2. If two or fewer links in Trigger 1 remain active, the switch temporarily disables all control ports. This action causes a failover event on Server 1 and Server 2 NIC 1.

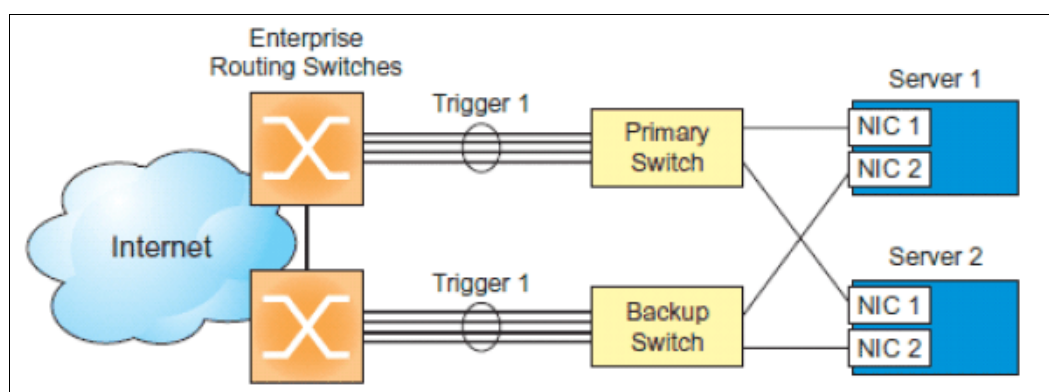


Figure 3-7 Basic Layer 2 Failover

This feature is also referred to as *Uplink Failure Detection*. The switch constantly monitors the port or LAG to the Core Network. When a failure is detected, the switch disables the pre-configured ports that are connected to the servers.



Layer 3 Overview

This chapter explains the Layer 3 fundamental networking protocols and terminology that are used in the rest of this book.

This chapter includes the following sections:

- ▶ Overview of Layer 3
- ▶ Static routes
- ▶ Default gateways
- ▶ Equal-cost multi-path (ECMP) static routes
- ▶ Routing Information Protocol v2 (RIPv2)
- ▶ Enhanced Interior Gateway Routing Protocol (EIGRP)
- ▶ Open Shortest Path First (OSPF) for IPv4
- ▶ Border Gateway Protocol (BGP)
- ▶ IPv6
- ▶ Open Shortest Path First for IPv6 (OSPFv3)
- ▶ FHRP (First Hop Redundancy Protocols) VRRP and HSRP

4.1 Overview of Layer 3

Without Layer 3 IP routing at the switch level, all cross-subnet traffic is relayed to the default gateway. That gateway provides the necessary IP address information and sends the data back down to the Layer 2 switch. Placing Layer 3 IP routing on the switch allows for cross-subnet traffic within the switch, freeing up the upstream router to handle just in-bound and out-bound traffic. Because IBM System networking switches use ASICs for forwarding Layer 3 packets, cross-subnet traffic can be routed within the switch at wirespeed Layer 2 performance rates. This configuration eases the load on the local router, and saves the network administrator from having to reconfigure each endpoint with new IP addresses. It is also achieved without any loss of performance.

4.2 Static routes

It is possible to manually configure static routes to forward IP packets. The entry specifies a network and the IP address of the gateway router, or the next “hop” in the network.

4.3 Default gateways

IBM System Networking switches can be configured with up to four IPv4 gateways:

- ▶ Gateway 1: data traffic
- ▶ Gateway 2: data traffic
- ▶ Gateway 3: management traffic for interface 127
- ▶ Gateway 4: management traffic for interface 128

It is possible to assign different gateway destinations to different VLANs on the switch. Using multiple gateways for the same IP address route is also used to configure ECMP.

4.4 Equal-cost multi-path (ECMP) static routes

ECMP is a forwarding mechanism that can be used to equally distribute load across multiple paths. ECMP is configured by assigning multiple gateways to the same IP route. ECMP routes allow the switch to choose between several next hops toward a destination. The switch runs periodic health checks (ping) on each ECMP gateway. If a gateway fails, it is removed from the routing table.

4.5 Routing Information Protocol v2 (RIPv2)

The goal of any routing protocol is to populate a devices routing table with valid, loop-free routes. Routing protocols have become essential in large and complex networks. RIPv2-enabled routers share and track available routes. RIPv2 is a distance vector protocol. Routers that use distance vector protocols do not know the entire path to a destination, or the topology of a network. Instead, they just have information as to which port to use and the distance away it is. Routers must synchronize (converge) their routing tables at regular intervals to prevent loops from occurring. RIPv2 has the advantage of being easy to configure. However, because routers must share their entire routing tables regularly, RIPv2 does not scale well. Both System Networking switches and Cisco switches support RIPv2.

Cisco Nexus switches support RIPv2 without the need for a L3 license. RIPv2 is Internet Standard STD56, RFC 2453 (<http://www.ietf.org/rfc/rfc2453>, 1998).

4.6 Enhanced Interior Gateway Routing Protocol (EIGRP)

EIGRP is a Cisco-proprietary advanced distance vector routing protocol. It has optimizations to minimize routing instability that is incurred during topology changes and to maximize bandwidth and processing power of the router. The Diffusing Update Algorithm (DUAL) ensures loop-free operation and provides for fast convergence. Unlike other distance vector routing protocols, EIGRP does not synchronize its tables with periodic updates. Instead, EIGRP exchanges full routing tables only when it is establishing new neighbors. After neighbor tables are shared, only updates get exchanged. EIGRP uses three tables:

- ▶ Neighbor table, which contains details of directly connected routers.
- ▶ Topology table, which contains an aggregation of the routing tables from all directly connected neighbor routers. It contains a list of destination networks with their associated metrics. Every destination in the topology table has a successor route (fastest route) and a feasible successor route (next fastest route) identified and stored in the table. Each route is designated as either passive, meaning that the route state is stable, or active meaning that the router is actively updating details of the route and should not be used.
- ▶ Routing table, which is populated by the successor and feasible successors if identified.

System networking switches do not support EIGRP. Any connected Cisco devices that share EIGRP routes must redistribute them using a supported routing protocol like RIPv2 or OSPF.

4.7 Open Shortest Path First (OSPF) for IPv4

OSPF is the most widely used interior routing protocol in large enterprise networks. It is defined as OSPF version 2 for IPv4 in RFC 2328 (<http://www.ietf.org/rfc/rfc2328>, 1998). OSPF networks scale well because they can be logically divided into routing areas. Each area is identified by a 32-bit number expressed as a decimal or often as an octet dotted decimal number similar to an IP address. By convention area 0 (or 0.0.0.0) represents the core or backbone region. Each additional area must be connected to area 0.0.0.0.

4.7.1 OSPF area types

Areas inject summary routing information into the backbone, which then distributes it to other areas as needed. OSPF defines the following types of areas (shown in Figure 4-1 on page 48):

- ▶ Stub area: An area that is connected to only one other area. External route information is not distributed into stub areas.
- ▶ Not-So-Stubby-Area (NSSA): Similar to a stub area but with more capabilities. Routes originating from within the NSSA can be propagated to adjacent transit and backbone areas. External routes from outside the stub area can be advertised within the NSSA, but are not distributed into other areas.
- ▶ Transit Area: An area that allows area summary information to be exchanged between routing devices. The backbone (area 0), any area that contains a virtual link to connect two areas, and any area that is not a stub area or an NSSA are considered transit areas.

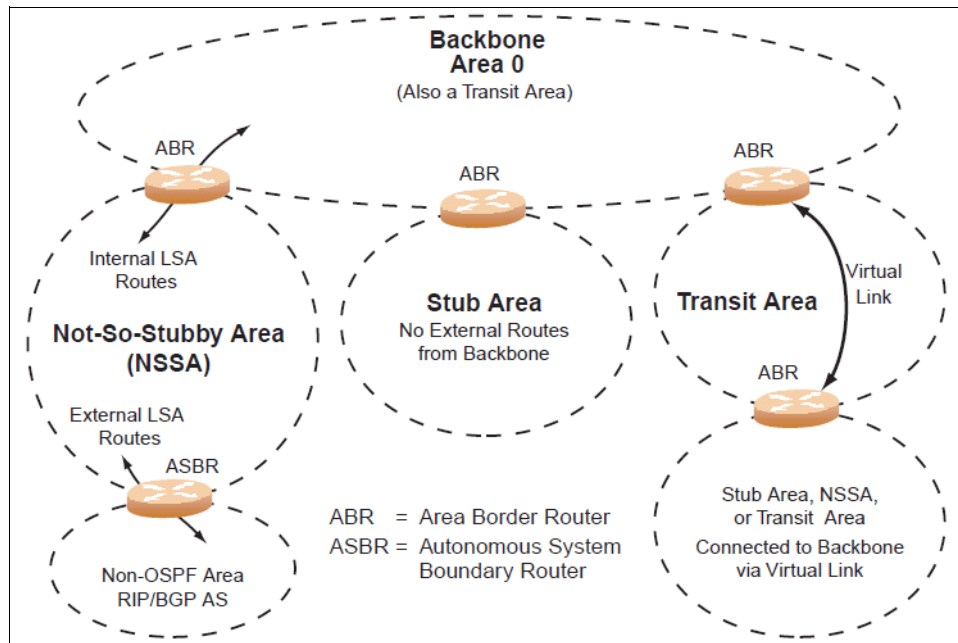


Figure 4-1 OSPF area types

4.7.2 Neighbors and adjacencies

In areas with two or more routing devices, *neighbors* and *adjacencies* are formed. Neighbors are routing devices that maintain information about each others' health. To establish neighbor relationships, routing devices periodically send hello packets out of each of their interfaces. All routing devices that share a common network segment appear in the same area, and have the same health parameters (*hello* and *dead* intervals), authentication parameters, area number, and area stub-flag respond to each other's hello packets and become neighbors.

Neighbors continue to send periodic hello packets to advertise their health to neighbors. In turn, they listen to hello packets to determine the health of their neighbors and to establish contact with new neighbors. On broadcast networks (like Ethernet), the hello process is used for electing one of the neighbors as the area's designated router (DR) and one as the area's backup designated router (BDR). The DR is next to all other neighbors and acts as the central contact for database exchanges. Each neighbor sends its database information to the DR, which relays the information to the other neighbors.

The BDR is next to all other neighbors (including the DR). Each neighbor sends its database information to the BDR as with the DR, but the BDR merely stores this data and does not distribute it. If the DR fails, the BDR takes over the task of distributing database information to the other neighbors.

4.7.3 Link State Database (LSDB)

OSPF is a link-state routing protocol. A *link* represents an interface (or routable path) from the routing device. By establishing an adjacency with the DR, each routing device in an OSPF area maintains an identical LSDB describing the network topology for its area.

Each routing device transmits a link-state advertisement (LSA) on each of its *active* interfaces. LSAs are entered into the LSDB of each routing device. OSPF uses flooding to distribute LSAs between routing devices. Interfaces can also be *passive*. Passive interfaces

send LSAs to active interfaces, but do not receive LSAs, hello packets, or any other OSPF protocol information from active interfaces. Passive interfaces behave as stub networks. They allow OSPF routing devices to be aware of devices that otherwise participate in OSPF (either because they do not support it, or because the administrator chooses to restrict OSPF traffic exchange or transit).

When LSAs result in changes to the routing device's LSDB, the routing device forwards the changes to the adjacent neighbors (the DR and BDR) for distribution to the other neighbors.

OSPF routing updates occur only when changes occur, instead of periodically. For each new route, if an adjacent neighbor is interested in that route, an update message that contains the new route is sent to the neighbor. For each route removed from the route table, if the route is already sent to an adjacent neighbor, an update message that contains the route to withdraw is sent.

4.7.4 OSPF router types

As shown in Figure 4-2, OSPF uses the following types of routing devices:

- ▶ Internal router (IR): A router that has all of its interfaces within the same area. IRs maintain LSDBs identical to the LSDBs of other routing devices within the local area.
- ▶ Area border router (ABR): A router that has interfaces in multiple areas. ABRs maintain one LSDB for each connected area and disseminate routing information between areas.
- ▶ Autonomous system boundary router (ASBR): A router that acts as a gateway between the OSPF domain and non-OSPF domains, such as RIP, BGP, and static routes.

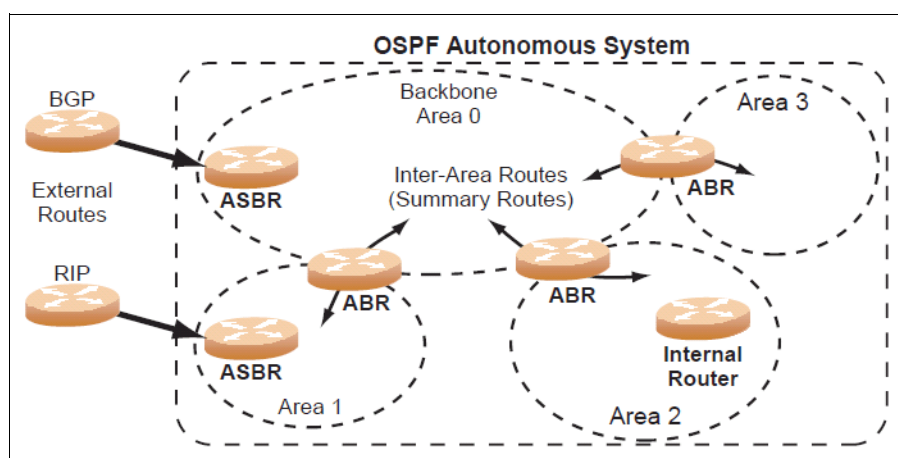


Figure 4-2 OSPF router types

4.7.5 Shortest path first

The routing devices use a link-state algorithm (Dijkstra's algorithm) to calculate the shortest path to all known destinations. This path is based on the cumulative cost that is required to reach the destination. The cost of an individual interface in OSPF is an indication of the processing that is required to send packets across it. The cost is inversely proportional to the bandwidth of the interface. A lower cost indicates a higher bandwidth.

4.8 Border Gateway Protocol (BGP)

BGP version 4 is an open standard. After several iterations, it is now published as RFC 4271 (<http://www.ietf.org/rfc/rfc4271.txt>, 2006). BGP is the lifeblood of the internet. It exchanges routing information between all the major Internet Service Providers (ISPs). It is an Exterior Gateway Protocol (EGP), which means it exchanges routing information between autonomous systems (ASs). This is different from Interior Gateway Protocols (IGP) such as RIPv2, EIGRP and OSPF, which support routing within an AS.

An AS is defined as “a connected group of one or more IP prefixes run by one or more network operators which has a single and clearly defined routing policy”. This in real terms tends to be an ISP network together with all of its downstream customer networks. BGP routers talk to one another over a permanent TCP connection on port 179. BGP communication between two routers within the same AS is called Interior BGP (iBGP), and between two ASs it is called Exterior BGP (eBGP). On smaller networks, BGP routers within an AS must form a complete mesh with each other. BGP requires that every AS has a 16-bit Autonomous System Number (ASN). ASNs can have values from 0 - 65535. RFC 4893 introduced 32-bit AS numbers, which IANA (<http://www.iana.org/>) has begun to allocate. The ASN is a globally unique identifier. BGP keeps a list of every AS (ASN) that a path passes through. This enables the router to eliminate paths with loops by deleting those that have the same ASN more than once. Unlike IGPs, BGP does not support multipath routing by default. If there are two or more paths to a destination, BGP ensures only one is actually used. There is a list of weighted steps that are used to determine which routes are preferred and which routes are removed.

An iBGP is a type of internal routing protocol you can use to do active routing inside your network. It also carries AS path information, which is important when your system is an ISP or doing BGP transit. The iBGP peers must maintain reciprocal sessions to every other iBGP router in the same AS (in a full-mesh manner) to propagate route information throughout the AS.

If the iBGP session shown between the two routers in AS 20 is not present (Figure 4-3), the top router does not learn the route to AS 50, and the bottom router does not learn the route to AS 11. This occurs even though the two AS 20 routers are connected through the IBM System Networking switch.




Figure 4-3 Diagram showing the importance of iBGP

Typically, an AS has one or more *border routers*, which are peer routers that exchange routes with other ASs, and an internal routing scheme that enables routers in that AS to reach every other router and destination within that AS. When you advertise routes to border routers on other autonomous systems, you are committing to carry data to the IPv4 space represented

in the route that is advertised. For example, if you advertise 192.204.4.0/24, you are declaring that if another router sends you data destined for any address in 192.204.4.0/24, you know how to carry that data to its destination.

4.9 IPv6

The IPv6 protocol is an RFC standard 2460 (<http://www.ietf.org/rfc/rfc2460>, 1998). The IPv4 protocol was developed back in the early 1980s. IPv6 has several improvements over IPv4 and resolved some unforeseen issues.

4.9.1 Address size

The obvious place to start is the size of the address space. Table 4-1 compares the absolute address spaces of IPv4 and IPv6. In reality, for both IPv4 and IPv6, not all addresses are available for host allocation or routing.

Table 4-1 Relative address space of IPv4 and IPv6

Protocol	Bits available for addresses	Absolute address space
IPv4	2^{32}	4,294,967,296
IPv6	2^{128}	340,282,366,920,938,463,374,607,431,768,211,456

4.9.2 Address usage

IPv6 design aim was not just to create a massive pool of IP addresses, but to also provide systematic, hierarchical allocation of addresses, and efficient route aggregation. A typical global IPv6 address looks like this: 2001:15f8:106:208:202:55ff:fe54:af3a. The first 64 bits describe the subnet identifier: 2001:15f8:106:208::/64. The last 64 bits are the host ID 202:55ff:fe54:af3a, which is usually derived from Layer 2 MAC address. The first 3 bits of an IPv6 address are reserved to define the type of IPv6 address used.

4.9.3 Address hierarchy

IPv6 has a hierarchy of address block allocation. The 64 bits used for the Subnet identifier are broken down further as follows:

- ▶ Top-Level Aggregation Identifier (TLA ID, 13 bits) assigned to major service providers.
- ▶ Next-Level Aggregation Identifier (NLA ID, 24 bits) assigned to minor service providers.
- ▶ Site-Level Aggregation Identifier (SLA ID, 16 bits) assigned to organizations/companies. The 16 bits provide 645,535 subnets.

This configuration is ideal for routing performance/management because core routers only must route based on the TLA ID and so on.

4.9.4 Address autoconfiguration/plug-and-play

When a host is enabled for IPv6, it automatically creates a tentative link-local address. When the host is connected to a port on a subnet, it confirms the uniqueness of the address by using a ping. Using the example above, the link-local address is fe80::202:55ff:fe54:af3a/64. All devices on a subnet can communicate by using their unique link-local addresses. The

local router can publish a global IPv6 prefix and a default route (to itself) to any hosts on the same subnet. This system has some advantages over IPv4 systems:

- ▶ A host always has the same unique global and link-local IPv6 address.
- ▶ Address allocation is built into the system, with no separate DHCP server issues.
- ▶ Both host and server devices can have their addresses auto configured.
- ▶ Router advertisements that are received at the same time also gives the host its default route.
- ▶ Changing IP addresses is achieved at the router with no loss of connectivity.
- ▶ In the absence of an IPv6 enabled router, link-local addresses can be used to communicate across a single LAN, for example, for printing.

4.10 Open Shortest Path First for IPv6 (OSPFv3)

OSPFv3 works in a similar way to OSPFv2. It is defined in RFC standard 5340 (<http://tools.ietf.org/html/rfc5340>, 2008). The two protocols are not compatible, however, because OSPFv3 is dedicated to sharing IPv6 routes whereas OSPFv2 is purely for IPv4. Both protocols can run together on the same device. Neighbor adjacencies in OSPFv3 are established and maintained by using the link-local addresses, and not configured IPv4 addresses. Although OSPFv3 deals entirely with IPv6 addresses, you must define an arbitrary 32-bit router-id expressed in dotted decimal (IPv4) format. Neither Cisco Nexus or IBM system networking switches currently support the other main IPv6 routing protocol RIPng.

4.11 FHRP (First Hop Redundancy Protocols) VRRP and HSRP

In a high-availability network topology, no device can create a single point of failure for the network or force a single point-of-failure to any other part of the network. This situation means that your network remains in service despite the failure of any single device. Achieving this goal usually requires redundancy for all vital network components.

FHRP protocols are designed to protect the default gateway address that is used on a subnetwork by allowing two or more routers to provide backup for that address. During a failure of the active router, a backup router takes over the function of that address. FHRP protocols have also been applied to other services that require redundancy for a single IP address. Virtual Router Redundancy Protocol (VRRP) and Hot Standby Router Protocol (HSRP) are similar FHRP protocols. VRRP is an open standard RFC 2338 (<http://www.ietf.org/rfc/rfc2281.txt>, 1998). HSRP is a Cisco innovation that was introduced in 1998 and described in RFC 2281 (<http://www.ietf.org/rfc/rfc2281>, 1998).

Table 4-2 shows the FHRP.

Table 4-2 Protocols

Protocol	HSRP	VRRP
Availability	Cisco only (including Nexus)	Open standard Cisco Nexus and IBM System Networking
RFC	2281	3768

Protocol	HSRP	VRRP
Virtual IP address	Separate from physical address	Can use same IP address as physical port
master/backups	One master, all others backup	One master, one standby, all others listening
hello packets	multicast ip 224.0.0.2 (v1) multicast ip 224.0.0.102 (v2)	multicast ip 224.0.0.18
MAC address	00-00-0C-07-AC-XX, where XX is the virtual group ID	00-00-5E-00-01-XX, where XX is the virtual group ID
Failback configuration	Use preempt command to force master to own Virtual IP address when available	Master claims Virtual IP address when available

For both HSRP and VRRP, the principles are the same. The virtual router consists of a user-configured virtual router identifier (VRID) and an IPv4 address. The VRID is used to build a virtual router MAC Address. The five highest-order octets of the virtual router MAC address are provided by the standard MAC prefix (either 00-00-5E-00-01 or 00-00-0C-07-AC). The lowest order octet is formed from the VRID.

One of the physical routers is elected as the virtual router master, based on a number of priority criteria. This master router assumes control of the virtual router IPv4 address. The master router forwards packets that are sent to the virtual router and responds to Address Resolution Protocol (ARP) requests. The master also sends out periodic advertisements to inform backup routers that it is alive and what its priority is.

If the master fails, one of the backup routers takes control of the virtual router IPv4 address and actively processes traffic addressed to it. Because the backup router uses the same MAC address, hosts do not need to send ARP requests and packets are processed with a minimum of disruption.

Configuration tip: Generally, configure all HSRP or VRRP options, such as priority, preempt, and authentication, before configuring the virtual IP address. Doing so minimizes disruption and state changes in the network.

A backup router can stop receiving advertisements for one of two reasons: The master is down, or all communications links between the master and the backup are down. If the master fails, the preferred solution is for the backup (or one of the backups, if there is more than one) becomes the master.

Two masters: If the master is working correctly but communication between the master and the backup fails, there can be two masters within the virtual router. To prevent this situation from happening, configure redundant links to be used between the switches that form a virtual router.

4.11.1 Active-active redundancy

In an active-active configuration as shown Figure 4-4, two switches provide redundancy for each other, with both active at the same time. Each switch processes traffic on a different subnet. When a failure occurs, the remaining switch can process traffic on all subnets.




Figure 4-4 Diagram demonstrating an active-active VRRP redundancy setup

Although this example shows only two switches, there is no limit of the number of switches that can be used in a redundant configuration. It is possible to implement an active-active configuration across all the VRRP-capable switches in a LAN. Each VRRP-capable switch in an active-active configuration is autonomous. Switches in a virtual router do not need to be identically configured.

4.11.2 VRRP high availability with VLAGs

VRRP can be used with VLAGs and LACP-capable servers and switches to provide seamless redundancy as shown in Figure 4-5.

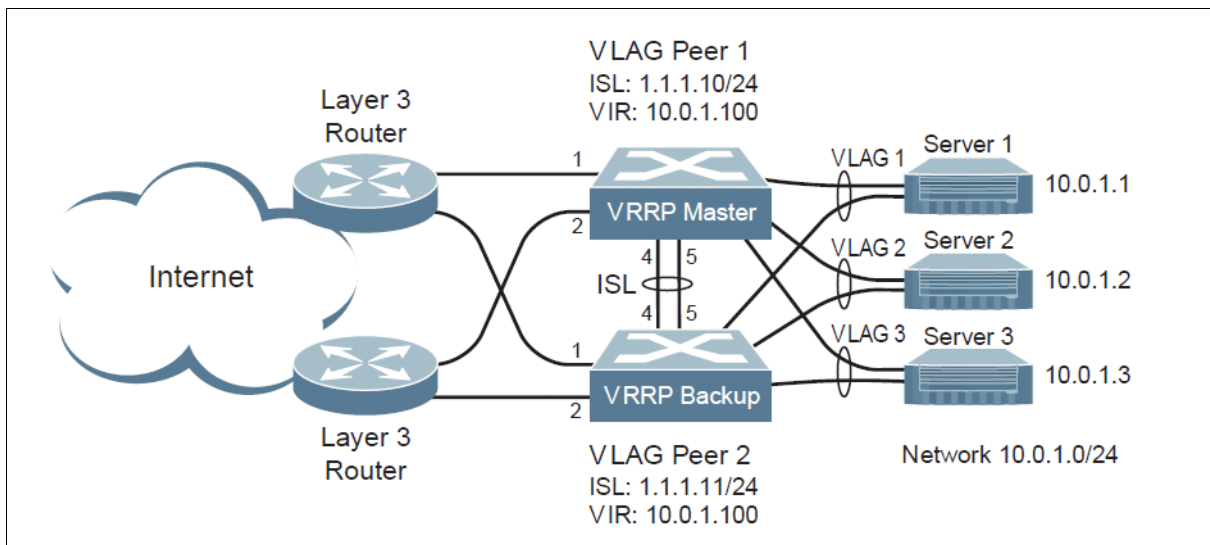


Figure 4-5 Active-active configuration using VRRP and VLAGs



Interoperability Use Cases: Connecting to a Cisco Network

This chapter provides industry-standard interoperability use cases with an upstream Cisco network. These use cases are useful whether you are getting ready to deploy and use the embedded switches in the back of the IBM Flex System Enterprise Chassis, the embedded switches in the back of the IBM BladeCenter chassis, or RackSwitches from the IBM System Networking portfolio.

This chapter includes the following sections:

- ▶ Introduction
- ▶ High availability overview
- ▶ Fully redundant with virtualized chassis technology (VSS/vPC/vLAG)
- ▶ Fully redundant with traditional spanning-tree
- ▶ Fully redundant with Open Shortest Path First (OSPF)

5.1 Introduction

The use cases described in this chapter were selected primarily based on input from IBM System Networking Consulting Engineers. They are configurations that have been observed most often in the field during customer engagements.

Note: Although these implementation scenarios have been tested and verified to be compatible with an upstream Cisco network in a lab environment, these are not the only design options available to the network architect. Use them as general guidance only. Consult with your IBM Account Representative to engage the Worldwide System Networking Consulting Engineers for more in-depth design discussion if a unique topology is required.

Before describing the scenarios, this chapter describes traditional, highly available network implementations. It describes their unique characteristics as a background as to why the scenarios are recommended.

5.2 High availability overview

Customers often require continuous access to their network-based resources and applications. Providing high availability (HA) for client network resources can be a complex task that involves fitting multiple pieces together on a hardware and software level. The focus is to provide high availability access to the network infrastructure.

Network infrastructure availability can be achieved by using various techniques and technologies. Most are widely used standards, and can be deployed with everything from rack-mount servers to full iDataPlex racks. However, some are specific to the IBM Flex System Enterprise Chassis. This section reviews the most common technologies that can be implemented in an Enterprise Chassis environment to provide high availability to the network infrastructure.

A typical LAN infrastructure consists of server NICs, client NICs, and network devices, such as Ethernet switches and cables, that connect them. Specific to the Enterprise Chassis, the potential failure areas for node network access include port failures (both on switches and the node adapters), the midplane, and the I/O modules.

The first step in achieving high availability is to provide physical redundancy of components that are connected to the infrastructure as a whole. Providing this redundancy typically means that the following measures are taken:

- ▶ Deploy node NICs in pairs
- ▶ Deploy top of rack switches or embedded switch modules in pairs
- ▶ Connect the pair of node NICs to separate I/O modules in the Enterprise Chassis
- ▶ Provide connections from each I/O module to a redundant upstream infrastructure

After physical redundancy requirements are met, consider the logical elements to use this physical redundancy. The following logical features aid in high availability:

- ▶ NIC teaming/bonding on the server or compute node
- ▶ Layer 2 (L2) failover (also known as *trunk failover*) on the I/O modules
- ▶ Rapid Spanning Tree Protocol for looped environments
- ▶ Virtual Link Aggregation on upstream devices connected to the I/O modules

- ▶ Virtual Router Redundancy Protocol for redundant upstream default gateway
- ▶ Routing Protocols (such as RIP or OSPF) on the I/O modules, if L2 adjacency is not a requirement

5.2.1 Looped and blocking design

One of the most traditional designs for chassis HA server-based deployments is the looped and blocking design as shown in Figure 5-1.

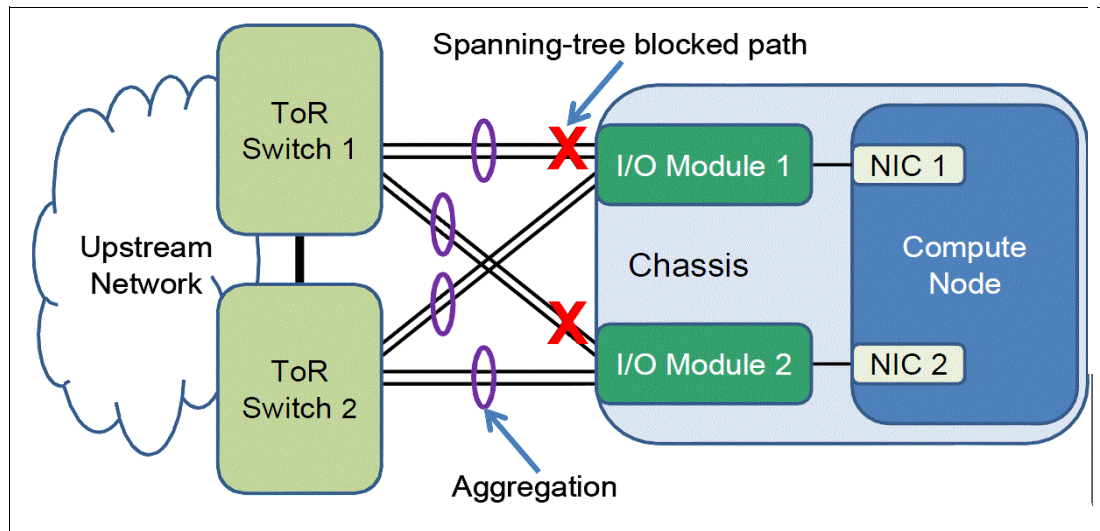


Figure 5-1 Looped and blocking design, no host NIC teaming

The looped and blocking design shows each I/O module in the Enterprise Chassis with two direct aggregations to a pair of upstream Top-of-Rack (ToR) switches. The specific number and speed of the external ports that are used for link aggregation depend on your redundancy and bandwidth requirements. This topology is a bit complicated, and is suggested for environments in which hosts need network redundancy, but they are not themselves running any NIC teaming. Although this choice offers complete network-level redundancy out of the chassis, the potential exists to lose half of the available links and bandwidth because of the Spanning Tree Protocol (STP) blocking them.

Important: Because of possible issues with looped designs in general, use loop-free topologies if you can still offer hosts the high availability access necessary to function.

5.2.2 Non-looped, single upstream device design

An alternative to the looped and blocking design in Figure 5-1 on page 57 is the non-looped, single upstream device HA design as shown in Figure 5-2.

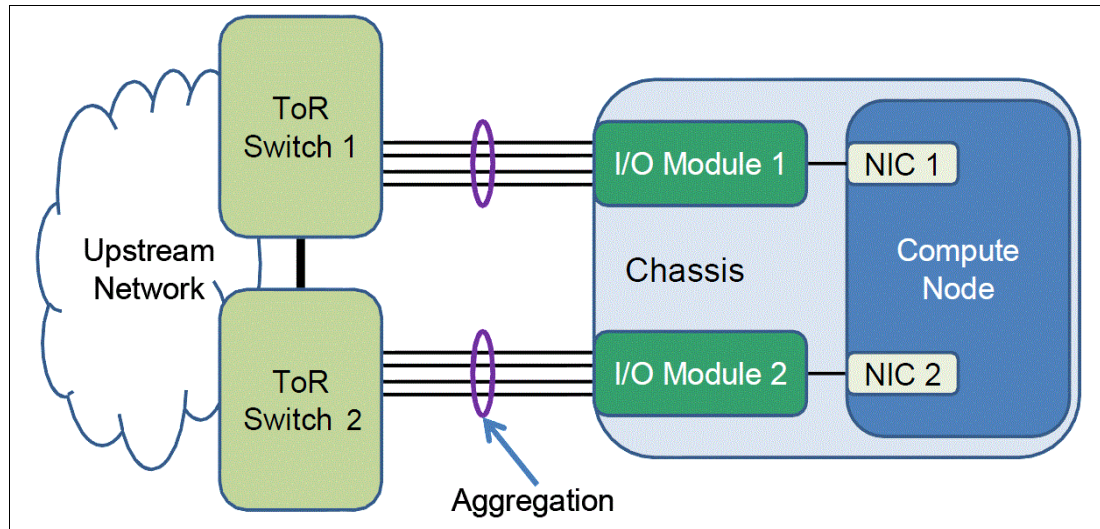


Figure 5-2 Non-looped, single upstream device design, with host NIC teaming

Figure 5-2 shows each I/O module in the Enterprise Chassis directly connected to a single ToR switch through aggregated links. This topology is highly useful when servers or compute nodes use some form of NIC teaming. To ensure that the nodes correctly detect uplink failures from the I/O modules, Layer 2 Failover must be enabled and configured on the I/O modules. If the uplinks go down with Layer 2 Failover enabled, the internal ports to the compute nodes are automatically shut down by the I/O module. NIC teaming/bonding is also used to fail the traffic over to the other NIC in the team, ensuring near seamless recovery for the nodes.

The combination of this architecture, NIC teaming on the host, and Layer 2 Failover on the I/O modules provides a highly available environment with no loops, and thus no wasted bandwidth to spanning-tree blocked links.

5.2.3 Non-looped, multiple upstream devices design

With the recent advent of virtualized chassis and virtual port-channeling technology from networking vendors (including IBM), a third general topology becomes available, which is illustrated in Figure 5-3.

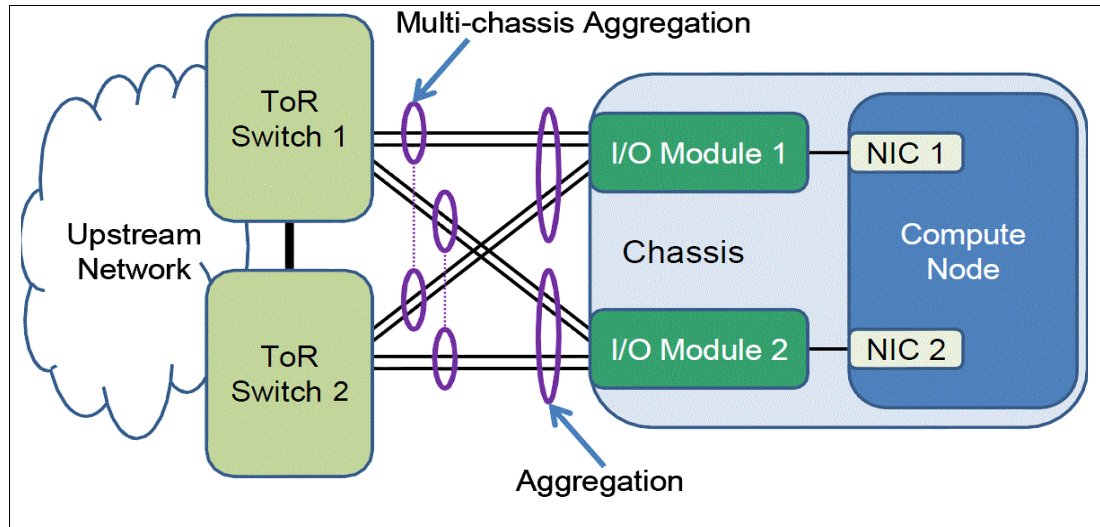


Figure 5-3 Non-looped, multiple upstream devices design, with hosts that can run either teamed or non-teamed NIC cards

The non-looped, multiple upstream devices design combines the best of both the looped and blocking design and the non-looped, single upstream device design in a robust, stable implementation. It is suitable for use with hosts that have either teamed or non-teamed NICs.

Offering the maximum bandwidth and high availability of the three topologies covered, this design requires the ToR switches to appear as a single logical switch to each I/O module in the Enterprise Chassis. This technology is vendor-specific at the time of this writing. However, the products of most major vendors support this function, including IBM System Networking products. The I/O modules in the implementation scenarios deploy the IBM Virtual Link Aggregation Group (vLAG) technology to the upstream ToR switch infrastructure to be displayed as a single, virtualized entity.

The designs that are reviewed in this section all assume that the L2/L3 boundary for the network is at or above the ToR switches in the diagrams. Ultimately, each environment must be analyzed to understand all the requirements and to ensure that the best design is selected and deployed.

5.3 Fully redundant with virtualized chassis technology (VSS/vPC/vLAG)

This implementation scenario incorporates switch virtualization features that allow a downstream switch to be connected to two upstream, virtualized switches through aggregated links, or port-channels. Inter-switch links (ISLs) between the same or similar products on the aggregation or access-layer provide a loop-free design that is both redundant and fully available in terms of bandwidth to the eventual downstream nodes. The switches are peers of one another, and synchronize their logical view of the access layer port structure. They internally prevent implicit loops. You this design if you want to use a best-practice implementation on a Cisco network that uses next generation networking features such as Cisco's Virtual Switching System (VSS) and Virtual Port Channel (vPC) technologies.

This approach has the following advantages:

- ▶ Active/Active uplinks helps to avoid the wasted bandwidth that is associated with links blocked by spanning tree
- ▶ Maximum redundancy and fault tolerance
- ▶ Extremely fast convergence times

This approach had the following disadvantages:

- ▶ Requires more expensive upstream equipment that supports virtualization features, and a network architect that is familiar with the implementation details
- ▶ More cabling and connections are necessary, increasing costs
- ▶ Careful implementation and planning are required to ensure correct operation

5.3.1 Components used

- ▶ Cisco Nexus 5548UP (Qty. 2)
- ▶ IBM G8264 RackSwitch™ (Qty. 2)
- ▶ IBM Flex System Fabric EN4093/R 10Gb Scalable Switch (Qty. 2)

5.3.2 Network topology and physical setup

Figure 5-4 shows the network topology for the fully redundant scenario with virtualized chassis technology (VSS/vPC/vLAG).

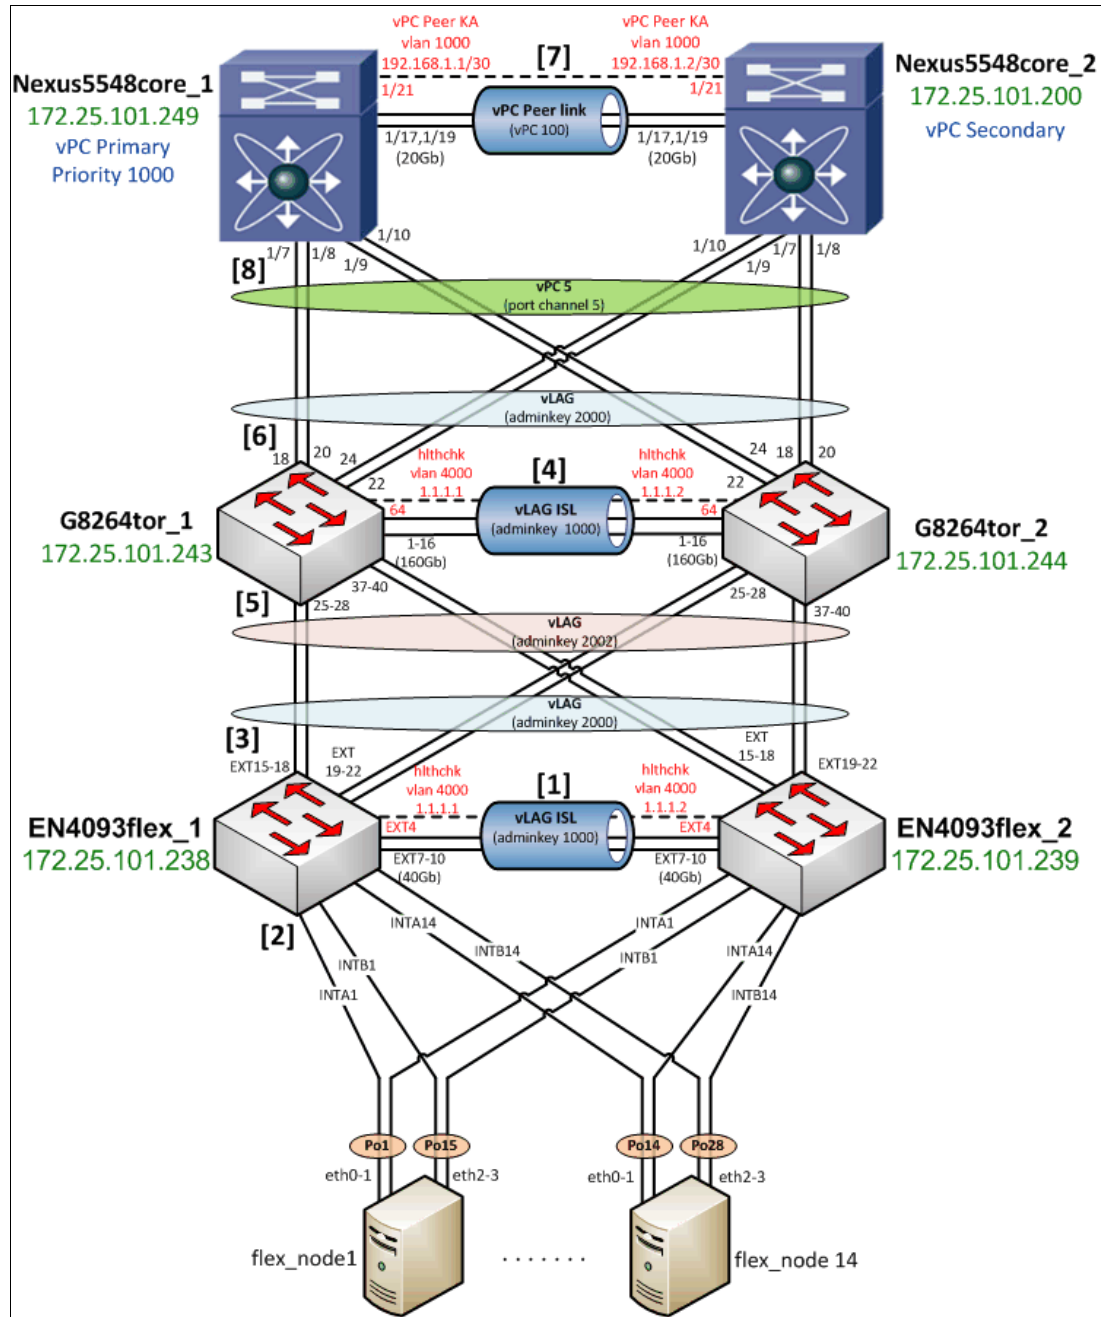


Figure 5-4 Network topology diagram for fully redundant scenario with virtualized chassis technology (VSS/vPC/vLAG)

Start by verifying the physical cabling between the EN4093/R switches and G8264's. The lab environment included four IBM QSFP+ DAC Break Out Cables from the EN4093/R switches to the upstream G8264's. This configuration requires that the EN4093/R switches be licensed for these particular features so that the ports can be used.

- ▶ Four 1m IBM QSFP+-to-QSFP+ Cables were used to form the 160 Gb ISL between the G8264 switches.
- ▶ 10Gb SFP+ DAC cables were used for all other connections in the diagram.

5.3.3 EN4093flex_1 configuration

Begin the implementation with the IBM Flex System Fabric EN4093/R switches, working up the diagram in Figure 5-4 on page 61. Each step provides the commands necessary and are reflective of the numbering schema in the diagram to aid the user in what is being configured.

General configuration

1. Create the ISL Healthcheck, ISL data, and Data VLANs as shown in Example 5-1, giving them descriptive names, assigning them to spanning-tree groups, and enabling them. You can elect to allow the switch itself to create STP instances for you. The example shows manually creating them instead.

Example 5-1 Creating ISL hlthchk, DATA, and ISL VLANs on EN4093flex_1

```
configure terminal
vlan 4000
    enable
    name "ISL hlthchk"
    stg 125
    exit
vlan 4092
    enable
    name "DATA"
    stg 126
    exit
vlan 4094
    enable
    name "ISL"
    stg 127
    exit
```

2. Assign IP addresses for both the ISL Healthcheck and Data VLANs as shown in Example 5-2. Doing so allows you to verify connectivity between the various pieces of equipment when verifying the configuration. In this example, interface ip 40 represents the vLAG Health Check IP address, and interface ip 92 represents an address on the Data VLAN that uses the prefix 10.1.4. The last octet is borrowed from the network diagram's Management address to quickly aid in the identification of which piece of equipment you are verifying connectivity to.

Example 5-2 Creating IP interfaces and assigning VLANs and IP addresses on EN4093flex_1

```
configure terminal
interface ip 40
    ip address 1.1.1.1 255.255.255.0
    vlan 4000
    enable
    exit
```

```
interface ip 92
  ip address 10.1.4.238 255.255.255.0
  vlan 4092
  enable
  exit
```

Configuring ISL between EN4093flex switches (step 1)

3. Configure the eventual ISL in Example 5-3 between the EN4093/R switches by configuring them to have a default (untagged) VLAN of 4094. Set an LACP key of 1000 to bundle the ports together in an aggregation, with 802.1q tagging enabled so that L2 VLAN traffic can traverse the ISL. Carry Data VLAN 4092 over these links.

Example 5-3 Initial ISL configuration on EN4093flex_1

```
configure terminal
interface port ext7-ext10
  pvid 4094
  tagging
  exit
vlan 4092
  member ext7-ext10
  exit
interface port ext7-ext10
  lacp key 1000
  lacp mode active
  exit
```

4. Create the dedicated health check VLAN and physical interface in Example 5-4 to be used for heartbeats between the EN4093/R switches. This example uses EXT4 as a dedicated interface and VLAN 4000 to serve as the health check for the ISL.

Example 5-4 Creating vLAG health check on EN4093flex_1

```
configure terminal
vlan 4000
  name "ISL hlthchk"
  enable
  exit
interface port ext4
  pvid 4000
  exit
```

5. Disable STP between the EN4093/R switches and activate a vLAG between them so that they appear as a single entity to upstream and downstream infrastructure as shown in Example 5-5, referencing the LACP key configured in the previous step.

Example 5-5 Disabling STP and activating ISL vLAG on EN4093flex_1

```
configure terminal
no spanning-tree stp 127 enable
vlag tier-id 1
vlag isl vlan 4094
vlag isl adminkey 1000
vlag hlthchk peer-ip 1.1.1.2
vlag enable
```

Configuring downstream internal node ports (step 2)

6. Configure the downstream node interfaces in Example 5-6 to have a default (untagged) VLAN of 4092, with 802.1q tagging enabled. Add the ability for all member ports to be on VLAN 4092.

Example 5-6 Downstream internal node port configuration on EN4093flex

```
configure terminal
interface port inta1-intb14
    pvid 4092
    tagging
    spanning-tree edge
    exit
vlan 4092
    member inta1-intb14
exit
```

7. For redundancy, create two port-channels on each of the 14 nodes. Each port-channel aggregates two ports, one from each EN4093flex switch. Have port channels 1-14 match the “A” internally labeled ports, and port channels 15-28 match the “B” ports as shown in Example 5-7.

Example 5-7 Node-facing port channel creation and vLAG activation, on EN4093flex_1

```
configure terminal
portchannel 1 port inta1
portchannel 1 enable
vlag portchannel 1 enable
portchannel 15 port intb1
portchannel 15 enable
vlag portchannel 15 enable
portchannel 2 port inta2
portchannel 2 enable
vlag portchannel 2 enable
portchannel 16 port intb2
portchannel 16 enable
vlag portchannel 16 enable
portchannel 3 port inta3
portchannel 3 enable
vlag portchannel 3 enable
portchannel 17 port intb3
portchannel 17 enable
vlag portchannel 17 enable
portchannel 4 port inta4
portchannel 4 enable
vlag portchannel 4 enable
portchannel 18 port intb4
portchannel 18 enable
vlag portchannel 18 enable
portchannel 5 port inta5
portchannel 5 enable
vlag portchannel 5 enable
portchannel 19 port intb5
portchannel 19 enable
vlag portchannel 19 enable
portchannel 6 port inta6
```

```
portchannel 6 enable
vlag portchannel 6 enable
portchannel 20 port intb6
portchannel 20 enable
vlag portchannel 20 enable
portchannel 7 port inta7
portchannel 7 enable
vlag portchannel 7 enable
portchannel 21 port intb7
portchannel 21 enable
vlag portchannel 21 enable
portchannel 8 port inta8
portchannel 8 enable
vlag portchannel 8 enable
portchannel 22 port intb8
portchannel 22 enable
vlag portchannel 22 enable
portchannel 9 port inta9
portchannel 9 enable
vlag portchannel 9 enable
portchannel 23 port intb9
portchannel 23 enable
vlag portchannel 23 enable
portchannel 10 port inta10
portchannel 10 enable
vlag portchannel 10 enable
portchannel 24 port intb10
portchannel 24 enable
vlag portchannel 24 enable
portchannel 11 port inta11
portchannel 11 enable
vlag portchannel 11 enable
portchannel 25 port intb11
portchannel 25 enable
vlag portchannel 25 enable
portchannel 12 port inta12
portchannel 12 enable
vlag portchannel 12 enable
portchannel 26 port intb12
portchannel 26 enable
vlag portchannel 26 enable
portchannel 13 port inta13
portchannel 13 enable
vlag portchannel 13 enable
portchannel 27 port intb13
portchannel 27 enable
vlag portchannel 27 enable
portchannel 14 port inta14
portchannel 14 enable
vlag portchannel 14 enable
portchannel 28 port intb14
portchannel 28 enable
vlag portchannel 28 enable
```

Configuring upstream, G8264tor facing ports, and layer 2 failover (step 3)

8. Configure the upstream ports with a default (untagged) VLAN of 4092 (Data vlan), tag the PVID and use an LACP key of 2000 to bundle the ports together as shown in Example 5-8.

Example 5-8 Upstream G8264 tor facing ports configuration on EN4093flex_1

```
configure terminal
interface port ext15-ext22
    pvid 4092
    tagging
    tag-pvid
    exit
vlan 4092
    member ext15-ext22
    exit
interface port ext15-ext22
    lacp key 2000
    lacp mode active
    exit
```

9. Activate the vLAG feature for the upstream EN4093/R ports so that the G8264s see the EN4093s as a single, virtualized entity as shown in Example 5-9. Use adminkey 2000, which represents the LACP key that is bundling ports EXT15-22 together as one.

Example 5-9 Activating the upstream G8264tor-facing vLAG on EN4093flex_1

```
configure terminal
vlag adminkey 2000 enable
```

10. Enable Layer-2 failover in Example 5-10, which shuts down the links to the compute nodes if the uplinks for the EN4093/R switch fail. This ensures that the downstream node is aware of the upstream failure and can fail traffic over to the other NIC in the node. In the example, the other NIC is connected to the other EN4093 switch in the Enterprise Chassis, ensuring that redundancy is maintained.

Example 5-10 Enabling layer 2 failover for the compute nodes on EN4093flex_1

```
configure terminal
failover trigger 1 mmon monitor admin-key 2000
failover trigger 1 mmon control member INTA1-INTB14
failover trigger 1 enable
failover enable
```

Repeat this configuration for EN4093_flex2 on the other I/O module. The only difference between the EN4093flex_1 switch and EN4093flex_2 switch is the vLAG health check peer address and the Data and ISL hlthchk VLAN IP addresses. To verify the EN4093flex switch configuration, run the **show** commands that are outlined in 5.3.7, “Verification and show command output” on page 73.

5.3.4 G8264tor_1 configuration

Next, configure RackSwitch G8264.

General configuration

1. Create the ISL Healthcheck, ISL data, and Data VLANs as shown in Example 5-11, giving them descriptive names, assigning them to spanning-tree groups, and enabling them.

Example 5-11 Creating ISL hlthchk, Data, and ISL VLANs on G8264tor_1

```
configure terminal
vlan 4000
    enable
    name "ISL hlthchk"
    stg 125
    exit
vlan 4092
    enable
    name "Data"
    stg 126
    exit
vlan 4094
    enable
    name "ISL"
    stg 127
    exit
```

2. Assign IP addresses for the ISL Healthcheck, Data VLANs, and management VLAN in Example 5-12. "Interface ip 128" represents the management IP address that is referenced in the Network Topology diagram. IP gateway 4 is the upstream router interface for the 172 management network.

Example 5-12 Creating IP interfaces and assigning VLANs and IP addresses, configuring management interface on G8264tor_1

```
configure terminal
interface ip 40
    ip address 1.1.1.1 255.255.255.0
    vlan 4000
    enable
    exit
interface ip 92
    ip address 10.1.4.243 255.255.255.0
    vlan 4092
    enable
    exit
interface ip 128
    ip address 172.25.101.243 255.255.0.0
    enable
    exit
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
```

Configuring ISL between G8264tor switches (step 4)

3. Configure the ISL between the G8264 switches as shown in Example 5-13. Make the default (untagged) VLAN 4094 (ISL VLAN). Assign an LACP key of 1000 to bundle the ports together in an aggregation, with 802.1q tagging enabled so that L2 VLAN traffic can traverse the ISL. Allow VLAN 4092 (data VLAN) over these links.

Example 5-13 Initial ISL configuration on G8264tor_1

```
configure terminal
interface port 1-16
    pvid 4094
    tagging
    exit
vlan 4092
    member 1-16
    exit
interface port 1-16
    lacp key 1000
    lacp mode active
    exit
```

4. Disable STP between the G8264 switches and activate a vLAG between them so that they appear as a single entity to upstream and downstream infrastructure as shown in Example 5-14. Reference the LACP key that was configured in the previous step.

Example 5-14 Disabling STP and activating ISL vLAG on G8264tor_1

```
configure terminal
no spanning-tree stp 127 enable
vlag tier-id 2
vlag isl vlan 4094
vlag isl adminkey 1000
vlag hlthchk peer-ip 1.1.1.2
vlag enable
```

Configuring downstream EN4093flex facing ports (step 5)

5. Configure the downstream EN4093flex facing ports as shown in Example 5-15. Make the default (untagged) VLAN 4092 (data VLAN), with 802.1q tagging enabled. Add the ability for all member ports to be on VLAN 4092.

Example 5-15 Configuring downstream EN4093flex facing ports on G8264tor_1

```
configure terminal
interface port 25-28,37-40
    pvid 4092
    tagging
    tag-pvid
    exit
vlan 4092
    member 25-28,37-40
    exit
interface port 25-28,37-40
    lacp key 2002
    lacp mode active
    exit
```

6. Activate the vLAG feature for the downstream EN4093flex facing ports so that the EN4093s see the G8264s as a single, virtualized entity as shown in Example 5-16. Use adminkey 2002, which represents the LACP key bundling ports 25-28, and 37-40 together as one.

Example 5-16 Activating downstream EN4093flex facing vLAG on G8264tor_1

```
configure terminal
vlag adminkey 2002 enable
```

Configuring upstream Nexus5548core facing ports (step 6)

7. Configure the upstream Nexus5548core facing ports as shown in Example 5-17 with a default (untagged) VLAN of 4092 (data VLAN). Tag the PVID, and use an LACP key of 2000 to bundle the ports together in an aggregation.

Example 5-17 Configuring upstream Nexus5548core facing ports on G8264tor_1

```
configure terminal
interface port 18,20,22,24
    pvid 4092
    tagging
    tag-pvid
    exit
vlan 4092
    member 18,20,22,24
    exit
interface port 18,20,22,24
    lacp key 2000
    lacp mode active
    exit
```

8. Activate the vLAG feature for the ports that are bundled with LACP key 2000, which the Nexus pair sees as a single, virtualized entity as shown in Example 5-18.

Example 5-18 Activating the upstream Nexus5548core facing vLAG, adminkey 2000 on G8264tor_1

```
configure terminal
vlag adminkey 2000 enable
```

Now repeat this configuration for G8264tor_2. The only difference between the G8264tor_1 switch and the G8264tor_2 switch is the vLAG health check peer address and the Data, management and ISL hlthchk VLAN IP addresses. To verify G8264tor switch configuration, run the **show** commands that are outlined in 5.3.7, “Verification and show command output” on page 73.

5.3.5 Nexus5548core_1 vPC primary switch configuration

Next, configure the Cisco Nexus 5548 primary core switch.

General configuration

1. Be sure that the features shown in Example 5-19 are enabled.

Example 5-19 Enabling NX-OS feature sets on Nexus5548core_1

```
configure terminal
feature interface-vlan
feature lacp
feature vpc
feature lldp
```

2. Create the DATA (4092) and vPC_PEER_LINK (1000) VLANs. Set the spanning-tree priority for the DATA (4092) VLAN to be half that of the Nexus5548core_2 switch. Because Nexus5548core_1 switch has a lower spanning-tree priority, it becomes the root bridge for layer 2 spanning-tree function as shown in Example 5-20.

Example 5-20 Data (4092) and vPC_PEER_LINK (1000) VLAN creation/STP priority configuration on Nexus5548core_1

```
configure terminal
vlan 4092
    name DATA_VLAN
vlan 1000
    name vPC_PEER_LINK
spanning-tree vlan 4092 priority 8192
```

Configuring virtual port channel (vPC) on Nexus5548core_1 (step 7)

3. Configure a VRF (virtual routing and forwarding) for the vPC peer link in Example 5-21. Build vPC domain 100, which will be bound to port-channel100 in the next section. Because the Nexus box has a Layer-3 card and license, create the Switched Virtual Interfaces (SVIs) for the Data and ISL VLANs.

Example 5-21 Configuring vPC domain on Nexus5548core_1

```
configure terminal
vrf context VPCKeepAlive
vpc domain 100
    role priority 1000
    peer-keepalive destination 192.168.1.2 source 192.168.1.1 vrf VPCKeepAlive
interface Vlan4092
    no shutdown
    ip address 10.1.4.249/24
interface Vlan1000
    no shutdown
    vrf member VPCKeepAlive
    ip address 192.168.1.1/30
```

4. Configure the physical interfaces that comprise the vPC peer link between the Nexus 5548-1 and 5548-2 switches as shown in Example 5-22. Use port-channel100 and Dynamic Link Aggregation Control Protocol (LACP).

Example 5-22 vPC peer-link physical and logical interface configuration on Nexus5548core_1

```
configure terminal
interface Ethernet1/17
    description vPC Peer link to Nexus5548core_2
    switchport mode trunk
```

```

switchport trunk allowed vlan 4092
channel-group 100 mode active
interface Ethernet1/19
description vPC Peer link to Nexus5548core_2
switchport mode trunk
switchport trunk allowed vlan 4092
channel-group 100 mode active
interface port-channel100
description "vPC Peer Link"
switchport mode trunk
vpc peer-link
switchport trunk allowed vlan 4092
spanning-tree port type network

```

5. Set up the vPC peer keepalive link to monitor the partners' health status as shown in Example 5-23. Increase the keepalive robustness with a separate, dedicated physical link for keepalives in a dedicated VRF so that the system cannot mis-direct traffic that is routed to the vPC peer keepalive address.

Example 5-23 vPC peer keepalive link configuration on Nexus5548core_1

```

configure terminal
interface Ethernet1/21
no shutdown
description vPC Keep alive
switchport access vlan 1000

```

Configuring downstream G8264tor facing ports (step 8)

6. For the Nexus 5548 primary switch, configure the downstream physical and logical interfaces in Example 5-24 to all be on the same virtual port-channel by using LACP aggregation. This configuration ensures that the Nexus pair presents itself as a single, logical entity to the G8264s.

Example 5-24 Downstream G8264tor facing interfaces on Nexus5548core_1

```

configure terminal
interface Ethernet1/7-10
switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
channel-group 5 mode active
interface port-channel5
switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
vpc 5

```

5.3.6 Cisco Nexus 5548core_2 vPC secondary switch configuration

Lastly, configure the Cisco Nexus5548core_2 vPC secondary switch.

General configuration

7. Be sure that the features shown in Example 5-25 are enabled.

Example 5-25 Enable Cisco Nexus feature sets on Nexus5548core_2

```
configure terminal
feature interface-vlan
feature lacp
feature vpc
feature lldp
```

8. Create the VLANs shown in Example 5-26. Set the spanning-tree priority of the Data VLAN to twice that of Nexus5548core_1, ensuring that the Nexus5548core_2 switch is the backup from a spanning-tree standpoint.

Example 5-26 Data (4092) and vPC_PEER_LINK (1000) VLAN creation/STP priority configuration on Nexus5548core_2

```
configure terminal
vlan 4092
    name DATA_VLAN
vlan 1000
    name vPC_PEER_LINK
spanning-tree vlan 4092 priority 16384
```

Configuring virtual port channel (vPC) on Nexus5548core_2 (step 7)

9. Configure a VRF for the vPC peer link as shown in Example 5-27. Create the SVIs for the Data and ISL VLANs.

Example 5-27 Configure vPC domain on Nexus5548core_2

```
configure terminal
vrf context VPCKeepAlive
vpc domain 100
    peer-keepalive destination 192.168.1.1 source 192.168.1.2 vrf VPCKeepAlive
interface Vlan4092
    no shutdown
    ip address 10.1.4.200/24
interface Vlan1000
    no shutdown
    vrf member VPCKeepAlive
    ip address 192.168.1.2/30
```

10. Configure the physical interfaces that comprise the vPC peer link between the Nexus 5548 switches as shown in Example 5-28. Use port-channel100 and LACP.

Example 5-28 vPC peer-link physical and logical interface configuration on Nexus5548_core2

```
configure terminal
interface Ethernet1/17
    description vPC Peer link to Nexus5548core_1
    switchport mode trunk
    switchport trunk allowed vlan 4092
    channel-group 100 mode active
interface Ethernet1/19
    description vPC Peer link to Nexus5548core_1
    switchport mode trunk
```

```
switchport trunk allowed vlan 4092
channel-group 100 mode active
interface port-channel100
description "VPC Peer Link"
switchport mode trunk
vpc peer-link
switchport trunk allowed vlan 4092
spanning-tree port type network
```

11. Set up the vPC peer keepalive to monitor health status between the Nexus pair as shown in Example 5-29.

Example 5-29 vPC peer keepalive link configuration on Nexus5548core_2

```
configure terminal
interface Ethernet1/21
no shutdown
description VPC KeepAlive
switchport access vlan 1000
```

Configuring downstream G8264tor facing ports (step 8)

12. For the Nexus5548core_2 switch, configure the downstream G8264tor facing physical and logical interfaces in Example 5-30 to all be on the same virtual port-channel by using LACP aggregation. This configuration ensures that the Nexus pair presents itself as a single, logical entity to the G8264s.

Example 5-30 Downstream G8264tor facing interface configuration on Nexus5548core_2

```
configure terminal
interface Ethernet1/7-10
switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
channel-group 5 mode active
interface port-channel5
switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
vpc 5
```

5.3.7 Verification and show command output

The following section lists output from common **show** commands that can aid the network architect in the implementation of this scenario. Ping verification of the various IP addresses that are configured on the equipment for the Data VLAN is also done to show that all of the devices can reach each other successfully.

As in the implementation section, helpful commands are listed from the EN4093/R switches, working your way up the Network Topology diagram to the Cisco Nexus pair.

EN4093/R output

This section shows output from the switch with hostname EN4093flex_1. Similar or identical output exists for the switch with hostname EN4093flex_2.

Show version

The command output in Example 5-31 shows information about the switch, and the associated code/firmware level at the time.

Example 5-31 EN4093flex_1 show version output

```
System Information at 23:04:56 Fri Oct 12, 2012
Time zone: No timezone configured
Daylight Savings Time Status: Disabled

IBM Flex System Fabric EN4093 10Gb Scalable Switch

Switch has been up for 1 day, 2 hours, 1 minute and 21 seconds.
Last boot: 21:05:54 Thu Oct 11, 2012 (reset from Telnet/SSH)

MAC address: 6c:ae:8b:bf:6d:00    IP (If 40) address: 1.1.1.1
Internal Management Port MAC Address: 6c:ae:8b:bf:6d:ef
Internal Management Port IP Address (if 128): 172.25.101.238
External Management Port MAC Address: 6c:ae:8b:bf:6d:fe
External Management Port IP Address (if 127):
Software Version 7.3.1.0          (FLASH image1), active configuration.


Hardware Part Number      : 49Y4272
Hardware Revision        : 02
Serial Number            : Y250VT24M099
Manufacturing Date (WWYY) : 1712
PCBA Part Number         : BAC-00072-01
PCBA Revision            : 0
PCBA Number              : 00
Board Revision           : 02
PLD Firmware Version     : 1.5


Temperature Warning       : 32 C (Warn at 60 C/Recover at 55 C)
Temperature Shutdown     : 32 C (Shutdown at 65 C/Recover at 60 C)
Temperature Inlet        : 27 C
Temperature Exhaust      : 33 C


Power Consumption         : 54.300 W (12.244 V, 4.435 A)

Switch is in I/O Module Bay 1
```

Show vlan

Example 5-32 shows output regarding VLAN assignment for all the various ports on the switch.

Example 5-32 EN4093flex_1 show vlan output

VLAN	Name	Status	MGT	Ports
----	-----	-----	-----	-----
1	Default VLAN	ena	dis	EXT1-EXT3 EXT5 EXT6
4000	ISL hlthchk	ena	dis	EXT4
4092	DATA	ena	dis	INTA1-INTB14 EXT7-EXT10 EXT15-EXT22

```

4094 ISL                               ena    dis  EXT7-EXT10
4095 Mgmt VLAN                         ena    ena  EXTM MGT1

```

Show interface status

Because there is only one compute node in the chassis (in slot 1), all the other internal ports are listed as down from a link perspective in the output shown in Example 5-33.

Example 5-33 EN4093flex_1 show interface status output

Alias	Port	Speed	Duplex	Flow Ctrl		Link	Name
				--TX--	--RX--		
INTA1	1	1000	full	no	no	up	INTA1
INTA2	2	1G/10G	full	yes	yes	down	INTA2
INTA3	3	1G/10G	full	yes	yes	down	INTA3
INTA4	4	1G/10G	full	yes	yes	down	INTA4
INTA5	5	1G/10G	full	yes	yes	down	INTA5
INTA6	6	1G/10G	full	yes	yes	down	INTA6
INTA7	7	1G/10G	full	yes	yes	down	INTA7
INTA8	8	1G/10G	full	yes	yes	down	INTA8
INTA9	9	1G/10G	full	yes	yes	down	INTA9
INTA10	10	1G/10G	full	yes	yes	down	INTA10
INTA11	11	1G/10G	full	yes	yes	down	INTA11
INTA12	12	1G/10G	full	yes	yes	down	INTA12
INTA13	13	1G/10G	full	yes	yes	down	INTA13
INTA14	14	1G/10G	full	yes	yes	down	INTA14
INTB1	15	1000	full	no	no	up	INTB1
INTB2	16	1G/10G	full	yes	yes	down	INTB2
INTB3	17	1G/10G	full	yes	yes	down	INTB3
INTB4	18	1G/10G	full	yes	yes	down	INTB4
INTB5	19	1G/10G	full	yes	yes	down	INTB5
INTB6	20	1G/10G	full	yes	yes	down	INTB6
INTB7	21	1G/10G	full	yes	yes	down	INTB7
INTB8	22	1G/10G	full	yes	yes	down	INTB8
INTB9	23	1G/10G	full	yes	yes	down	INTB9
INTB10	24	1G/10G	full	yes	yes	down	INTB10
INTB11	25	1G/10G	full	yes	yes	down	INTB11
INTB12	26	1G/10G	full	yes	yes	down	INTB12
INTB13	27	1G/10G	full	yes	yes	down	INTB13
INTB14	28	1G/10G	full	yes	yes	down	INTB14
EXT1	43	10000	full	no	no	up	EXT1
EXT2	44	10000	full	no	no	up	EXT2
EXT3	45	10000	full	no	no	up	EXT3
EXT4	46	10000	full	no	no	up	ISL h1thchk
EXT5	47	1G/10G	full	no	no	down	EXT5
EXT6	48	1G/10G	full	no	no	down	EXT6
EXT7	49	10000	full	no	no	up	ISL
EXT8	50	10000	full	no	no	up	ISL
EXT9	51	10000	full	no	no	up	ISL
EXT10	52	10000	full	no	no	up	ISL
EXT15	57	10000	full	no	no	up	Link to g8264tor_1
EXT16	58	10000	full	no	no	up	Link to g8264tor_1
EXT17	59	10000	full	no	no	up	Link to g8264tor_1
EXT18	60	10000	full	no	no	up	Link to g8264tor_1
EXT19	61	10000	full	no	no	up	Link to g8264tor_2

EXT20	62	10000	full	no	no	up	Link to g8264tor_2
EXT21	63	10000	full	no	no	up	Link to g8264tor_2
EXT22	64	10000	full	no	no	up	Link to g8264tor_2
EXTM	65	1000	half	yes	yes	down	EXTM
MGT1	66	1000	full	yes	yes	up	MGT1

Show lldp remote-device

Example 5-34 command output illustrates the physical topology and verifies that cables are plugged into the ports that are specified in both the Network Topology diagram, and the configuration specified in the appendix.

Example 5-34 EN4093flex_1 show lldp remote-device output

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
EXT16	3	08 17 f4 33 9d 00	25	G8264TOR-1
EXT15	4	08 17 f4 33 9d 00	26	G8264TOR-1
EXT18	5	08 17 f4 33 9d 00	27	G8264TOR-1
EXT17	6	08 17 f4 33 9d 00	28	G8264TOR-1
EXT21	7	08 17 f4 33 75 00	25	G8264TOR-2
EXT19	8	08 17 f4 33 75 00	26	G8264TOR-2
EXT22	9	08 17 f4 33 75 00	27	G8264TOR-2
EXT20	10	08 17 f4 33 75 00	28	G8264TOR-2
EXT4	12	6c ae 8b bf fe 00	46	en4093flex_2
EXT7	13	6c ae 8b bf fe 00	49	en4093flex_2
EXT8	14	6c ae 8b bf fe 00	50	en4093flex_2
EXT9	15	6c ae 8b bf fe 00	51	en4093flex_2
EXT10	16	6c ae 8b bf fe 00	52	en4093flex_2

Show vlag isl

Example 5-35 shows command output about the status of the ISL between the EN4093/R switches, and the ports that comprise the ISL itself.

Example 5-35 EN4093flex_1 show vlag isl output

ISL_ID	ISL_Vlan	ISL_Trunk	ISL_Members	Link_State	Trunk_State
65	4094	Adminkey 1000	EXT7	UP	UP
			EXT8	UP	UP
			EXT9	UP	UP
			EXT10	UP	UP

Show vlag information

The command output in Example 5-36 on page 77 shows that the vLAG between the EN4093/R switches and G8264 switches is up and operational as referenced by the LACP admin key of 2000. The ISL between the EN4093/R switches is up as well.

EN4093flex_1 is acting as the admin and operational role of PRIMARY. For centralized vLAG functions, such as vLAG STP, one of the vLAG switch must control the protocol operations. Select which switch controls the centralized vLAG function by performing role election. The switch with the primary role controls the centralized operation. Role election is non-preemptive. That is, if a primary already exists, another switch coming up remains as secondary even if it can become primary based on the role election logic.

Role election is determined by comparing the local vLAG system priority and local system MAC address. The switch with the smaller priority value is the vLAG primary switch. If the priorities are the same, the switch with the smaller system MAC address is the vLAG primary switch. It is possible to configure vLAG priority to anything between <0-65535>. The priority was left at the default value of 0 in all examples.

Example 5-36 EN4093flex_1 show vlag information output

```
vLAG Tier ID: 1
vLAG system MAC: 08:17:f4:c3:dd:00
Local MAC 6c:ae:8b:bf:6d:00 Priority 0 Admin Role PRIMARY (Operational Role
PRIMARY)
Peer MAC 6c:ae:8b:bf:fe:00 Priority 0
Health local 1.1.1.1 peer 1.1.1.2 State UP
ISL trunk id 65
ISL state Up
Startup Delay Interval: 120s (Finished)
```

vLAG 65: config with admin key 2000, associated trunk 66, state formed

Show vlag adminkey 2000

The output in Example 5-37 shows that the vLAG is formed and enabled by using LACP reference key 2000.

Example 5-37 EN4093flex_1 show vlag adminkey 2000 output

```
vLAG is enabled on admin key 2000
Current LACP params for EXT15: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT16: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT17: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT18: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT19: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT20: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT21: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT22: active, Priority 32768, Admin Key 2000, Min-Links 1
```

Show lacp information state up

The command output in Example 5-38 shows which ports are participating in an LACP aggregation, and which reference keys are used on those specific interfaces.

Example 5-38 EN4093flex_1 show lacp information state up

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
<hr/>									
EXT7	active	1000	1000	yes	32768	49	65	up	1
EXT8	active	1000	1000	yes	32768	49	65	up	1
EXT9	active	1000	1000	yes	32768	49	65	up	1
EXT10	active	1000	1000	yes	32768	49	65	up	1
EXT15	active	2000	2000	yes	32768	57	66	up	1

EXT16	active	2000	2000	yes	32768	57	66	up	1
EXT17	active	2000	2000	yes	32768	57	66	up	1
EXT18	active	2000	2000	yes	32768	57	66	up	1
EXT19	active	2000	2000	yes	32768	57	66	up	1
EXT20	active	2000	2000	yes	32768	57	66	up	1
EXT21	active	2000	2000	yes	32768	57	66	up	1
EXT22	active	2000	2000	yes	32768	57	66	up	1

Show failover trigger 1

The failover output in Example 5-39 shows which ports are monitored, and which ports are shut down if an issue is encountered. In this example, the upstream to G8264 links are monitored with LACP reference key 2000. The control ports are the downstream internal I/O module ports that are used by the compute nodes.

Example 5-39 EN4093flex_1 show failover output

```

Failover: On
VLAN Monitor: OFF

Trigger 1 Manual Monitor: Enabled
Trigger 1 limit: 0
Monitor State: Up
Member      Status
-----
adminkey 2000
EXT15      Operational
EXT16      Operational
EXT17      Operational
EXT18      Operational
EXT19      Operational
EXT20      Operational
EXT21      Operational
EXT22      Operational
Control State: Auto Controlled
Member      Status
-----
INTA1      Operational
INTA2      Operational
INTA3      Operational
INTA4      Operational
INTA5      Operational
INTA6      Operational
INTA7      Operational
INTA8      Operational
INTA9      Operational
INTA10     Operational
INTA11     Operational
INTA12     Operational
INTA13     Operational
INTA14     Operational
INTB1      Operational
INTB2      Operational
INTB3      Operational
INTB4      Operational
INTB5      Operational

```

INTB6	Operational
INTB7	Operational
INTB8	Operational
INTB9	Operational
INTB10	Operational
INTB11	Operational
INTB12	Operational
INTB13	Operational
INTB14	Operational

Trigger 2: Disabled

Trigger 3: Disabled

Trigger 4: Disabled

Trigger 5: Disabled

Trigger 6: Disabled

Trigger 7: Disabled

Trigger 8: Disabled

Ping output for equipment on VLAN 4092

To verify connectivity, issue **ping** commands to devices on VLAN 4092 (Data VLAN) as shown in Example 5-40. IP address 10.4.1.10 represents a compute node with an operating system installed, flex_node1 on the Network Topology diagram.

Example 5-40 Ping verification for equipment on VLAN 4092

```
en4093flex_1#ping 10.1.4.10 data-port
Connecting via DATA port.
[host 10.1.4.10, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl 255,
tos 0]
10.1.4.10: #1 ok, RTT 1 msec.
10.1.4.10: #2 ok, RTT 0 msec.
10.1.4.10: #3 ok, RTT 1 msec.
10.1.4.10: #4 ok, RTT 0 msec.
10.1.4.10: #5 ok, RTT 0 msec.
Ping finished.
```

```
en4093flex_1#ping 10.1.4.239 data-port
Connecting via DATA port.
[host 10.1.4.239, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.239: #1 ok, RTT 4 msec.
10.1.4.239: #2 ok, RTT 1 msec.
10.1.4.239: #3 ok, RTT 2 msec.
10.1.4.239: #4 ok, RTT 3 msec.
10.1.4.239: #5 ok, RTT 1 msec.
Ping finished.
```

```
en4093flex_1#ping 10.1.4.243 data-port
Connecting via DATA port.
```

```
[host 10.1.4.243, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.243: #1 ok, RTT 1 msec.
10.1.4.243: #2 ok, RTT 1 msec.
10.1.4.243: #3 ok, RTT 2 msec.
10.1.4.243: #4 ok, RTT 8 msec.
10.1.4.243: #5 ok, RTT 6 msec.
Ping finished.
```

```
en4093flex_1#ping 10.1.4.244 data-port
Connecting via DATA port.
[host 10.1.4.244, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.244: #1 ok, RTT 1 msec.
10.1.4.244: #2 ok, RTT 2 msec.
10.1.4.244: #3 ok, RTT 1 msec.
10.1.4.244: #4 ok, RTT 2 msec.
10.1.4.244: #5 ok, RTT 0 msec.
Ping finished.
```

```
en4093flex_1#ping 10.1.4.249 data-port
Connecting via DATA port.
[host 10.1.4.241, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.241: #1 ok, RTT 2 msec.
10.1.4.241: #2 ok, RTT 1 msec.
10.1.4.241: #3 ok, RTT 2 msec.
10.1.4.241: #4 ok, RTT 1 msec.
10.1.4.241: #5 ok, RTT 3 msec.
Ping finished.
```

```
en4093flex_1#ping 10.1.4.200 data-port
Connecting via DATA port.
[host 10.1.4.241, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.241: #1 ok, RTT 2 msec.
10.1.4.241: #2 ok, RTT 2 msec.
10.1.4.241: #3 ok, RTT 2 msec.
10.1.4.241: #4 ok, RTT 1 msec.
10.1.4.241: #5 ok, RTT 3 msec.
Ping finished
```

G8264 output

This section lists output from the switch with hostname G8264tor_1. Similar or identical output exists for the switch with hostname G8264tor_2.

Show version

Example 5-41 output shows information about the switch and the associated code/firmware level.

Example 5-41 G8264tor_1 show version output

```
System Information at 20:30:07 Thu Oct 18, 2012
Time zone: No timezone configured
Daylight Savings Time Status: Disabled
```

IBM Networking Operating System RackSwitch G8264

Switch has been up for 1 day, 20 hours, 28 minutes and 18 seconds.
Last boot: 6:05:44 Thu Feb 7, 2001 (reset from console)

MAC address: 08:17:f4:33:9d:00 IP (If 20) address: 10.10.20.2
Management Port MAC Address: 08:17:f4:33:9d:fe
Management Port IP Address (if 128): 172.25.101.243
Hardware Revision: 0
Hardware Part No: BAC-00065-00
Switch Serial No: US71120007
Manufacturing date: 11/13

Software Version 7.4.1.0 (FLASH image1), active configuration.

Temperature Mother Top: 26 C
Temperature Mother Bottom: 32 C
Temperature Daughter Top: 26 C
Temperature Daughter Bottom: 30 C

Warning at 75 C and Recover at 90 C

Fan 1 in Module 1: RPM= 8463 PWM= 15(5%) Front-To-Back
Fan 2 in Module 1: RPM= 3976 PWM= 15(5%) Front-To-Back
Fan 3 in Module 2: RPM= 8667 PWM= 15(5%) Front-To-Back
Fan 4 in Module 2: RPM= 4115 PWM= 15(5%) Front-To-Back
Fan 5 in Module 3: RPM= 7894 PWM= 15(5%) Front-To-Back
Fan 6 in Module 3: RPM= 4195 PWM= 15(5%) Front-To-Back
Fan 7 in Module 4: RPM= 8852 PWM= 15(5%) Front-To-Back
Fan 8 in Module 4: RPM= 3976 PWM= 15(5%) Front-To-Back

System Fan Airflow: Front-To-Back

Power Supply 1: OK
Power Supply 2: OK

Power Faults: ()
Fan Faults: ()
Service Faults: ()

Show vlan

The output in Example 5-42 shows VLAN assignment for all of the ports on the switch.

Example 5-42 G8264tor_1 show vlan output

VLAN	Name	Status	Ports
----	-----	-----	-----
1	Default VLAN	ena	17-63
4000	ISL hlthchk	ena	64
4092	DATA	ena	1-16 18 20 22 24-28 37-40

4094	ISL	ena	1-16
4095	Mgmt VLAN	ena	MGT

Show interface status

The output in Example 5-43 displays the interface status information.

Example 5-43 G8264tor_1 show interface status output

Alias	Port	Speed	Duplex	Flow Ctrl		Link	Name
				--TX--	--RX--		
1	1	10000	full	no	no	up	ISL
2	2	10000	full	no	no	up	ISL
3	3	10000	full	no	no	up	ISL
4	4	10000	full	no	no	up	ISL
5	5	10000	full	no	no	up	ISL
6	6	10000	full	no	no	up	ISL
7	7	10000	full	no	no	up	ISL
8	8	10000	full	no	no	up	ISL
9	9	10000	full	no	no	up	ISL
10	10	10000	full	no	no	up	ISL
11	11	10000	full	no	no	up	ISL
12	12	10000	full	no	no	up	ISL
13	13	10000	full	no	no	up	ISL
14	14	10000	full	no	no	up	ISL
15	15	10000	full	no	no	up	ISL
16	16	10000	full	no	no	up	ISL
17	17	1G/10G	full	no	no	down	17
18	18	10000	full	no	no	up	VLAG to
Nexus5548Core_1							
19	19	1G/10G	full	no	no	down	19
20	20	10000	full	no	no	up	VLAG to
Nexus5548Core_1							
21	21	1G/10G	full	no	no	down	21
22	22	10000	full	no	no	up	VLAG to
Nexus5548Core_2							
23	23	1G/10G	full	no	no	down	23
24	24	10000	full	no	no	up	VLAG to
Nexus5548Core_2							
25	25	10000	full	no	no	up	Link to EN4093-1
26	26	10000	full	no	no	up	Link to EN4093-1
27	27	10000	full	no	no	up	Link to EN4093-1
28	28	10000	full	no	no	up	Link to EN4093-1
29	29	1G/10G	full	no	no	down	29
30	30	1G/10G	full	no	no	down	30
31	31	1G/10G	full	no	no	down	31
32	32	1G/10G	full	no	no	down	32
33	33	1G/10G	full	no	no	down	33
34	34	1G/10G	full	no	no	down	34
35	35	1G/10G	full	no	no	down	35
36	36	1G/10G	full	no	no	down	36
37	37	10000	full	no	no	up	Link to EN4093-2
38	38	10000	full	no	no	up	Link to EN4093-2
39	39	10000	full	no	no	up	Link to EN4093-2
40	40	10000	full	no	no	up	Link to EN4093-2

41	41	1G/10G	full	no	no	down	41
42	42	1G/10G	full	no	no	down	42
43	43	1G/10G	full	no	no	down	43
44	44	1G/10G	full	no	no	down	44
45	45	1G/10G	full	no	no	down	45
46	46	1G/10G	full	no	no	down	46
47	47	1G/10G	full	no	no	down	47
48	48	1G/10G	full	no	no	down	48
49	49	1G/10G	full	no	no	down	49
50	50	1G/10G	full	no	no	down	50
51	51	1G/10G	full	no	no	down	51
52	52	1G/10G	full	no	no	down	52
53	53	1G/10G	full	no	no	down	53
54	54	1G/10G	full	no	no	down	54
55	55	1G/10G	full	no	no	down	55
56	56	1G/10G	full	no	no	down	56
57	57	1G/10G	full	no	no	down	57
58	58	1G/10G	full	no	no	down	58
59	59	1G/10G	full	no	no	down	59
60	60	1G/10G	full	no	no	down	60
61	61	1G/10G	full	no	no	down	61
62	62	1G/10G	full	no	no	down	62
63	63	1G/10G	full	no	no	down	63
64	64	10000	full	no	no	up	ISL h1thchk
MGT	65	1000	full	yes	yes	up	MGT

Show lldp remote-device

The command output in Example 5-44 shows the physical topology and verifies that cables are plugged into the ports specified in both the Network Topology diagram, and the configuration specified in the appendix.

Example 5-44 G8264tor_1 show lldp remote-device output

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
1	2	08 17 f4 33 75 00	1	G8264TOR-2
2	3	08 17 f4 33 75 00	2	G8264TOR-2
3	4	08 17 f4 33 75 00	3	G8264TOR-2
4	5	08 17 f4 33 75 00	4	G8264TOR-2
5	6	08 17 f4 33 75 00	5	G8264TOR-2
6	7	08 17 f4 33 75 00	6	G8264TOR-2
26	8	6c ae 8b bf 6d 00	57	en4093flex_1
18	9	54 7f ee 2d 36 0e	Eth1/7	Nexus5548core_1
25	10	6c ae 8b bf 6d 00	58	en4093flex_1
7	11	08 17 f4 33 75 00	7	G8264TOR-2
28	12	6c ae 8b bf 6d 00	59	en4093flex_1
27	13	6c ae 8b bf 6d 00	60	en4093flex_1
8	14	08 17 f4 33 75 00	8	G8264TOR-2
37	15	6c ae 8b bf fe 00	57	en4093flex_2
39	16	6c ae 8b bf fe 00	58	en4093flex_2
9	17	08 17 f4 33 75 00	9	G8264TOR-2
20	18	54 7f ee 2d 36 0f	Eth1/8	Nexus5548core_1
38	19	6c ae 8b bf fe 00	59	en4093flex_2
10	20	08 17 f4 33 75 00	10	G8264TOR-2
40	21	6c ae 8b bf fe 00	60	en4093flex_2

22	22	00 05 73 bc 02 70	Eth1/9	Nexus5548core_2
24	23	00 05 73 bc 02 71	Eth1/10	Nexus5548core_2
11	24	08 17 f4 33 75 00	11	G8264TOR-2
12	25	08 17 f4 33 75 00	12	G8264TOR-2
13	26	08 17 f4 33 75 00	13	G8264TOR-2
14	27	08 17 f4 33 75 00	14	G8264TOR-2
15	28	08 17 f4 33 75 00	15	G8264TOR-2
16	29	08 17 f4 33 75 00	16	G8264TOR-2
64	30	08 17 f4 33 75 00	64	G8264TOR-2

Show vlag isl

The command output in Example 5-45 shows the status of the ISL between the G8264 switches, and the ports that comprise the ISL itself.

Example 5-45 G8264tor_1 show vlag isl output

ISL_ID	ISL_Vlan	ISL_Trunk	ISL_Members	Link_State	Trunk_State
67	4094	Adminkey 1000	1	UP	UP
			2	UP	UP
			3	UP	UP
			4	UP	UP
			5	UP	UP
			6	UP	UP
			7	UP	UP
			8	UP	UP
			9	UP	UP
			10	UP	UP
			11	UP	UP
			12	UP	UP
			13	UP	UP
			14	UP	UP
			15	UP	UP
			16	UP	UP

Show vlag information

Example 5-46 on page 85 output shows that the downstream vLAG between the G8264 and EN4093 switches is up and operational as referenced by the LACP admin key of 2002. Also shown is the upstream vLAG between the G8264 and Nexus switches, referenced by the LACP admin key of 2000. The ISL between the G8264 switches is up as well.

G8264tor_1 is acting as the admin and operational role of SECONDARY. For centralized vLAG functions, such as vLAG STP, one of the vLAG switches must control the protocol operations. To select the switch that controls the centralized vLAG function, perform role election. The switch with the primary role controls the centralized operation. Role election is non-preemptive. That is, a primary already exists, another switch that is coming up remains as secondary even if it can become primary based on the role election logic.

Role election is determined by comparing the local vLAG system priority and local system MAC address. The switch with the smaller priority value becomes the vLAG primary switch. If the priorities are the same, the switch with smaller system MAC address becomes the vLAG primary switch. You can configure vLAG priority to anything between <0-65535>. For the examples, the priority was left at the default value of 0.

Example 5-46 G8264tor_1 show vlag information output

```
vLAG Tier ID: 2
vLAG system MAC: 08:17:f4:c3:dd:01
Local MAC 08:17:f4:33:9d:00 Priority 0 Admin Role SECONDARY (Operational Role
SECONDARY)
Peer MAC 08:17:f4:33:75:00 Priority 0
Health local 1.1.1.1 peer 1.1.1.2 State UP
ISL trunk id 67
ISL state Up
Startup Delay Interval: 120s (Finished)
```

vLAG 65: config with admin key 2000, associated trunk 65, state formed

vLAG 66: config with admin key 2002, associated trunk 66, state formed

Show vlag adminkey 2002

The output in Example 5-47 shows that the downstream vLAG towards the EN4093/R switches is formed and enabled by using LACP reference key 2002.

Example 5-47 G8264tor_1 show vlag adminkey 2002 output

```
vLAG is enabled on admin key 2002
Current LACP params for 25: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 26: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 27: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 28: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 37: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 38: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 39: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 40: active, Priority 32768, Admin Key 2002, Min-Links 1
```

Show vlag adminkey 2000

The output in Example 5-48 shows that the upstream vLAG towards the Nexus switches is formed and enabled by using LACP reference key 2000.

Example 5-48 G8264tor_1 show vlag adminkey 2000 output

```
vLAG is enabled on admin key 2000
Current LACP params for 18: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for 20: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for 22: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for 24: active, Priority 32768, Admin Key 2000, Min-Links 1
```

Show lacp information state up

Example 5-49 shows which ports are participating in an LACP aggregation, and which reference keys are used on those specific interfaces.

Example 5-49 G8264tor_1 show lacp information state up

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
1	active	1000	1000	yes	32768	1	67	up	1
2	active	1000	1000	yes	32768	1	67	up	1
3	active	1000	1000	yes	32768	1	67	up	1
4	active	1000	1000	yes	32768	1	67	up	1
5	active	1000	1000	yes	32768	1	67	up	1
6	active	1000	1000	yes	32768	1	67	up	1
7	active	1000	1000	yes	32768	1	67	up	1
8	active	1000	1000	yes	32768	1	67	up	1
9	active	1000	1000	yes	32768	1	67	up	1
10	active	1000	1000	yes	32768	1	67	up	1
11	active	1000	1000	yes	32768	1	67	up	1
12	active	1000	1000	yes	32768	1	67	up	1
13	active	1000	1000	yes	32768	1	67	up	1
14	active	1000	1000	yes	32768	1	67	up	1
15	active	1000	1000	yes	32768	1	67	up	1
16	active	1000	1000	yes	32768	1	67	up	1
18	active	2000	2000	yes	32768	20	65	up	1
20	active	2000	2000	yes	32768	20	65	up	1
22	active	2000	2000	yes	32768	20	65	up	1
24	active	2000	2000	yes	32768	20	65	up	1
25	active	2002	2002	yes	32768	26	66	up	1
26	active	2002	2002	yes	32768	26	66	up	1
27	active	2002	2002	yes	32768	26	66	up	1
28	active	2002	2002	yes	32768	26	66	up	1
37	active	2002	2002	yes	32768	26	66	up	1
38	active	2002	2002	yes	32768	26	66	up	1
39	active	2002	2002	yes	32768	26	66	up	1
40	active	2002	2002	yes	32768	26	66	up	1

Ping output for equipment on VLAN 4092

To verify connectivity, issue **ping** commands to devices on VLAN 4092 (Data VLAN) as shown in Example 5-50. IP address 10.4.1.10 represents a compute node with an operating system installed, flex_node1 on the Network Topology diagram.

Example 5-50 Ping verification for equipment on VLAN 4092

```
G8264TOR-1#ping 10.1.4.10 data-port
Connecting via DATA port.
[host 10.1.4.10, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl 255,
tos 0]
10.1.4.10: #1 ok, RTT 1 msec.
10.1.4.10: #2 ok, RTT 0 msec.
10.1.4.10: #3 ok, RTT 0 msec.
10.1.4.10: #4 ok, RTT 0 msec.
10.1.4.10: #5 ok, RTT 0 msec.
Ping finished.

G8264TOR-1#ping 10.1.4.249 data-port
```

```
Connecting via DATA port.
[host 10.1.4.249, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.249: #1 ok, RTT 1 msec.
10.1.4.249: #2 ok, RTT 0 msec.
10.1.4.249: #3 ok, RTT 1 msec.
10.1.4.249: #4 ok, RTT 0 msec.
10.1.4.249: #5 ok, RTT 0 msec.
Ping finished.
```

```
G8264TOR-1#ping 10.1.4.238 data-port
Connecting via DATA port.
[host 10.1.4.238, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.238: #1 ok, RTT 4 msec.
10.1.4.238: #2 ok, RTT 1 msec.
10.1.4.238: #3 ok, RTT 1 msec.
10.1.4.238: #4 ok, RTT 1 msec.
10.1.4.238: #5 ok, RTT 0 msec.
Ping finished.
```

Nexus output

This section lists output from the switch with hostname Nexus5548core_1. Similar or identical output exists for the switch with hostname Nexus5548core_2.

Show version

The output in Example 5-51 shows information about the switch and the associated code/firmware level.

Example 5-51 Nexus5548core_1 show version output

```
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Documents: http://www.cisco.com/en/US/products/ps9372/tsd\_products\_support\_series\_home.html
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software are covered under the GNU Public
License. A copy of the license is available at
http://www.gnu.org/licenses/gpl.html.
```

```
Software
  BIOS:      version 3.5.0
  loader:    version N/A
  kickstart: version 5.2(1)N1(1b)
  system:    version 5.2(1)N1(1b)
  power-seq: Module 1: version v1.0
             Module 3: version v5.0
  uC:        version v1.2.0.1
  SFP uC:    Module 1: v1.0.0.0
  BIOS compile time:      02/03/2011
  kickstart image file is: bootflash:///n5000-uk9-kickstart.5.2.1.N1.1b.bin
  kickstart compile time: 9/17/2012 11:00:00 [09/17/2012 18:38:53]
```

system image file is: bootflash:///n5000-uk9.5.2.1.N1.1b.bin
system compile time: 9/17/2012 11:00:00 [09/17/2012 20:38:22]

Hardware

cisco Nexus5548 Chassis ("02 32X10GE/Modular Universal Platform Supervisor")
Intel(R) Xeon(R) CPU with 8263848 kB of memory.
Processor Board ID FOC15424504

Device name: Nexus5548core_1
bootflash: 2007040 kB

Kernel uptime is 0 day(s), 22 hour(s), 32 minute(s), 3 second(s)

Last reset

Reason: Unknown
System version: 5.2(1)N1(1b)
Service:

plugin

Core Plugin, Ethernet Plugin

Show vlan

Example 5-52 displays the VLAN assignments for all of the ports on the switch.

Example 5-52 Nexus5548core_1 show vlan output

VLAN	Name	Status	Ports

1	default	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/11, Eth1/12 Eth1/13, Eth1/14, Eth1/15 Eth1/16, Eth1/18, Eth1/20 Eth1/22, Eth1/23, Eth1/24 Eth1/25, Eth1/26, Eth1/27 Eth1/28, Eth1/29, Eth1/30 Eth1/31, Eth1/32
1000	vPC_PEER_LINK	active	Eth1/21
4092	DATA_VLAN	active	Po5, Po100, Eth1/7, Eth1/8 Eth1/9, Eth1/10, Eth1/17 Eth1/19

Show interface status

Example 5-53 shows the full interface table, listing port status, speed, and so on.

Example 5-53 Nexus5548core_1 show interface status output

Port	Name	Status	Vlan	Duplex	Speed	Type

Eth1/1	--	sfpAbsent	1	full	10G	--
Eth1/2	--	sfpAbsent	1	full	10G	--
Eth1/3	--	sfpAbsent	1	full	10G	--
Eth1/4	--	sfpAbsent	1	full	10G	--
Eth1/5	--	sfpAbsent	1	full	10G	--

Eth1/6	--	sfpAbsent	1	full	10G	--
Eth1/7	VPC to G8264s	connected	trunk	full	10G	10Gbase-(un
Eth1/8	VPC to G8264s	connected	trunk	full	10G	10Gbase-(un
Eth1/9	VPC to G8264s	connected	trunk	full	10G	10Gbase-(un
Eth1/10	VPC to G8264s	connected	trunk	full	10G	10Gbase-(un
Eth1/11	--	sfpAbsent	1	full	10G	--
Eth1/12	--	sfpAbsent	1	full	10G	--
Eth1/13	--	sfpAbsent	1	full	10G	--
Eth1/14	--	sfpAbsent	1	full	10G	--
Eth1/15	--	sfpAbsent	1	full	10G	--
Eth1/16	--	sfpAbsent	1	full	10G	--
Eth1/17	vPC Peer link to N	connected	trunk	full	10G	10Gbase-(un
Eth1/18	--	sfpAbsent	1	full	10G	--
Eth1/19	vPC Peer link to N	connected	trunk	full	10G	10Gbase-(un
Eth1/20	--	sfpAbsent	1	full	10G	--
Eth1/21	vPC Keep alive	connected	1000	full	10G	10Gbase-(un
Eth1/22	--	sfpAbsent	1	full	10G	--
Eth1/23	--	sfpAbsent	1	full	10G	--
Eth1/24	--	sfpAbsent	1	full	10G	--
Eth1/25	--	sfpAbsent	1	full	10G	--
Eth1/26	--	sfpAbsent	1	full	10G	--
Eth1/27	--	sfpAbsent	1	full	10G	--
Eth1/28	--	sfpAbsent	1	full	10G	--
Eth1/29	--	sfpAbsent	1	full	10G	--
Eth1/30	--	sfpAbsent	1	full	10G	--
Eth1/31	--	sfpAbsent	1	full	10G	--
Eth1/32	--	sfpAbsent	1	full	10G	--
Po5	--	connected	trunk	full	10G	--
Po100	vPC Peer Link	connected	trunk	full	10G	--
mgmt0	--	connected	routed	full	1000	--

Show lldp neighbors

Example 5-54 lists the LLDP information and serves as a means to verify physical connectivity.

Example 5-54 Nexus5548core_1 show lldp neighbors output

Capability codes:

(R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device

(W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

Device ID	Local Intf	Hold-time	Capability	Port ID
G8264TOR-1	Eth1/7	120	BR	18
G8264TOR-1	Eth1/8	120	BR	20
G8264TOR-2	Eth1/9	120	BR	22
G8264TOR-2	Eth1/10	120	BR	24
Nexus5548core_2	Eth1/17	120	B	Eth1/17
Nexus5548core_2	Eth1/19	120	B	Eth1/19
Nexus5548core_2	Eth1/21	120	B	Eth1/21

Total entries displayed: 7

Show vpc

Example 5-55 shows output about the vPC feature in effect between the Nexus pair. In this example, the vPC peer link is established through Port-channel 100, and configuration consistency is exchanged over the vPC peer keep-alive link. Port-channel 5, vPC 5 is the downstream link aggregation group to the downstream G8264 pair, which is seen as a single entity by the Nexus pair.

Example 5-55 Nexus5548core_1 show vpc output

Legend:

(*) - local vPC is down, forwarding via vPC peer-link

```
vPC domain id          : 100
Peer status             : peer adjacency formed ok
vPC keep-alive status   : peer is alive
Configuration consistency status : success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role                : primary
Number of vPCs configured : 1
Peer Gateway            : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status    : Disabled
```

vPC Peer-link status

id	Port	Status	Active vlans
1	Po100	up	4092

vPC status

id	Port	Status	Consistency	Reason	Active vlans
5	Po5	up	success	success	4092

Show vpc peer-keepalive

Example 5-56 displays the status of the vPC peer-keepalive link.

Example 5-56 vPC peer-keepalive status

```
vPC keep-alive status      : peer is alive
--Peer is alive for        : (68229) seconds, (353) msec
--Send status              : Success
--Last send at             : 2012.10.16 20:19:46 950 ms
--Sent on interface        : Vlan1000
--Receive status           : Success
--Last receive at          : 2012.10.16 20:19:47 91 ms
--Received on interface    : Vlan1000
--Last update from peer    : (0) seconds, (454) msec
```

vPC Keep-alive parameters

```
--Destination      : 192.168.1.2
--Keepalive interval : 1000 msec
--Keepalive timeout  : 5 seconds
```



```
--Keepalive hold timeout      : 3 seconds
--Keepalive vrf               : VPCKeepAlive
--Keepalive udp port          : 3200
--Keepalive tos                : 192
```

Ping output for equipment on VLAN 4092

To verify connectivity, issue **ping** commands to the devices on VLAN 4092 (Data VLAN) as shown in Example 5-57.

Example 5-57 Ping verification for equipment on VLAN 4092

```
Nexus5548core_1# ping 10.1.4.243
PING 10.1.4.243 (10.1.4.243): 56 data bytes
64 bytes from 10.1.4.243: icmp_seq=0 ttl=254 time=1.008 ms
64 bytes from 10.1.4.243: icmp_seq=1 ttl=254 time=1.919 ms
64 bytes from 10.1.4.243: icmp_seq=2 ttl=254 time=0.856 ms
64 bytes from 10.1.4.243: icmp_seq=3 ttl=254 time=6.261 ms
64 bytes from 10.1.4.243: icmp_seq=4 ttl=254 time=9.596 ms

--- 10.1.4.243 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.856/3.928/9.596 ms

Nexus5548core_1# ping 10.1.4.238
PING 10.1.4.238 (10.1.4.238): 56 data bytes
64 bytes from 10.1.4.238: icmp_seq=0 ttl=254 time=7.571 ms
64 bytes from 10.1.4.238: icmp_seq=1 ttl=254 time=2.426 ms
64 bytes from 10.1.4.238: icmp_seq=2 ttl=254 time=0.817 ms
64 bytes from 10.1.4.238: icmp_seq=3 ttl=254 time=1.25 ms
64 bytes from 10.1.4.238: icmp_seq=4 ttl=254 time=5.628 ms

--- 10.1.4.238 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.817/3.538/7.571 ms

Nexus5548core_1# ping 10.1.4.10
PING 10.1.4.10 (10.1.4.10): 56 data bytes
64 bytes from 10.1.4.10: icmp_seq=0 ttl=63 time=0.586 ms
64 bytes from 10.1.4.10: icmp_seq=1 ttl=63 time=0.648 ms
64 bytes from 10.1.4.10: icmp_seq=2 ttl=63 time=0.561 ms
64 bytes from 10.1.4.10: icmp_seq=3 ttl=63 time=2.068 ms
64 bytes from 10.1.4.10: icmp_seq=4 ttl=63 time=9.057 ms

--- 10.1.4.10 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.561/2.583/9.057 ms
```

5.3.8 Full configuration files

This section displays the configuration of all of the devices in the Network Topology diagram.

EN4093flex-1

Example 5-58 shows the configuration for the EN4093flex-1 switch.

Example 5-58 EN4093-1 switch configuration file

```
version "7.3.1"
switch-type "IBM Flex System Fabric EN4093 10Gb Scalable Switch"
!
!

snmp-server name "en4093flex_1"
!
!
hostname "en4093flex_1"
!
!
interface port INTA1
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port INTB1
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT4
    name "ISL hlthchk"
    pvid 4000
    exit
!
interface port EXT7
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT8
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT9
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT10
    name "ISL"
    tagging
    pvid 4094
    exit
```

```

!
interface port EXT15
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT16
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT17
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT18
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT19
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT20
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT21
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT22
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092

```

```

    exit
!
vlan 1
    member INTA2-INTA14,INTB2-INTB14,EXT1-EXT3,EXT5-EXT6
    no member INTA1,INTB1,EXT4,EXT7-EXT10,EXT15-EXT22
!
vlan 4000
    enable
    name "ISL hlthchk"
    member EXT4
!
vlan 4092
    enable
    name "DATA"
    member INTA1,INTB1,EXT7-EXT10,EXT15-EXT22
!
vlan 4094
    enable
    name "ISL"
    member EXT7-EXT10
!
!
spanning-tree stp 125 vlan 4000
!
spanning-tree stp 126 vlan 4092
!
no spanning-tree stp 127 enable
spanning-tree stp 127 vlan 4094
!
!
interface port EXT7
    lacp mode active
    lacp key 1000
!
interface port EXT8
    lacp mode active
    lacp key 1000
!
interface port EXT9
    lacp mode active
    lacp key 1000
!
interface port EXT10
    lacp mode active
    lacp key 1000
!
interface port EXT15
    lacp mode active
    lacp key 2000
!
interface port EXT16
    lacp mode active
    lacp key 2000
!
interface port EXT17

```

```

        lacp mode active
        lacp key 2000
    !
interface port EXT18
    lacp mode active
    lacp key 2000
!
interface port EXT19
    lacp mode active
    lacp key 2000
!
interface port EXT20
    lacp mode active
    lacp key 2000
!
interface port EXT21
    lacp mode active
    lacp key 2000
!
interface port EXT22
    lacp mode active
    lacp key 2000
!
failover enable
failover trigger 1 mmon monitor admin-key 2000
failover trigger 1 mmon control member INTA1-INTB14
failover trigger 1 enable
!
!
!
vlag enable
vlag tier-id 1
vlag isl vlan 4094
vlag hlthchk peer-ip 1.1.1.2
vlag isl adminkey 1000
vlag adminkey 2000 enable
!
!
!
!
!
!
!
!
lldp enable
!
interface ip 40
    ip address 1.1.1.1 255.255.255.0
    vlan 4000
    enable
    exit
!
interface ip 92
    ip address 10.1.4.238 255.255.255.0

```

```

    vlan 4092
    enable
    exit
!
!
!
!
!
ntp enable
ntp ipv6 primary-server fe80::211:25ff:fec3:9b69 MGT
ntp interval 15
ntp authenticate
ntp primary-key 8811
!
ntp message-digest-key 8811 md5-ekey
1e389d20083088209635f6e3cb802bd2b52a41c0125c9904874d06d2a3af9d16341b4054daa0d14523
ca25ad2e9ec7d8ef2248b85c18a59a2436918a0ee41cea
!
ntp trusted-key 8811
!
end

```

EN4093flex_2

Example 5-59 lists the configuration for the EN4093flex_2 switch.

Example 5-59 EN4093flex_2 switch configuration

```

version "7.3.1"
switch-type "IBM Flex System Fabric EN4093 10Gb Scalable Switch"
!
!

snmp-server name "en4093flex_2"
!
!
hostname "en4093flex_2"
!
!
interface port INTA1
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port INTB1
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT4
    name "ISL hlthchk"
    pvid 4000
    exit
!

```

```

interface port EXT7
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT8
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT9
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT10
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT15
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT16
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT17
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT18
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT19
    name "Link to g8264tor_2"
    tagging

```

```

        tag-pvid
        pvid 4092
        exit
    !
interface port EXT20
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT21
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT22
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
vlan 1
    member INTA2-INTA14,INTB2-INTB14,EXT1-EXT3,EXT5-EXT6
    no member INTA1,INTB1,EXT4,EXT7-EXT10,EXT15-EXT22
!
vlan 4000
    enable
    name "ISL hlthchk"
    member EXT4
!
vlan 4092
    enable
    name "DATA"
    member INTA1,INTB1,EXT7-EXT10,EXT15-EXT22
!
vlan 4094
    enable
    name "ISL"
    member EXT7-EXT10
!
!
spanning-tree stp 125 vlan 4000
!
spanning-tree stp 126 vlan 4092
!
no spanning-tree stp 127 enable
spanning-tree stp 127 vlan 4094
!
!
no logging console
!

```



```

interface port EXT7
    lacp mode active
    lacp key 1000
!
interface port EXT8
    lacp mode active
    lacp key 1000
!
interface port EXT9
    lacp mode active
    lacp key 1000
!
interface port EXT10
    lacp mode active
    lacp key 1000
!
interface port EXT15
    lacp mode active
    lacp key 2000
!
interface port EXT16
    lacp mode active
    lacp key 2000
!
interface port EXT17
    lacp mode active
    lacp key 2000
!
interface port EXT18
    lacp mode active
    lacp key 2000
!
interface port EXT19
    lacp mode active
    lacp key 2000
!
interface port EXT20
    lacp mode active
    lacp key 2000
!
interface port EXT21
    lacp mode active
    lacp key 2000
!
interface port EXT22
    lacp mode active
    lacp key 2000
!
failover enable
failover trigger 1 mmon monitor admin-key 2000
failover trigger 1 mmon control member INTA1-INTB14
failover trigger 1 enable
!
!
!

```

```

vlag enable
vlag tier-id 1
vlag isl vlan 4094
vlag hlthchk peer-ip 1.1.1.1
vlag isl adminkey 1000
vlag adminkey 2000 enable
!
!
!
!
!
!
!
!
!
!
lldp enable
!
interface ip 40
    ip address 1.1.1.2 255.255.255.0
    vlan 4000
    enable
    exit
!
interface ip 92
    ip address 10.1.4.239 255.255.255.0
    vlan 4092
    enable
    exit
!
!
!
!
!
ntp enable
ntp ipv6 primary-server fe80::211:25ff:fec3:9b69 MGT
ntp interval 15
ntp authenticate
ntp primary-key 8811
!
ntp message-digest-key 8811 md5-ekey
ef9d8bb6cf808aa2b6b6e2f70c3029501c9b293eb41d60e5ebbd0fbbd72171ed3c867d24b9976e2052
771345e26681dc63a675b9033673c9923707f9d0f1c078
!
ntp trusted-key 8811
!
end

```

G8264tor_1

Example 5-60 shows the configuration for the G8264tor_1 switch.

Example 5-60 G8264tor_1 switch configuration

```

version "7.4.1"
switch-type "IBM Networking Operating System RackSwitch G8264"
!

```

```

!
ssh enable
!

!
!
no system dhcp
no system default-ip mgt
hostname "G8264TOR-1"
!
!
interface port 1
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 2
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 3
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 4
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 5
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 6
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 7
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 8
    name "ISL"

```

```

        tagging
        pvid 4094
        exit
    !
interface port 9
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 10
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 11
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 12
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 13
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 14
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 15
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 16
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 18
    name "VLAG to Nexus5548Core_1"
    tagging

```

```

        tag-pvid
        pvid 4092
        exit
    !
interface port 20
    name "VLAG to Nexus5548Core_1"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 22
    name "VLAG to Nexus5548Core_2"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 24
    name "VLAG to Nexus5548Core_2"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 25
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 26
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 27
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 28
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 37
    name "Link to EN4093-2"

```

```

        tagging
        tag-pvid
        pvid 4092
        exit
    !
interface port 38
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 39
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 40
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 64
    name "ISL hlthchk"
    pvid 4000
    exit
    !
vlan 1
    member 17-63
    no member 1-16,64
    !
vlan 4000
    enable
    name "ISL hlthchk"
    member 64
    !
vlan 4092
    enable
    name "DATA"
    member 1-16,18,20,22,24-28,37-40
    !
vlan 4094
    enable
    name "ISL"
    member 1-16
    !
    !
    !
spanning-tree stp 125 vlan 4000
    !
spanning-tree stp 126 vlan 4092

```

```

!
no spanning-tree stp 127 enable
spanning-tree stp 127 vlan 4094
!
!
interface port 1
    lacp mode active
    lacp key 1000
!
interface port 2
    lacp mode active
    lacp key 1000
!
interface port 3
    lacp mode active
    lacp key 1000
!
interface port 4
    lacp mode active
    lacp key 1000
!
interface port 5
    lacp mode active
    lacp key 1000
!
interface port 6
    lacp mode active
    lacp key 1000
!
interface port 7
    lacp mode active
    lacp key 1000
!
interface port 8
    lacp mode active
    lacp key 1000
!
interface port 9
    lacp mode active
    lacp key 1000
!
interface port 10
    lacp mode active
    lacp key 1000
!
interface port 11
    lacp mode active
    lacp key 1000
!
interface port 12
    lacp mode active
    lacp key 1000
!
interface port 13
    lacp mode active

```

```

        lacp key 1000
    !
interface port 14
    lacp mode active
    lacp key 1000
    !
interface port 15
    lacp mode active
    lacp key 1000
    !
interface port 16
    lacp mode active
    lacp key 1000
    !
interface port 18
    lacp mode active
    lacp key 2000
    !
interface port 20
    lacp mode active
    lacp key 2000
    !
interface port 22
    lacp mode active
    lacp key 2000
    !
interface port 24
    lacp mode active
    lacp key 2000
    !
interface port 25
    lacp mode active
    lacp key 2002
    !
interface port 26
    lacp mode active
    lacp key 2002
    !
interface port 27
    lacp mode active
    lacp key 2002
    !
interface port 28
    lacp mode active
    lacp key 2002
    !
interface port 37
    lacp mode active
    lacp key 2002
    !
interface port 38
    lacp mode active
    lacp key 2002
    !
interface port 39

```



```

        lacp mode active
        lacp key 2002
    !
interface port 40
    lacp mode active
    lacp key 2002
    !
    !
    !
vlag enable
vlag tier-id 2
vlag isl vlan 4094
vlag hlthchk peer-ip 1.1.1.2
vlag isl adminkey 1000
vlag adminkey 2000 enable
vlag adminkey 2002 enable
    !
    !
    !
    !
    !
    !
    !
    !
    !
interface ip 40
    ip address 1.1.1.1 255.255.255.0
    vlan 4000
    enable
    exit
    !
interface ip 92
    ip address 10.1.4.243 255.255.255.0
    vlan 4092
    enable
    exit
    !
interface ip 128
    ip address 172.25.101.243 255.255.0.0
    enable
    exit
    !
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
    !
    !
    !
    !
    !
end

```

G8264tor_2

Example 5-61 shows the configuration for the G8264tor_2 switch.

Example 5-61 G8264tor_2 switch configuration

```
version "7.4.1"
switch-type "IBM Networking Operating System RackSwitch G8264"
!
!
ssh enable
!

!
!
no system dhcp
no system default-ip mgt
hostname "G8264TOR-2"
!
!
interface port 1
    name "ISL"
    tagging
    exit
!
interface port 2
    name "ISL"
    tagging
    exit
!
interface port 3
    name "ISL"
    tagging
    exit
!
interface port 4
    name "ISL"
    tagging
    exit
!
interface port 5
    name "ISL"
    tagging
    exit
!
interface port 6
    name "ISL"
    tagging
    exit
!
interface port 7
    name "ISL"
    tagging
    exit
!
interface port 8
    name "ISL"
```

```

        tagging
        exit
    !
interface port 9
    name "ISL"
    tagging
    exit
    !
interface port 10
    name "ISL"
    tagging
    exit
    !
interface port 11
    name "ISL"
    tagging
    exit
    !
interface port 12
    name "ISL"
    tagging
    exit
    !
interface port 13
    name "ISL"
    tagging
    exit
    !
interface port 14
    name "ISL"
    tagging
    exit
    !
interface port 15
    name "ISL"
    tagging
    exit
    !
interface port 16
    name "ISL"
    tagging
    exit
    !
interface port 18
    name "VLAG to Nexus5548Core_2"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 20
    name "VLAG to Nexus5548Core_2"
    tagging
    tag-pvid
    pvid 4092

```

```

        exit
    !
interface port 22
    name "VLAG to Nexus5548Core_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 24
    name "VLAG to Nexus5548Core_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 25
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 26
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 27
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 28
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 37
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 38
    name "Link to EN4093-2"
    tagging
    tag-pvid

```

```

        pvid 4092
    exit
!
interface port 39
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 40
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 64
    name "ISL hlthchk"
    pvid 4000
    exit
!
vlan 1
    member 1-63
    no member 64
!
vlan 4000
    enable
    name "ISL hlthchk"
    member 64
!
vlan 4092
    enable
    name "DATA"
    member 1-16,18,20,22,24-28,37-40
!
vlan 4094
    enable
    name "ISL"
    member 1-16
!
!
!
spanning-tree stp 125 vlan 4000
!
spanning-tree stp 126 vlan 4092
!
no spanning-tree stp 127 enable
spanning-tree stp 127 vlan 4094
!
!
interface port 1
    lacp mode active
    lacp key 1000
!

```

```

interface port 2
    lacp mode active
    lacp key 1000
!
interface port 3
    lacp mode active
    lacp key 1000
!
interface port 4
    lacp mode active
    lacp key 1000
!
interface port 5
    lacp mode active
    lacp key 1000
!
interface port 6
    lacp mode active
    lacp key 1000
!
interface port 7
    lacp mode active
    lacp key 1000
!
interface port 8
    lacp mode active
    lacp key 1000
!
interface port 9
    lacp mode active
    lacp key 1000
!
interface port 10
    lacp mode active
    lacp key 1000
!
interface port 11
    lacp mode active
    lacp key 1000
!
interface port 12
    lacp mode active
    lacp key 1000
!
interface port 13
    lacp mode active
    lacp key 1000
!
interface port 14
    lacp mode active
    lacp key 1000
!
interface port 15
    lacp mode active
    lacp key 1000

```

```
!  
interface port 16  
    lacp mode active  
    lacp key 1000  
!  
interface port 18  
    lacp mode active  
    lacp key 2000  
!  
interface port 20  
    lacp mode active  
    lacp key 2000  
!  
interface port 22  
    lacp mode active  
    lacp key 2000  
!  
interface port 24  
    lacp mode active  
    lacp key 2000  
!  
interface port 25  
    lacp mode active  
    lacp key 2002  
!  
interface port 26  
    lacp mode active  
    lacp key 2002  
!  
interface port 27  
    lacp mode active  
    lacp key 2002  
!  
interface port 28  
    lacp mode active  
    lacp key 2002  
!  
interface port 37  
    lacp mode active  
    lacp key 2002  
!  
interface port 38  
    lacp mode active  
    lacp key 2002  
!  
interface port 39  
    lacp mode active  
    lacp key 2002  
!  
interface port 40  
    lacp mode active  
    lacp key 2002  
!  
!  
!
```

```

vlag enable
vlag tier-id 2
vlag isl vlan 4094
vlag hlthchk peer-ip 1.1.1.1
vlag isl adminkey 1000
vlag adminkey 2000 enable
vlag adminkey 2002 enable
!
!
!
!
!
!
!
!
!
!
interface ip 40
    ip address 1.1.1.2 255.255.255.0
    vlan 4000
    enable
    exit
!
interface ip 92
    ip address 10.1.4.244 255.255.255.0
    vlan 4092
    enable
    exit
!
interface ip 128
    ip address 172.25.101.244 255.255.0.0
    enable
    exit
!
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
!
!
!
!
!
!
end

```

Nexus5548core_1 switch

Example 5-62 shows the configuration of the Nexus5548core_1 switch.

Example 5-62 Nexus5548core_1 switch configuration

```

!Command: show startup-config
!Time: Tue Oct 16 20:43:19 2012
!Startup config saved at: Tue Oct 16 20:42:45 2012

version 5.2(1)N1(1b)

```



```

logging level feature-mgr 0
hostname Nexus5548core_1

feature telnet
cfs ipv4 distribute
cfs eth distribute
feature interface-vlan
feature lacp
feature vpc
feature lldp

username admin password 5 $1$huQeFTJf$dYim2oGvqYAGk3THH5KP.0 role network-admin

banner motd #Nexus 5000 Switch
#

no ip domain-lookup
class-map type qos class-fcoe
class-map type queuing class-fcoe
    match qos-group 1
class-map type queuing class-all-flood
    match qos-group 2
class-map type queuing class-ip-multicast
    match qos-group 2
class-map type network-qos class-fcoe
    match qos-group 1
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
snmp-server user admin network-admin auth md5 0x50d80b5959ad2a911a11fcaa8453db8a
priv 0x50d80b5959ad2a911a11fcaa8453db8a localizedkey

vrf context management
    ip route 0.0.0.0/0 172.25.1.1
vrf context VPCKeepAlive
vlan 1
vlan 1000
    name vPC_PEER_LINK
vlan 4092
    name DATA_VLAN
spanning-tree vlan 1000 priority 24576
spanning-tree vlan 4092 priority 8192
vpc domain 100
    role priority 1000
    peer-keepalive destination 192.168.1.2 source 192.168.1.1 vrf VPCKeepAlive
    delay restore 150
port-profile default max-ports 512

interface Vlan1

interface Vlan1000
    no shutdown
    vrf member VPCKeepAlive

```

```

ip address 192.168.1.1/30

interface Vlan4092
  no shutdown
  ip address 10.1.4.249/24

interface port-channel5
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  speed auto
  vpc 5

interface port-channel100
  description vPC Peer Link
  switchport mode trunk
  switchport trunk allowed vlan 4092
  spanning-tree port type network
  vpc peer-link

interface Ethernet1/1

interface Ethernet1/2

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7
  description VPC to G8264s
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  channel-group 5 mode active

interface Ethernet1/8
  description VPC to G8264s
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  channel-group 5 mode active

interface Ethernet1/9
  description VPC to G8264s
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  channel-group 5 mode active

interface Ethernet1/10
  description VPC to G8264s

```

```
    switchport mode trunk
    switchport trunk native vlan 4092
    switchport trunk allowed vlan 4092
    channel-group 5 mode active

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17
    description vPC Peer link to Nexus5548core_2
    switchport mode trunk
    switchport trunk allowed vlan 4092
    speed auto
    channel-group 100 mode active

interface Ethernet1/18

interface Ethernet1/19
    description vPC Peer link to Nexus5548core_2
    switchport mode trunk
    switchport trunk allowed vlan 4092
    speed auto
    channel-group 100 mode active

interface Ethernet1/20

interface Ethernet1/21
    description vPC Keep alive
    switchport access vlan 1000

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29
```

```

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface mgmt0
  ip address 172.25.101.249/16
cli alias name wr copy run start
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.2.1.N1.1b.bin
boot system bootflash:/n5000-uk9.5.2.1.N1.1b.bin

```

Nexus5548core_2 switch

Example 5-63 shows the configuration of the Nexus5548core_2 switch.

Example 5-63 Nexus5548core_2 switch configuration

```

!Command: show startup-config
!Time: Tue Oct 16 20:05:31 2012
!Startup config saved at: Tue Oct 16 20:05:24 2012

version 5.2(1)N1(1b)
logging level feature-mgr 0
hostname Nexus5548core_2

feature telnet
cfs ipv4 distribute
cfs eth distribute
feature interface-vlan
feature lacp
feature vpc
feature lldp

username admin password 5 $1$W5m0kb.B$kFgCTs1WQy/ElfbozmrDt/ role network-admin

banner motd #Nexus 5000 Switch
#

no ip domain-lookup
class-map type qos class-fcoe
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2

```

```
snmp-server user admin network-admin auth md5 0xf6e8ccc23aa981dc5c6c28cfa16eb886
priv 0xf6e8ccc23aa981dc5c6c28cfa16eb886 localizedkey
```

```
vrf context management
  ip route 0.0.0.0/0 172.25.1.1
vrf context VPCKeepAlive
vlan 1
vlan 1000
  name vPC_PEER_LINK
vlan 4092
  name DATA_VLAN
spanning-tree vlan 4092 priority 16384
vpc domain 100
  peer-keepalive destination 192.168.1.1 source 192.168.1.2 vrf VPCKeepAlive
port-profile default max-ports 512
```

```
interface Vlan1
```

```
interface Vlan1000
  no shutdown
  vrf member VPCKeepAlive
  ip address 192.168.1.2/30
```

```
interface Vlan4092
  no shutdown
  ip address 10.1.4.200/24
```

```
interface port-channel5
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  speed auto
  vpc 5
```

```
interface port-channel100
  description vPC Peer Link
  switchport mode trunk
  switchport trunk allowed vlan 4092
  spanning-tree port type network
  vpc peer-link
```

```
interface Ethernet1/1
```

```
interface Ethernet1/2
```

```
interface Ethernet1/3
```

```
interface Ethernet1/4
```

```
interface Ethernet1/5
```

```
interface Ethernet1/6
```

```
interface Ethernet1/7
```

```

description VPC to G8264s
switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
channel-group 5 mode active

interface Ethernet1/8
description VPC to G8264s
switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
channel-group 5 mode active

interface Ethernet1/9
description VPC to G8264s
switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
channel-group 5 mode active

interface Ethernet1/10
description VPC to G8264s
switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
channel-group 5 mode active

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17
description vPC Peer link to Nexus5548core_1
switchport mode trunk
switchport trunk allowed vlan 4092
speed auto
channel-group 100 mode active

interface Ethernet1/18

interface Ethernet1/19
description vPC Peer link to Nexus5548core_1
switchport mode trunk
switchport trunk allowed vlan 4092
speed auto
channel-group 100 mode active

```

```
interface Ethernet1/20

interface Ethernet1/21
  description vPC Keep alive
  switchport access vlan 1000

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet2/1

interface Ethernet2/2

interface Ethernet2/3

interface Ethernet2/4

interface Ethernet2/5

interface Ethernet2/6

interface Ethernet2/7

interface Ethernet2/8

interface Ethernet2/9

interface Ethernet2/10

interface Ethernet2/11

interface Ethernet2/12

interface Ethernet2/13

interface Ethernet2/14
```

```
interface Ethernet2/15

interface Ethernet2/16

interface mgmt0
  ip address 172.25.101.200/16

interface loopback1
  ip address 192.168.1.1/24
cli alias name wr copy run start
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.2.1.N1.1b.bin
boot system bootflash:/n5000-uk9.5.2.1.N1.1b.bin
```

5.4 Fully redundant with traditional spanning-tree

This section details the implementation of a fully redundant configuration that uses a traditional spanning-tree.

5.4.1 Topology and requirements

This implementation scenario uses a more traditional, classic network design with the spanning-tree protocol that serves as a protection against bridge or L2 loops. If you use upstream Cisco equipment, you might not be able to aggregate from a virtualized standpoint. For more information, see *Cisco Catalyst 6500 Virtual Switching System*, or *Cisco Virtual PortChannel on the Nexus platform*. If you are more comfortable with STP, you can choose this implementation scenario.

This approach has the following advantages:

- ▶ Almost ready to use if Per VLAN Rapid Spanning Tree protocol (PVRST+) is used on both Cisco (default selection in NX-OS) equipment and IBM equipment (default selection as of recent software versions of IBM Networking OS)
- ▶ Does not require extra steps or implementation experience in switch virtualization features and functionality to begin implementation
- ▶ Can be done with almost any datacenter-class upstream Cisco switch

This approach has the following disadvantages:

- ▶ Links are blocked by spanning-tree to prevent bridging loops, wasting valuable bandwidth
- ▶ Can require longer convergence times during a link failure
- ▶ Troubleshooting problems with spanning-tree can be more difficult for less experienced network architects

5.4.2 Components used

The following components are used in the example configuration:

- ▶ Cisco Nexus 5548UP (Qty. 2)
- ▶ IBM G8264 RackSwitch (Qty. 2)
- ▶ IBM Flex System Fabric EN4093/R 10Gb Scalable Switch (Qty. 2)

5.4.3 Network diagram and physical setup

Figure 5-5 shows the Network Topology diagram for the fully redundant scenario with spanning tree.

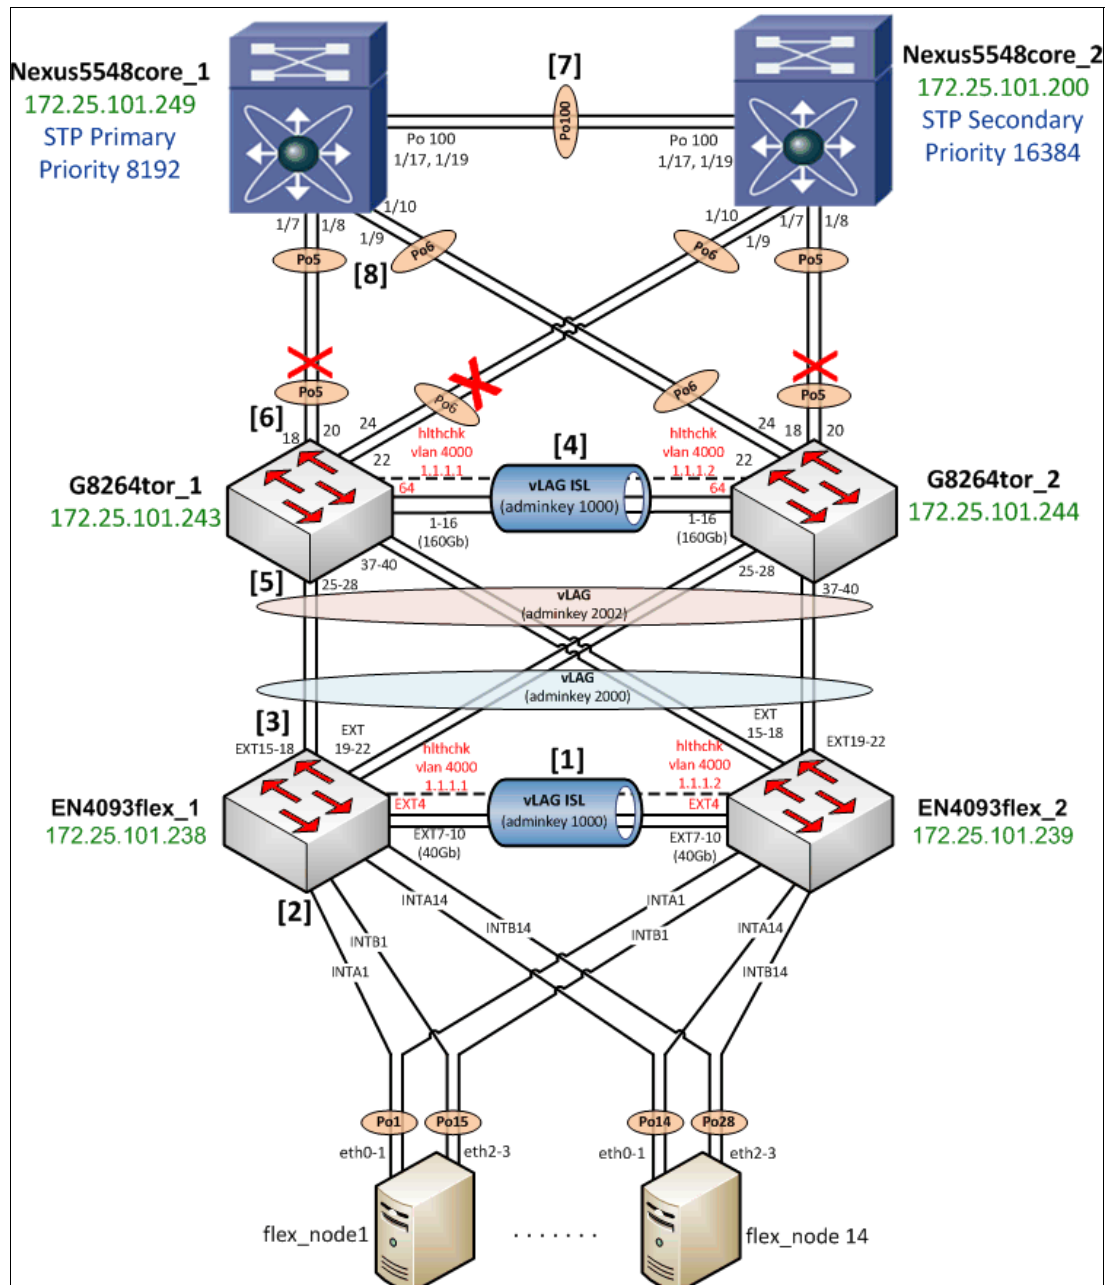


Figure 5-5 Network Topology diagram for fully redundant scenario using spanning-tree

Verify the physical cabling between the EN4093flex switches and G8264tor switches. The lab environment used four IBM QSFP+ DAC Break Out Cables from the EN4093/R switches to the upstream G8264s. This configuration requires that the EN4093/R switches be licensed for these particular features so that the ports can be used.

- ▶ Four 1m IBM QSFP+-to-QSFP+ Cables were used to form the 160 Gb ISL between the G8264 switches.
- ▶ 10Gb SFP+ DAC cables were used for all other connections in the diagram.

5.4.4 EN4093flex_1 configuration

Begin the implementation of this scenario on the IBM Flex System Fabric EN4093/R switches, then work up the diagram in Figure 5-5 on page 123. Each step provides the commands necessary, and lists the step number from the diagram.

General configuration

1. Create the ISL Healthcheck, ISL data, and Data VLANs as shown in Example 5-64, giving them descriptive names, assigning them to spanning-tree groups, and enabling them. You can elect to allow the switch itself to create STP instances for you. In this example, they were manually created instead.

Example 5-64 Create ISL hlthchk, Data, and ISL VLANs on EN4093flex_1

```
configure terminal
vlan 4000
    enable
    name "ISL hlthchk"
    stg 125
    exit
vlan 4092
    enable
    name "DATA"
    stg 126
    exit
vlan 4094
    enable
    name "ISL"
    stg 127
    exit
```

2. Assign IP addresses for both the ISL Healthcheck and Data VLANs in shown in Example 5-65. Doing so allows you to verify connectivity between the various pieces of equipment when verifying the configuration. In this example, interface ip 40 represents the vLAG Health Check IP address, and interface ip 92 represents an address on the Data VLAN that uses the prefix 10.1.4. The last octet is borrowed from the network diagram's Management address to aid in the identification of which piece of equipment you are verifying connectivity to.

Example 5-65 Creating IP interfaces and assigning VLANs and IP addresses on EN4093flex_1

```
configure terminal
interface ip 40
    ip address 1.1.1.1 255.255.255.0
    vlan 4000
    enable
    exit
```

```
interface ip 92
  ip address 10.1.4.238 255.255.255.0
  vlan 4092
  enable
  exit
```

Configuring ISL between EN4093flex switches (step 1)

3. Configure the eventual ISL in Example 5-66 between the EN4093/R switches by configuring them to have a default (untagged) VLAN of 4094. Set an LACP key of 1000 to bundle the ports together in an aggregation, with 802.1q tagging enabled so that L2 VLAN traffic can traverse the ISL. Carry Data VLAN 4092 over these links.

Example 5-66 Initial ISL configuration on EN4093flex_1

```
configure terminal
interface port ext7-ext10
  pvid 4094
  tagging
  exit
vlan 4092
  member ext7-ext10
  exit
interface port ext7-ext10
  lacp key 1000
  lacp mode active
  exit
```

4. Create the dedicated health check VLAN and physical interface in Example 5-67 to be used for heartbeats between the EN4093/R switches. In this example, EXT4 was chosen as a dedicated interface and VLAN 4000 to serve as the health check for the ISL.

Example 5-67 Creating vLAG health check on EN4093flex_1

```
configure terminal
vlan 4000
  name "ISL hlthchk"
  enable
  exit
interface port ext4
  pvid 4000
  exit
```

5. Disable STP between the EN4093/R switches and activate a vLAG between them so that they appear as a single entity to upstream and downstream infrastructure as shown in Example 5-68. Reference the LACP key that was configured in the previous step.

Example 5-68 Disabling STP and activating ISL vLAG on EN4093flex_1

```
configure terminal
no spanning-tree stp 127 enable
vlag tier-id 1
vlag isl vlan 4094
vlag isl adminkey 1000
vlag hlthchk peer-ip 1.1.1.2
vlag enable
```

Configuring downstream internal node ports (step 2)

6. Configure downstream node interfaces in Example 5-69 to have a default (untagged) VLAN of 4092 (data VLAN), with 802.1q tagging enabled. Add the ability for all member ports to be on VLAN 4092.

Example 5-69 Downstream Internal node port configuration, on EN4093flex_1

```
configure terminal
interface port inta1-intb14
    pvid 4092
    tagging
    spanning-tree edge
    exit
vlan 4092
    member inta1-intb14
exit
```

7. For redundancy, create two port-channels on each of the 14 nodes. Each port channel aggregates two ports, one from each EN4093flex switch. Port channels 1-14 match the “A” internally labeled ports, and port channels 15-28 match the “B” ports as shown in Example 5-70.

Example 5-70 Node-facing port channel creation and vLAG activation on EN4093flex_1

```
configure terminal
portchannel 1 port inta1
portchannel 1 enable
vlag portchannel 1 enable
portchannel 15 port intb1
portchannel 15 enable
vlag portchannel 15 enable
portchannel 2 port inta2
portchannel 2 enable
vlag portchannel 2 enable
portchannel 16 port intb2
portchannel 16 enable
vlag portchannel 16 enable
portchannel 3 port inta3
portchannel 3 enable
vlag portchannel 3 enable
portchannel 17 port intb3
portchannel 17 enable
vlag portchannel 17 enable
portchannel 4 port inta4
portchannel 4 enable
vlag portchannel 4 enable
portchannel 18 port intb4
portchannel 18 enable
vlag portchannel 18 enable
portchannel 5 port inta5
portchannel 5 enable
vlag portchannel 5 enable
portchannel 19 port intb5
portchannel 19 enable
vlag portchannel 19 enable
portchannel 6 port inta6
```

```
portchannel 6 enable
vlag portchannel 6 enable
portchannel 20 port intb6
portchannel 20 enable
vlag portchannel 20 enable
portchannel 7 port inta7
portchannel 7 enable
vlag portchannel 7 enable
portchannel 21 port intb7
portchannel 21 enable
vlag portchannel 21 enable
portchannel 8 port inta8
portchannel 8 enable
vlag portchannel 8 enable
portchannel 22 port intb8
portchannel 22 enable
vlag portchannel 22 enable
portchannel 9 port inta9
portchannel 9 enable
vlag portchannel 9 enable
portchannel 23 port intb9
portchannel 23 enable
vlag portchannel 23 enable
portchannel 10 port inta10
portchannel 10 enable
vlag portchannel 10 enable
portchannel 24 port intb10
portchannel 24 enable
vlag portchannel 24 enable
portchannel 11 port inta11
portchannel 11 enable
vlag portchannel 11 enable
portchannel 25 port intb11
portchannel 25 enable
vlag portchannel 25 enable
portchannel 12 port inta12
portchannel 12 enable
vlag portchannel 12 enable
portchannel 26 port intb12
portchannel 26 enable
vlag portchannel 26 enable
portchannel 13 port inta13
portchannel 13 enable
vlag portchannel 13 enable
portchannel 27 port intb13
portchannel 27 enable
vlag portchannel 27 enable
portchannel 14 port inta14
portchannel 14 enable
vlag portchannel 14 enable
portchannel 28 port intb14
portchannel 28 enable
vlag portchannel 28 enable
```

Configuring upstream G8264tor facing ports and layer 2 failover (step 3)

8. Configure the upstream G8264tor facing ports in Example 5-71 with a default (untagged) VLAN of 4092 (data VLAN), tag the PVID, and use an LACP key of 2000 to bundle the ports together in an aggregation.

Example 5-71 Upstream G8264tor facing port configuration on EN4093flex_1

```
configure terminal
interface port ext15-ext22
    pvid 4092
    tagging
    tag-pvid
    exit
vlan 4092
    member ext15-ext22
    exit
interface port ext15-ext22
    lacp key 2000
    lacp mode active
    exit
```

9. Activate the vLAG feature for the upstream EN4093/R ports so that the G8264s see the EN4093s as a single, virtualized entity as shown in Example 5-72. Use adminkey 2000, which represents the LACP key that bundles ports EXT15-22 together as one.

Example 5-72 Activating the upstream vLAG to the G8264 switches on EN4093flex_1

```
configure terminal
vlag adminkey 2000 enable
```

10. Enable Layer-2 failover in Example 5-73, which shuts down the links to the compute nodes if the uplinks for the EN4093/R switch fail. This ensures that the downstream node is aware of the upstream failure and can fail traffic over to the other NIC in the node. The other NIC in the example is connected to the other EN4093/R switch in the Enterprise Chassis, ensuring that redundancy is maintained.

Example 5-73 Enabling L2 failover for the compute nodes on EN4093flex_1

```
configure terminal
failover trigger 1 mmon monitor admin-key 2000
failover trigger 1 mmon control member INTA1-INTB14
failover trigger 1 enable
failover enable
```

Repeat this configuration for EN4093flex_2 on the other I/O module. The only difference between the EN4093flex_1 switch and the EN4093flex_2 switch is the vLAG health check peer address and the Data, and ISL h1thchk VLAN IP addresses. To verify EN4093flex switch configuration, run the **show** commands that are outlined in 5.4.8, “Verification and show command output” on page 135.

5.4.5 G8264tor_1 configuration

Next, configure the RackSwitch G8264.

General configuration

1. Create the ISL Healthcheck, ISL data, and Data VLANs as shown in Example 5-74. Give them descriptive names, assign them to spanning-tree groups, and enable them.

Example 5-74 Creating ISL hlthchk, Data, and ISL VLANs on G8264tor_1

```
configure terminal
vlan 4000
    enable
    name "ISL hlthchk"
    stg 125
    exit
vlan 4092
    enable
    name "Data"
    stg 126
    exit
vlan 4094
    enable
    name "ISL"
    stg 127
    exit
```

2. Assign IP addresses for the ISL Healthcheck, Data VLANs, and management VLAN as shown in Example 5-75. Interface ip 128 represents the management IP address that is referenced in the Network Topology diagram. IP gateway 4 is the upstream router interface for the 172 management network.

Example 5-75 Creating IP interfaces and assigning VLANs and IP addresses on G8264tor_1

```
configure terminal
interface ip 40
    ip address 1.1.1.1 255.255.255.0
    vlan 4000
    enable
    exit
interface ip 92
    ip address 10.1.4.243 255.255.255.0
    vlan 4092
    enable
    exit
interface ip 128
    ip address 172.25.101.243 255.255.0.0
    enable
    exit
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
```

Configuring ISL between G8264tor switches (step 4)

3. Configure the ISL between the G8264tor switches as shown in Example 5-76. Make the default (untagged) VLAN 4094, LACP key of 1000 to bundle the ports together in an aggregation. Enable 802.1q tagging so that L2 VLAN traffic can traverse the ISL. Carry Data VLAN 4092 over these links.

Example 5-76 Initial ISL configuration on G8264tor_1

```
configure terminal
interface port 1-16
    pvid 4094
    tagging
    exit
vlan 4092
    member 1-16
    exit
interface port 1-16
    lacp key 1000
    lacp mode active
    exit
```

4. Disable STP between the G8264 switches and activate a vLAG between them so that they appear as a single entity to upstream and downstream infrastructure as shown in Example 5-77. Reference the LACP key that was configured in the previous step.

Example 5-77 Disabling STP and activating ISL vLAG on G8264tor_1

```
configure terminal
no spanning-tree stp 127 enable
vlag tier-id 2
vlag isl vlan 4094
vlag isl adminkey 1000
vlag hlthchk peer-ip 1.1.1.2
vlag enable
```

Configuring downstream EN4093flex facing ports (step 5)

5. Configure the downstream ports towards the EN4093/R switches in Example 5-78 to have a default (untagged) VLAN of 4092 (data VLAN), with 802.1q tagging enabled. Add the ability for all member ports to be on VLAN 4092. Bundle ports 25-28 and 37-40 together in an LACP aggregation.

Example 5-78 Downstream EN4093flex facing port configuration on G8264tor_1

```
configure terminal
interface port 25-28,37-40
    pvid 4092
    tagging
    tag-pvid
    exit
vlan 4092
    member 25-28,37-40
    exit
interface port 25-28,37-40
    lacp key 2002
```



```
lACP mode active
exit
```

6. Activate the vLAG for the downstream EN4093flex facing ports so that the 4093s see the G8264s as a single, virtualized entity as shown in Example 5-79. Use adminkey 2002, which represents the LACP key that bundles ports 25-28, and 37-40 together as one.

Example 5-79 Activating downstream EN4093flex facing vLAG on G8264tor_1

```
configure terminal
vlag adminkey 2002 enable
```

Configuring upstream Nexus5548core facing ports (step 6)

7. Configure the upstream ports to the Nexus5548core switches in Example 5-80 with a default (untagged) VLAN of 4092 (data VLAN), tag the PVID, and provide a useful description on the interfaces.

Example 5-80 Upstream Nexus5548core facing port configuration on G8264tor_1

```
configure terminal
interface port 18,20
    name "Po5 to Nexus5548core_1"
    pvid 4092
    tagging
    tag-pvid
    exit
interface port 22,24
    name "Po5 to Nexus5548core_2"
    pvid 4092
    tagging
    tag-pvid
    exit
vlan 4092
    member 18,20,22,24
    exit
```

8. Activate link aggregation groups using static port-channeling as shown in Example 5-81. The example uses static port-channeling to illustrate that IBM System Networking equipment inter-operates with an upstream Cisco infrastructure with either LACP or static (no negotiation protocol) port-channeling.

Example 5-81 Creating port-channel interfaces on G8264tor_1

```
configure terminal
portchannel 5 port 18
portchannel 5 port 20
portchannel 5 enable
!
portchannel 6 port 22
portchannel 6 port 24
portchannel 6 enable
```

Repeat this configuration on the other top of rack switch, G8264tor_2. The only difference between the G8264tor_1 switch and the G8264tor_2 switch is the vLAG health check peer address and the Data, management and ISL hltchk VLAN IP addresses. To verify G8264tor

switch configuration, run the **show** commands that are outlined in 5.4.8, “Verification and show command output” on page 135.

5.4.6 Nexus5548core_1 STP primary switch configuration

Configure the Cisco Nexus5548core_1 STP primary switch.

General configuration

9. Enable NX-OS feature sets as shown in Example 5-82.

Example 5-82 Enabling Cisco NX-OS feature set on Nexus5548core_1

```
configure terminal
feature interface-vlan
feature lacp
feature lldp
```

10. Create vlan 4092 (data VLAN). Set the spanning-tree priority for the data VLAN to be half that of the Nexus5548core_2 switch (8192). Because Nexus5548core_1 switch has a lower spanning-tree priority, it becomes the root bridge for L2 functionality as shown in Example 5-83.

Example 5-83 Data VLAN configuration and spanning-tree priority configuration on Nexus5548core_1

```
configure terminal
vlan 4092
  name DATA_VLAN
spanning-tree vlan 4092 priority 8192
```

11. Because the Nexus box has a Layer-3 card and license, create the SVIs for the Data VLAN, which are useful during verification of this scenario’s implementation (Example 5-84).

Example 5-84 Create IP address for vlan 4092 (Data vlan) on Nexus5548core_1

```
configure terminal
interface Vlan4092
  no shutdown
  ip address 10.1.4.249/24
```

Configuring switch-to-switch link between the Nexus switches (step 7)

12. Configure the physical interfaces that comprising the switch-to-switch link between the Nexus 5548-1 and 5548-2 switches as shown in Example 5-85. Use port-channel100 and LACP.

Example 5-85 Switch-to-switch link physical and logical interface configuration on Nexus5548core_1

```
configure terminal
interface Ethernet1/17
  description Po100 to Nexus5548core_2
  switchport mode trunk
  switchport trunk allowed vlan 4092
  channel-group 100 mode active
interface Ethernet1/19
```

```
description Po100 to Nexus5548core_2
switchport mode trunk
switchport trunk allowed vlan 4092
channel-group 100 mode active
interface port-channel100
description Switch-to-Switch link
switchport mode trunk
switchport trunk allowed vlan 4092
spanning-tree port type network
```

Configuring downstream G8264tor facing ports (step 8)

13. For the Nexus 5548 primary switch, configure the downstream physical and logical interfaces in Example 5-86. Bundle interfaces Ethernet1/7 and Ethernet1/8 in static aggregation Po5, and interfaces Ethernet1/9 and Ethernet1/10 in static aggregation Po6.

Example 5-86 Downstream G8264tor facing port configuration on Nexus5548core_1

```
configure terminal
interface Ethernet1/7-8
description Po5 to G8264tor_1
switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
channel-group 5 mode on
interface Ethernet1/9-10
description Po6 to G8264tor_2
switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
channel-group 6 mode on
interface port-channel5
switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
speed auto
interface port-channel6
switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
speed auto
```

5.4.7 Nexus5548core_2 STP secondary switch configuration

Configure the Cisco Nexus5548core_2 STP secondary switch.

General configuration

14. Enable NX-OS feature sets as shown in Example 5-87.

Example 5-87 Enable NX-OS feature sets on Nexus5548core_2

```
configure terminal
feature interface-vlan
```

```
feature lACP
feature lldp
```

15. Create vlan 4092 (data VLAN) as shown in Example 5-88. Configure the spanning-tree priority for the data VLAN to be twice that of Nexus5548core_1 (8192). Because Nexus5548core_2 switch has a lower spanning-tree priority than Nexus5548core_1, it becomes the backup for the spanning tree protocol layer 2 function.

Example 5-88 Data vlan (4092) creation and spanning-tree priority configuration on Nexus5548core_2

```
configure terminal
vlan 4092
  name DATA_VLAN
spanning-tree vlan 4092 priority 16384
```

16. Create the SVIs for the Data VLAN (4092) as shown in Example 5-89, which are useful during verification of this scenario's implementation.

Example 5-89 Data VLAN ip address configuration on Nexus5548core_2

```
configure terminal
interface Vlan4092
  no shutdown
  ip address 10.1.4.200/24
```

Configuring switch-to-switch link between Nexus switches (step 7)

17. Configure the physical interfaces that comprise the switch-to-switch link between the Nexus5548core_1 and Nexus5548core_2 switches as shown in Example 5-90. Use port-channel100 and LACP.

Example 5-90 Switch-to-switch link physical and logical interface configuration on Nexus5548core_2

```
configure terminal
interface Ethernet1/17
  description Po100 to Nexus5548core_1
  switchport mode trunk
  switchport trunk allowed vlan 4092
  channel-group 100 mode active
interface Ethernet1/19
  description Po100 to Nexus5548core_1
  switchport mode trunk
  switchport trunk allowed vlan 4092
  channel-group 100 mode active
interface port-channel100
  description Switch-to-Switch link
  switchport mode trunk
  switchport trunk allowed vlan 4092
  spanning-tree port type network
```

Configuring downstream G8264tor facing ports (step 8)

18. Finally, configure the downstream physical and logical interfaces as shown in Example 5-91. Bundle interfaces Ethernet1/7 and Ethernet1/8 in static aggregation Po5, and interfaces Ethernet1/9 and Ethernet1/10 in static aggregation Po6.

Example 5-91 Downstream G8264tor facing port configuration on Nexus5548core_2

```
configure terminal
interface Ethernet1/7-8
  description Po5 to G8264tor_2
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  channel-group 5 mode on
interface Ethernet1/9-10
  description Po6 to G8264tor_1
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  channel-group 6 mode on
interface port-channel5
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  speed auto
interface port-channel6
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  speed auto
```

5.4.8 Verification and show command output

The following section lists output from common **show** commands that can aid you in the implementation of this scenario. Ping verification of the IP addresses configured on the equipment for the Data VLAN is also done to show that all of the devices can reach each other successfully.

As in the implementation section, the helpful commands are described starting with the EN4093/R switches, and working up the Network Topology diagram to the Cisco Nexus pair.

EN4093/R output

This section lists output from the switch with hostname EN4093flex_1. Similar or identical output exists for the switch with hostname EN4093flex_2.

Show version

The command output in Example 5-92 shows information about the switch used and the associated code/firmware level.

Example 5-92 EN4093flex_1 show version output

```
System Information at 23:04:56 Fri Oct 12, 2012
Time zone: No timezone configured
Daylight Savings Time Status: Disabled
```

IBM Flex System Fabric EN4093 10Gb Scalable Switch

Switch has been up for 1 day, 2 hours, 1 minute and 21 seconds.
Last boot: 21:05:54 Thu Oct 11, 2012 (reset from Telnet/SSH)

MAC address: 6c:ae:8b:bf:6d:00 IP (If 40) address: 1.1.1.1
Internal Management Port MAC Address: 6c:ae:8b:bf:6d:ef
Internal Management Port IP Address (if 128): 172.25.101.238
External Management Port MAC Address: 6c:ae:8b:bf:6d:fe
External Management Port IP Address (if 127):
Software Version 7.3.1.0 (FLASH image1), active configuration.

Hardware Part Number : 49Y4272
Hardware Revision : 02
Serial Number : Y250VT24M099
Manufacturing Date (WWYY) : 1712
PCBA Part Number : BAC-00072-01
PCBA Revision : 0
PCBA Number : 00
Board Revision : 02
PLD Firmware Version : 1.5

Temperature Warning : 32 C (Warn at 60 C/Recover at 55 C)
Temperature Shutdown : 32 C (Shutdown at 65 C/Recover at 60 C)
Temperature Inlet : 27 C
Temperature Exhaust : 33 C

Power Consumption : 54.300 W (12.244 V, 4.435 A)

Switch is in I/O Module Bay 1

Show vlan

Example 5-93 shows output about VLAN assignment for all of the ports on the switch.

Example 5-93 EN4093flex_1 show vlan output

VLAN	Name	Status	MGT	Ports
----	-----	-----	----	-----
1	Default VLAN	ena	dis	EXT1-EXT3 EXT5 EXT6
4000	ISL hlthchk	ena	dis	EXT4
4092	DATA	ena	dis	INTA1-INTB14 EXT7-EXT10 EXT15-EXT22
4094	ISL	ena	dis	EXT7-EXT10
4095	Mgmt VLAN	ena	ena	EXTM MGT1

Show interface status

Because there is only one compute node in the chassis (in slot 1), all the other internal ports are listed as “down” from a link perspective in the output in Example 5-94.

Example 5-94 EN4093flex_1 show interface status output

Alias	Port	Speed	Duplex	Flow Ctrl		Link	Name
				--TX--	--RX--		
INTA1	1	1000	full	no	no	up	INTA1
INTA2	2	1G/10G	full	yes	yes	down	INTA2
INTA3	3	1G/10G	full	yes	yes	down	INTA3
INTA4	4	1G/10G	full	yes	yes	down	INTA4
INTA5	5	1G/10G	full	yes	yes	down	INTA5
INTA6	6	1G/10G	full	yes	yes	down	INTA6
INTA7	7	1G/10G	full	yes	yes	down	INTA7
INTA8	8	1G/10G	full	yes	yes	down	INTA8
INTA9	9	1G/10G	full	yes	yes	down	INTA9
INTA10	10	1G/10G	full	yes	yes	down	INTA10
INTA11	11	1G/10G	full	yes	yes	down	INTA11
INTA12	12	1G/10G	full	yes	yes	down	INTA12
INTA13	13	1G/10G	full	yes	yes	down	INTA13
INTA14	14	1G/10G	full	yes	yes	down	INTA14
INTB1	15	1000	full	no	no	up	INTB1
INTB2	16	1G/10G	full	yes	yes	down	INTB2
INTB3	17	1G/10G	full	yes	yes	down	INTB3
INTB4	18	1G/10G	full	yes	yes	down	INTB4
INTB5	19	1G/10G	full	yes	yes	down	INTB5
INTB6	20	1G/10G	full	yes	yes	down	INTB6
INTB7	21	1G/10G	full	yes	yes	down	INTB7
INTB8	22	1G/10G	full	yes	yes	down	INTB8
INTB9	23	1G/10G	full	yes	yes	down	INTB9
INTB10	24	1G/10G	full	yes	yes	down	INTB10
INTB11	25	1G/10G	full	yes	yes	down	INTB11
INTB12	26	1G/10G	full	yes	yes	down	INTB12
INTB13	27	1G/10G	full	yes	yes	down	INTB13
INTB14	28	1G/10G	full	yes	yes	down	INTB14
EXT1	43	10000	full	no	no	up	EXT1
EXT2	44	10000	full	no	no	up	EXT2
EXT3	45	10000	full	no	no	up	EXT3
EXT4	46	10000	full	no	no	up	ISL h1thchk
EXT5	47	1G/10G	full	no	no	down	EXT5
EXT6	48	1G/10G	full	no	no	down	EXT6
EXT7	49	10000	full	no	no	up	ISL
EXT8	50	10000	full	no	no	up	ISL
EXT9	51	10000	full	no	no	up	ISL
EXT10	52	10000	full	no	no	up	ISL
EXT15	57	10000	full	no	no	up	Link to g8264tor_1
EXT16	58	10000	full	no	no	up	Link to g8264tor_1
EXT17	59	10000	full	no	no	up	Link to g8264tor_1
EXT18	60	10000	full	no	no	up	Link to g8264tor_1
EXT19	61	10000	full	no	no	up	Link to g8264tor_2
EXT20	62	10000	full	no	no	up	Link to g8264tor_2
EXT21	63	10000	full	no	no	up	Link to g8264tor_2
EXT22	64	10000	full	no	no	up	Link to g8264tor_2

EXTM	65	1000	half	yes	yes	down	EXTM
MGT1	66	1000	full	yes	yes	up	MGT1

Show lldp remote-device

The command output in Example 5-95 shows the physical topology, and verifies that cables are plugged into the ports specified in both the Network Topology diagram and the configuration specified in the appendix.

Example 5-95 EN4093flex_1 show lldp remote-device output

LLDP Remote Devices Information				
LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
EXT16	3	08 17 f4 33 9d 00	25	G8264TOR-1
EXT15	4	08 17 f4 33 9d 00	26	G8264TOR-1
EXT18	5	08 17 f4 33 9d 00	27	G8264TOR-1
EXT17	6	08 17 f4 33 9d 00	28	G8264TOR-1
EXT21	7	08 17 f4 33 75 00	25	G8264TOR-2
EXT19	8	08 17 f4 33 75 00	26	G8264TOR-2
EXT22	9	08 17 f4 33 75 00	27	G8264TOR-2
EXT20	10	08 17 f4 33 75 00	28	G8264TOR-2
EXT4	12	6c ae 8b bf fe 00	46	en4093flex_2
EXT7	13	6c ae 8b bf fe 00	49	en4093flex_2
EXT8	14	6c ae 8b bf fe 00	50	en4093flex_2
EXT9	15	6c ae 8b bf fe 00	51	en4093flex_2
EXT10	16	6c ae 8b bf fe 00	52	en4093flex_2

Show vlag isl

Example 5-96 shows command output about the status of the ISL between the EN4093/R switches, and the ports that comprise the ISL itself.

Example 5-96 EN4093flex_1 show vlag isl output

ISL_ID	ISL_Vlan	ISL_Trunk	ISL_Members	Link_State	Trunk_State
65	4094	Adminkey 1000	EXT7	UP	UP
			EXT8	UP	UP
			EXT9	UP	UP
			EXT10	UP	UP

Show vlag information

The command output in Example 5-97 on page 139 shows that the vLAG between the EN4093/R switches and G8264 switches is up and operational as referenced by the LACP admin key of 2000. The ISL between the EN4093/R switches is up as well.

EN4093flex_1 is acting as the admin and operational role of PRIMARY. For centralized vLAG functions, such as vLAG STP, one of the vLAG switches must control the protocol operations. To select the switch that controls the centralized vLAG function, perform role election. The switch with the primary role controls the centralized operation. Role election is non-preemptive. That is, if a primary already exists, another switch that is coming up remains as secondary even if it can become primary based on the role election logic.

Role election is determined by comparing the local vLAG system priority and local system MAC address. The switch with the smaller priority value becomes the vLAG primary switch. If

the priorities are the same, the switch with smaller system MAC address becomes the vLAG primary switch. You can configure vLAG priority to anything between <0-65535>. In these examples, priority was left at the default value of 0.

Example 5-97 EN4093flex_1 show vlag information output

```
vLAG Tier ID: 1
vLAG system MAC: 08:17:f4:c3:dd:00
Local MAC 6c:ae:8b:bf:6d:00 Priority 0 Admin Role PRIMARY (Operational Role
PRIMARY)
Peer MAC 6c:ae:8b:bf:fe:00 Priority 0
Health local 1.1.1.1 peer 1.1.1.2 State UP
ISL trunk id 65
ISL state Up
Startup Delay Interval: 120s (Finished)
```

vLAG 65: config with admin key 2000, associated trunk 66, state formed

Show vlag adminkey 2000

Example 5-98 output shows that the vLAG is formed and enabled by using LACP reference key 2000.

Example 5-98 EN4093flex_1 show vlag adminkey 2000 output

```
vLAG is enabled on admin key 2000
Current LACP params for EXT15: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT16: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT17: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT18: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT19: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT20: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT21: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT22: active, Priority 32768, Admin Key 2000, Min-Links 1
```

Show lacp information state up

The command output in Example 5-99 shows which ports are participating in an LACP aggregation, and which reference keys are used on those specific interfaces.

Example 5-99 EN4093flex_1 show lacp information state up

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
EXT7	active	1000	1000	yes	32768	49	65	up	1
EXT8	active	1000	1000	yes	32768	49	65	up	1
EXT9	active	1000	1000	yes	32768	49	65	up	1
EXT10	active	1000	1000	yes	32768	49	65	up	1
EXT15	active	2000	2000	yes	32768	57	66	up	1
EXT16	active	2000	2000	yes	32768	57	66	up	1
EXT17	active	2000	2000	yes	32768	57	66	up	1

EXT18	active	2000	2000	yes	32768	57	66	up	1
EXT19	active	2000	2000	yes	32768	57	66	up	1
EXT20	active	2000	2000	yes	32768	57	66	up	1
EXT21	active	2000	2000	yes	32768	57	66	up	1
EXT22	active	2000	2000	yes	32768	57	66	up	1

Show failover trigger 1

The failover output in Example 5-100 shows which ports are monitored, and which ports are shut down if an issue is encountered. In this example, the upstream to G8264 links are monitored with LACP reference key 2000. The control ports are the downstream internal I/O module ports that are used by the compute nodes.

Example 5-100 EN4093flex_1 show failover output

```

Failover: On
VLAN Monitor: OFF

Trigger 1 Manual Monitor: Enabled
Trigger 1 limit: 0
Monitor State: Up
Member      Status
-----
adminkey 2000
EXT15      Operational
EXT16      Operational
EXT17      Operational
EXT18      Operational
EXT19      Operational
EXT20      Operational
EXT21      Operational
EXT22      Operational
Control State: Auto Controlled
Member      Status
-----
INTA1      Operational
INTA2      Operational
INTA3      Operational
INTA4      Operational
INTA5      Operational
INTA6      Operational
INTA7      Operational
INTA8      Operational
INTA9      Operational
INTA10     Operational
INTA11     Operational
INTA12     Operational
INTA13     Operational
INTA14     Operational
INTB1      Operational
INTB2      Operational
INTB3      Operational
INTB4      Operational
INTB5      Operational
INTB6      Operational
INTB7      Operational

```

INTB8	Operational
INTB9	Operational
INTB10	Operational
INTB11	Operational
INTB12	Operational
INTB13	Operational
INTB14	Operational

Trigger 2: Disabled

Trigger 3: Disabled

Trigger 4: Disabled

Trigger 5: Disabled

Trigger 6: Disabled

Trigger 7: Disabled

Trigger 8: Disabled

Ping output for equipment on VLAN 4092

To verify connectivity, issue **ping** commands to devices on VLAN 4092 (Data VLAN) as shown in Example 5-101. IP address 10.4.1.10 represents a Compute Node with an operating system installed, flex_node1 on the Network Topology diagram.

Example 5-101 Ping verification for equipment on VLAN 4092

```
en4093flex_1#ping 10.1.4.10 data-port
Connecting via DATA port.
[host 10.1.4.10, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl 255,
tos 0]
10.1.4.10: #1 ok, RTT 1 msec.
10.1.4.10: #2 ok, RTT 0 msec.
10.1.4.10: #3 ok, RTT 1 msec.
10.1.4.10: #4 ok, RTT 0 msec.
10.1.4.10: #5 ok, RTT 0 msec.
Ping finished.
```

```
en4093flex_1#ping 10.1.4.239 data-port
Connecting via DATA port.
[host 10.1.4.239, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.239: #1 ok, RTT 4 msec.
10.1.4.239: #2 ok, RTT 1 msec.
10.1.4.239: #3 ok, RTT 2 msec.
10.1.4.239: #4 ok, RTT 3 msec.
10.1.4.239: #5 ok, RTT 1 msec.
Ping finished.
```

```
en4093flex_1#ping 10.1.4.243 data-port
Connecting via DATA port.
[host 10.1.4.243, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
```

```
10.1.4.243: #1 ok, RTT 1 msec.  
10.1.4.243: #2 ok, RTT 1 msec.  
10.1.4.243: #3 ok, RTT 2 msec.  
10.1.4.243: #4 ok, RTT 8 msec.  
10.1.4.243: #5 ok, RTT 6 msec.  
Ping finished.
```

```
en4093flex_1#ping 10.1.4.244 data-port  
Connecting via DATA port.  
[host 10.1.4.244, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl  
255, tos 0]  
10.1.4.244: #1 ok, RTT 1 msec.  
10.1.4.244: #2 ok, RTT 2 msec.  
10.1.4.244: #3 ok, RTT 1 msec.  
10.1.4.244: #4 ok, RTT 2 msec.  
10.1.4.244: #5 ok, RTT 0 msec.  
Ping finished.
```

```
en4093flex_1#ping 10.1.4.249 data-port  
Connecting via DATA port.  
[host 10.1.4.241, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl  
255, tos 0]  
10.1.4.241: #1 ok, RTT 2 msec.  
10.1.4.241: #2 ok, RTT 1 msec.  
10.1.4.241: #3 ok, RTT 2 msec.  
10.1.4.241: #4 ok, RTT 1 msec.  
10.1.4.241: #5 ok, RTT 3 msec.  
Ping finished.
```

```
en4093flex_1#ping 10.1.4.200 data-port  
Connecting via DATA port.  
[host 10.1.4.241, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl  
255, tos 0]  
10.1.4.241: #1 ok, RTT 2 msec.  
10.1.4.241: #2 ok, RTT 2 msec.  
10.1.4.241: #3 ok, RTT 2 msec.  
10.1.4.241: #4 ok, RTT 1 msec.  
10.1.4.241: #5 ok, RTT 3 msec.  
Ping finished
```

G8264 output

This section lists output from the switch with hostname G8264tor_1. Similar or identical output exists for the switch with hostname G8264tor_2 unless otherwise noted.

Show version

Example 5-102 shows information about the switch used, and the associated code/firmware level.

Example 5-102 G8264tor_1 show version output

```
System Information at 20:30:07 Thu Oct 18, 2012  
Time zone: No timezone configured  
Daylight Savings Time Status: Disabled
```

```
IBM Networking Operating System RackSwitch G8264
```

Switch has been up for 1 day, 20 hours, 28 minutes and 18 seconds.
Last boot: 6:05:44 Thu Feb 7, 2001 (reset from console)

MAC address: 08:17:f4:33:9d:00 IP (If 20) address: 10.10.20.2
Management Port MAC Address: 08:17:f4:33:9d:fe
Management Port IP Address (if 128): 172.25.101.243
Hardware Revision: 0
Hardware Part No: BAC-00065-00
Switch Serial No: US71120007
Manufacturing date: 11/13

Software Version 7.4.1.0 (FLASH image1), active configuration.

Temperature Mother Top: 26 C
Temperature Mother Bottom: 32 C
Temperature Daughter Top: 26 C
Temperature Daughter Bottom: 30 C

Warning at 75 C and Recover at 90 C

Fan 1 in Module 1: RPM= 8463 PWM= 15(5%) Front-To-Back
Fan 2 in Module 1: RPM= 3976 PWM= 15(5%) Front-To-Back
Fan 3 in Module 2: RPM= 8667 PWM= 15(5%) Front-To-Back
Fan 4 in Module 2: RPM= 4115 PWM= 15(5%) Front-To-Back
Fan 5 in Module 3: RPM= 7894 PWM= 15(5%) Front-To-Back
Fan 6 in Module 3: RPM= 4195 PWM= 15(5%) Front-To-Back
Fan 7 in Module 4: RPM= 8852 PWM= 15(5%) Front-To-Back
Fan 8 in Module 4: RPM= 3976 PWM= 15(5%) Front-To-Back

System Fan Airflow: Front-To-Back

Power Supply 1: OK
Power Supply 2: OK

Power Faults: ()
Fan Faults: ()
Service Faults: ()

Show vlan

Example 5-103 shows VLAN assignment for all of the ports on the switch.

Example 5-103 G8264tor_1 show vlan output

VLAN	Name	Status	Ports
1	Default VLAN	ena	17-63
4000	ISL hlthchk	ena	64
4092	DATA	ena	1-16 18 20 22 24-28 37-40
4094	ISL	ena	1-16
4095	Mgmt VLAN	ena	MGT

Show interface status

Because there is only one compute node in the chassis (in slot 1), all the other internal ports are listed as “down” from a link perspective in the output in Example 5-104.

Example 5-104 G8264tor_1 show interface status output

Alias	Port	Speed	Duplex	Flow Ctrl		Link	Name
				--TX--	--RX--		
1	1	10000	full	no	no	up	ISL
2	2	10000	full	no	no	up	ISL
3	3	10000	full	no	no	up	ISL
4	4	10000	full	no	no	up	ISL
5	5	10000	full	no	no	up	ISL
6	6	10000	full	no	no	up	ISL
7	7	10000	full	no	no	up	ISL
8	8	10000	full	no	no	up	ISL
9	9	10000	full	no	no	up	ISL
10	10	10000	full	no	no	up	ISL
11	11	10000	full	no	no	up	ISL
12	12	10000	full	no	no	up	ISL
13	13	10000	full	no	no	up	ISL
14	14	10000	full	no	no	up	ISL
15	15	10000	full	no	no	up	ISL
16	16	10000	full	no	no	up	ISL
17	17	1G/10G	full	no	no	down	17
18	18	10000	full	no	no	up	Po5 to
Nexus5548Core_1							
19	19	1G/10G	full	no	no	down	19
20	20	10000	full	no	no	up	Po5 to
Nexus5548Core_1							
21	21	1G/10G	full	no	no	down	21
22	22	10000	full	no	no	up	Po6 to
Nexus5548Core_2							
23	23	1G/10G	full	no	no	down	23
24	24	10000	full	no	no	up	Po6 to
Nexus5548Core_2							
25	25	10000	full	no	no	up	Link to EN4093-1
26	26	10000	full	no	no	up	Link to EN4093-1
27	27	10000	full	no	no	up	Link to EN4093-1
28	28	10000	full	no	no	up	Link to EN4093-1
29	29	1G/10G	full	no	no	down	29
30	30	1G/10G	full	no	no	down	30
31	31	1G/10G	full	no	no	down	31
32	32	1G/10G	full	no	no	down	32
33	33	1G/10G	full	no	no	down	33
34	34	1G/10G	full	no	no	down	34
35	35	1G/10G	full	no	no	down	35
36	36	1G/10G	full	no	no	down	36
37	37	10000	full	no	no	up	Link to EN4093-2
38	38	10000	full	no	no	up	Link to EN4093-2
39	39	10000	full	no	no	up	Link to EN4093-2
40	40	10000	full	no	no	up	Link to EN4093-2
41	41	1G/10G	full	no	no	down	41
42	42	1G/10G	full	no	no	down	42
43	43	1G/10G	full	no	no	down	43

44	44	1G/10G	full	no	no	down	44
45	45	1G/10G	full	no	no	down	45
46	46	1G/10G	full	no	no	down	46
47	47	1G/10G	full	no	no	down	47
48	48	1G/10G	full	no	no	down	48
49	49	1G/10G	full	no	no	down	49
50	50	1G/10G	full	no	no	down	50
51	51	1G/10G	full	no	no	down	51
52	52	1G/10G	full	no	no	down	52
53	53	1G/10G	full	no	no	down	53
54	54	1G/10G	full	no	no	down	54
55	55	1G/10G	full	no	no	down	55
56	56	1G/10G	full	no	no	down	56
57	57	1G/10G	full	no	no	down	57
58	58	1G/10G	full	no	no	down	58
59	59	1G/10G	full	no	no	down	59
60	60	1G/10G	full	no	no	down	60
61	61	1G/10G	full	no	no	down	61
62	62	1G/10G	full	no	no	down	62
63	63	1G/10G	full	no	no	down	63
64	64	10000	full	no	no	up	ISL h1thchk
MGT	65	1000	full	yes	yes	up	MGT

Show lldp remote-device

The command output in Example 5-105 shows the physical topology and verifies that cables are plugged into the ports specified in both the Network Topology diagram, and the configuration specified in the appendix.

Example 5-105 G8264tor_1 show lldp remote-device output

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
1	1	08 17 f4 33 75 00	1	G8264TOR-2
2	2	08 17 f4 33 75 00	2	G8264TOR-2
3	3	08 17 f4 33 75 00	3	G8264TOR-2
4	4	08 17 f4 33 75 00	4	G8264TOR-2
5	6	08 17 f4 33 75 00	5	G8264TOR-2
6	7	08 17 f4 33 75 00	6	G8264TOR-2
7	8	08 17 f4 33 75 00	7	G8264TOR-2
8	9	08 17 f4 33 75 00	8	G8264TOR-2
9	10	08 17 f4 33 75 00	9	G8264TOR-2
10	11	08 17 f4 33 75 00	10	G8264TOR-2
11	12	08 17 f4 33 75 00	11	G8264TOR-2
12	13	08 17 f4 33 75 00	12	G8264TOR-2
18	14	54 7f ee 2d 36 0e	Eth1/7	Nexus5548core_1
13	15	08 17 f4 33 75 00	13	G8264TOR-2
20	16	54 7f ee 2d 36 0f	Eth1/8	Nexus5548core_1
22	17	00 05 73 bc 02 70	Eth1/9	Nexus5548core_2
14	18	08 17 f4 33 75 00	14	G8264TOR-2
24	19	00 05 73 bc 02 71	Eth1/10	Nexus5548core_2
25	20	6c ae 8b bf 6d 00	58	en4093flex_1
15	21	08 17 f4 33 75 00	15	G8264TOR-2
26	22	6c ae 8b bf 6d 00	57	en4093flex_1
27	23	6c ae 8b bf 6d 00	60	en4093flex_1
16	24	08 17 f4 33 75 00	16	G8264TOR-2

28	25	6c ae 8b bf 6d 00	59	en4093flex_1
37	26	6c ae 8b bf fe 00	57	en4093flex_2
38	27	6c ae 8b bf fe 00	59	en4093flex_2
39	28	6c ae 8b bf fe 00	58	en4093flex_2
40	29	6c ae 8b bf fe 00	60	en4093flex_2
64	30	08 17 f4 33 75 00	64	G8264TOR-2

Show vlag isl

The command output in Example 5-106 shows the status of the ISL between the G8264 switches, and the ports that comprise the ISL itself.

Example 5-106 G8264tor_1 show vlag isl output

ISL_ID	ISL_Vlan	ISL_Trunk	ISL_Members	Link_State	Trunk_State
67	4094	Adminkey 1000	1	UP	UP
			2	UP	UP
			3	UP	UP
			4	UP	UP
			5	UP	UP
			6	UP	UP
			7	UP	UP
			8	UP	UP
			9	UP	UP
			10	UP	UP
			11	UP	UP
			12	UP	UP
			13	UP	UP
			14	UP	UP
			15	UP	UP
			16	UP	UP

Show vlag information

Example 5-107 output shows that the downstream vLAG between the G8264 and EN4093/R switches is up and operational as referenced by the LACP admin key of 2002. The ISL between the G8264 switches is up too.

G8264tor_1 is acting as the admin and operational role of SECONDARY. For centralized vLAG functions, such as vLAG STP, one of the vLAG switches must control the protocol operations. To select the switch that controls the centralized vLAG function, perform role election. The switch with the primary role controls the centralized operation. Role election is non-preemptive. That is, a primary already exists, another switch that is coming up remains as secondary even if it can become primary based on the role election logic.

Role election is determined by comparing the local vLAG system priority and local system MAC address. The switch with the smaller priority value becomes the vLAG primary switch. If priorities are the same, the switch with the smaller system MAC address becomes the vLAG primary switch. You can configure vLAG priority to anything between <0-65535>. Priority was left at the default value of 0 in all examples.

Example 5-107 G8264tor_1 show vlag information output

```
vLAG Tier ID: 2
vLAG system MAC: 08:17:f4:c3:dd:01
Local MAC 08:17:f4:33:9d:00 Priority 0 Admin Role SECONDARY (Operational Role
SECONDARY)
```



```

Peer MAC 08:17:f4:33:75:00 Priority 0
Health local 1.1.1.1 peer 1.1.1.2 State UP
ISL trunk id 67
ISL state Up
Startup Delay Interval: 120s (Finished)

```

vLAG 66: config with admin key 2002, associated trunk 66, state formed

Show vlag adminkey 2002

The output in Example 5-108 shows that the downstream vLAG towards the EN4093/R switches is formed and enabled by using LACP reference key 2002.

Example 5-108 G8264tor_1 show vlag adminkey 2002 output

```

vLAG is enabled on admin key 2002
Current LACP params for 25: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 26: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 27: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 28: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 37: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 38: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 39: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 40: active, Priority 32768, Admin Key 2002, Min-Links 1

```

Show lacp information state up

Example 5-109 shows which ports are participating in an LACP aggregation, and which reference keys are used on those specific interfaces.

Example 5-109 G8264tor_1 show lacp information state up

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
1	active	1000	1000	yes	32768	1	67	up	1
2	active	1000	1000	yes	32768	1	67	up	1
3	active	1000	1000	yes	32768	1	67	up	1
4	active	1000	1000	yes	32768	1	67	up	1
5	active	1000	1000	yes	32768	1	67	up	1
6	active	1000	1000	yes	32768	1	67	up	1
7	active	1000	1000	yes	32768	1	67	up	1
8	active	1000	1000	yes	32768	1	67	up	1
9	active	1000	1000	yes	32768	1	67	up	1
10	active	1000	1000	yes	32768	1	67	up	1
11	active	1000	1000	yes	32768	1	67	up	1
12	active	1000	1000	yes	32768	1	67	up	1
13	active	1000	1000	yes	32768	1	67	up	1
14	active	1000	1000	yes	32768	1	67	up	1
15	active	1000	1000	yes	32768	1	67	up	1
16	active	1000	1000	yes	32768	1	67	up	1

25	active	2002	2002	yes	32768	26	66	up	1
26	active	2002	2002	yes	32768	26	66	up	1
27	active	2002	2002	yes	32768	26	66	up	1
28	active	2002	2002	yes	32768	26	66	up	1
37	active	2002	2002	yes	32768	26	66	up	1
38	active	2002	2002	yes	32768	26	66	up	1
39	active	2002	2002	yes	32768	26	66	up	1
40	active	2002	2002	yes	32768	26	66	up	1

Show spanning-tree on G8264tor_1

Example 5-110 lists output from the **show spanning-tree** command on G8264tor_1. Note that the blocked links are reflected in the Network Topology diagram for VLAN 4092.

Example 5-110 G8264tor_1 show spanning-tree output

```
-----
Pvst+ compatibility mode enabled
-----
```

```
-----
Spanning Tree Group 1: On (PVRST)
VLANs: 1
-----
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel
8001 08:17:f4:33:9d:00      0      0   2    20    15
```

```
Parameters:  Priority  Hello  MaxAge  FwdDel  Aging  Topology Change Counts
              32769      2      20      15     300           14
```

	Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
18	(pc5)	128	990!+	FWD	DESG	8001-08:17:f4:33:9d:00	8404	P2P
20	(pc5)	128	990!+	FWD	DESG	8001-08:17:f4:33:9d:00	8404	P2P
22	(pc6)	128	990!+	FWD	DESG	8001-08:17:f4:33:9d:00	8405	P2P
24	(pc6)	128	990!+	FWD	DESG	8001-08:17:f4:33:9d:00	8405	P2P
25	(pc65)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c0	P2P
26	(pc65)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c0	P2P
27	(pc65)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c0	P2P
28	(pc65)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c0	P2P
37	(pc65)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c0	P2P
38	(pc65)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c0	P2P
39	(pc65)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c0	P2P
40	(pc65)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c0	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

```
-----
Spanning Tree Group 125: On (PVRST)
VLANs: 4000
-----
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel
807d 08:17:f4:33:75:00    2000     64   2    20    15
```

```
Parameters:  Priority  Hello  MaxAge  FwdDel  Aging  Topology Change Counts
              32893      2      20      15     300           1
```

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
64	128	2000!	FWD	ROOT	807d-08:17:f4:33:75:00	8040	P2P

! = Automatic path cost.

Spanning Tree Group 126: On (PVRST)
VLANs: 4092

Current Root:	Path-Cost	Port	Hello	MaxAge	FwdDel
2ffc 54:7f:ee:2d:36:41	1105	1	2	20	15

Parameters:	Priority	Hello	MaxAge	FwdDel	Aging	Topology Change Counts
	32894	2	20	15	300	17

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
1	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
2	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
3	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
4	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
5	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
6	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
7	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
8	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
9	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
10	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
11	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
12	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
13	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
14	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
15	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
16	(pc66)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
18	(pc5)	128	990!+ DISC	ALTN	2ffc-54:7f:ee:2d:36:41	9004	P2P
20	(pc5)	128	990!+ DISC	ALTN	2ffc-54:7f:ee:2d:36:41	9004	P2P
22	(pc6)	128	990!+ DISC	ALTN	4ffc-00:05:73:bc:02:bc	9005	P2P
24	(pc6)	128	990!+ DISC	ALTN	4ffc-00:05:73:bc:02:bc	9005	P2P
25	(pc65)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c0	P2P
26	(pc65)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c0	P2P
27	(pc65)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c0	P2P
28	(pc65)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c0	P2P
37	(pc65)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c0	P2P
38	(pc65)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c0	P2P
39	(pc65)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c0	P2P
40	(pc65)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c0	P2P

! = Automatic path cost.
+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 127: Off (PVRST), FDB aging timer 300
VLANs: 4094

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
1	(pc66)	0	0	FWD *			

```

2      (pc66)      0      0      FWD *
3      (pc66)      0      0      FWD *
4      (pc66)      0      0      FWD *
5      (pc66)      0      0      FWD *
6      (pc66)      0      0      FWD *
7      (pc66)      0      0      FWD *
8      (pc66)      0      0      FWD *
9      (pc66)      0      0      FWD *
10     (pc66)      0      0      FWD *
11     (pc66)      0      0      FWD *
12     (pc66)      0      0      FWD *
13     (pc66)      0      0      FWD *
14     (pc66)      0      0      FWD *
15     (pc66)      0      0      FWD *
16     (pc66)      0      0      FWD *

```

* = STP turned off for this port.

```

-----
Spanning Tree Group 128: Off (PVRST), FDB aging timer 300
VLANs:  4095

```

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
MGT	0	0	FWD *				

* = STP turned off for this port.

Show spanning-tree on G8264tor_2

Example 5-111 lists output from the **show spanning-tree** command on G8264tor_2. Note that the blocked links are reflected in the Network Topology diagram for VLAN 4092.

Example 5-111 G8264tor_2 show spanning-tree output

```

-----
Pvst+ compatibility mode enabled

```

```

-----
Spanning Tree Group 1: On (PVRST)
VLANs:  1

```

```

Current Root:          Path-Cost  Port Hello MaxAge FwdDel
8001 08:17:f4:33:75:00      0      0   2    20    15

```

```

Parameters:  Priority  Hello  MaxAge  FwdDel  Aging  Topology Change Counts
              32769      2      20      15      300      31

```

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
1 (pc65)	128	115!+	FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
2 (pc65)	128	115!+	FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
3 (pc65)	128	115!+	FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
4 (pc65)	128	115!+	FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
5 (pc65)	128	115!+	FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
6 (pc65)	128	115!+	FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
7 (pc65)	128	115!+	FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
8 (pc65)	128	115!+	FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P

9	(pc65)	128	115!+ FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
10	(pc65)	128	115!+ FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
11	(pc65)	128	115!+ FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
12	(pc65)	128	115!+ FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
13	(pc65)	128	115!+ FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
14	(pc65)	128	115!+ FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
15	(pc65)	128	115!+ FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
16	(pc65)	128	115!+ FWD	DESG	8001-08:17:f4:33:75:00	8440	P2P
18	(pc5)	128	990!+ FWD	DESG	8001-08:17:f4:33:75:00	8404	P2P
20	(pc5)	128	990!+ FWD	DESG	8001-08:17:f4:33:75:00	8404	P2P
22	(pc6)	128	990!+ FWD	DESG	8001-08:17:f4:33:75:00	8405	P2P
24	(pc6)	128	990!+ FWD	DESG	8001-08:17:f4:33:75:00	8405	P2P
25	(pc66)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c0	P2P
26	(pc66)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c0	P2P
27	(pc66)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c0	P2P
28	(pc66)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c0	P2P
37	(pc66)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c0	P2P
38	(pc66)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c0	P2P
39	(pc66)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c0	P2P
40	(pc66)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c0	P2P

! = Automatic path cost.
+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 125: On (PVRST)
VLANs: 4000

Current Root:	Path-Cost	Port	Hello	MaxAge	FwdDel
807d 08:17:f4:33:75:00	0	0	2	20	15

Parameters:	Priority	Hello	MaxAge	FwdDel	Aging	Topology Change Counts
	32893	2	20	15	300	1

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
-----	-----	-----	-----	-----	-----	-----	-----
64	128	2000!	FWD	DESG	807d-08:17:f4:33:75:00	8040	P2P

! = Automatic path cost.

Spanning Tree Group 126: On (PVRST)
VLANs: 4092

Current Root:	Path-Cost	Port	Hello	MaxAge	FwdDel
2ffc 54:7f:ee:2d:36:41	990	22	2	20	15

Parameters:	Priority	Hello	MaxAge	FwdDel	Aging	Topology Change Counts
	32894	2	20	15	300	10

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
-----	-----	-----	-----	-----	-----	-----	-----
1	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
2	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
3	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
4	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
5	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P

6	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
7	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
8	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
9	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
10	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
11	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
12	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
13	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
14	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
15	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
16	(pc65)	128	115!+ FWD	DESG	807e-08:17:f4:33:75:00	8440	P2P
18	(pc5)	128	990!+ DISC	ALTN	4ffc-00:05:73:bc:02:bc	9004	P2P
20	(pc5)	128	990!+ DISC	ALTN	4ffc-00:05:73:bc:02:bc	9004	P2P
22	(pc6)	128	990!+ FWD	ROOT	2ffc-54:7f:ee:2d:36:41	9005	P2P
24	(pc6)	128	990!+ FWD	ROOT	2ffc-54:7f:ee:2d:36:41	9005	P2P
25	(pc66)	128	200!+ FWD	DESG	807e-08:17:f4:33:75:00	84c0	P2P
26	(pc66)	128	200!+ FWD	DESG	807e-08:17:f4:33:75:00	84c0	P2P
27	(pc66)	128	200!+ FWD	DESG	807e-08:17:f4:33:75:00	84c0	P2P
28	(pc66)	128	200!+ FWD	DESG	807e-08:17:f4:33:75:00	84c0	P2P
37	(pc66)	128	200!+ FWD	DESG	807e-08:17:f4:33:75:00	84c0	P2P
38	(pc66)	128	200!+ FWD	DESG	807e-08:17:f4:33:75:00	84c0	P2P
39	(pc66)	128	200!+ FWD	DESG	807e-08:17:f4:33:75:00	84c0	P2P
40	(pc66)	128	200!+ FWD	DESG	807e-08:17:f4:33:75:00	84c0	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 127: Off (PVRST), FDB aging timer 300
VLANs: 4094

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
1	(pc65)	0	0	FWD *			
2	(pc65)	0	0	FWD *			
3	(pc65)	0	0	FWD *			
4	(pc65)	0	0	FWD *			
5	(pc65)	0	0	FWD *			
6	(pc65)	0	0	FWD *			
7	(pc65)	0	0	FWD *			
8	(pc65)	0	0	FWD *			
9	(pc65)	0	0	FWD *			
10	(pc65)	0	0	FWD *			
11	(pc65)	0	0	FWD *			
12	(pc65)	0	0	FWD *			
13	(pc65)	0	0	FWD *			
14	(pc65)	0	0	FWD *			
15	(pc65)	0	0	FWD *			
16	(pc65)	0	0	FWD *			

* = STP turned off for this port.

Spanning Tree Group 128: Off (PVRST), FDB aging timer 300
VLANs: 4095

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
------	------	------	-------	------	-------------------	----------	------

```
MGT          0          0    FWD *
* = STP turned off for this port.
```

Ping output for equipment on VLAN 4092

To verify connectivity, issue **ping** commands to devices on VLAN 4092 (Data VLAN) as shown in Example 5-112. IP address 10.4.1.10 represents a compute node with an operating system installed, flex_node1 on the Network Topology diagram.

Example 5-112 Ping verification for equipment on VLAN 4092

```
G8264TOR-1#ping 10.1.4.249 data-port
Connecting via DATA port.
[host 10.1.4.249, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.249: #1 ok, RTT 0 msec.
10.1.4.249: #2 ok, RTT 0 msec.
10.1.4.249: #3 ok, RTT 0 msec.
10.1.4.249: #4 ok, RTT 0 msec.
10.1.4.249: #5 ok, RTT 0 msec.
Ping finished.

G8264TOR-1#ping 10.1.4.238 data-port
Connecting via DATA port.
[host 10.1.4.238, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.238: #1 ok, RTT 4 msec.
10.1.4.238: #2 ok, RTT 2 msec.
10.1.4.238: #3 ok, RTT 0 msec.
10.1.4.238: #4 ok, RTT 1 msec.
10.1.4.238: #5 ok, RTT 1 msec.
Ping finished.

G8264TOR-1#ping 10.1.4.10 data-port
Connecting via DATA port.
[host 10.1.4.10, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl 255,
tos 0]
10.1.4.10: #1 ok, RTT 0 msec.
10.1.4.10: #2 ok, RTT 0 msec.
10.1.4.10: #3 ok, RTT 0 msec.
10.1.4.10: #4 ok, RTT 0 msec.
10.1.4.10: #5 ok, RTT 0 msec.
Ping finished.
```

Nexus output

This section lists output from the switch with hostname Nexus5548core_1. Similar or identical output exists for the switch with hostname Nexus5548core_2 unless otherwise noted.

Show version

Example 5-113 shows information about the switch, and the associated code/firmware level.

Example 5-113 Nexus5548core_1 show version output

```
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Documents: http://www.cisco.com/en/US/products/ps9372/tsd_products_support_series_home.html
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software are covered under the GNU Public
License. A copy of the license is available at
http://www.gnu.org/licenses/gpl.html.
```

Software

```
BIOS:      version 3.5.0
loader:    version N/A
kickstart: version 5.2(1)N1(1b)
system:    version 5.2(1)N1(1b)
power-seq: Module 1: version v1.0
           Module 3: version v5.0
uC:        version v1.2.0.1
SFP uC:    Module 1: v1.0.0.0
BIOS compile time:      02/03/2011
kickstart image file is: bootflash:///n5000-uk9-kickstart.5.2.1.N1.1b.bin
kickstart compile time: 9/17/2012 11:00:00 [09/17/2012 18:38:53]
system image file is:   bootflash:///n5000-uk9.5.2.1.N1.1b.bin
system compile time:    9/17/2012 11:00:00 [09/17/2012 20:38:22]
```

Hardware

```
cisco Nexus5548 Chassis ("02 32X10GE/Modular Universal Platform Supervisor")
Intel(R) Xeon(R) CPU          with 8263848 kB of memory.
Processor Board ID FOC15424504
```

```
Device name: Nexus5548core_1
bootflash:   2007040 kB
```

```
Kernel uptime is 0 day(s), 22 hour(s), 32 minute(s), 3 second(s)
```

Last reset

```
Reason: Unknown
System version: 5.2(1)N1(1b)
Service:
```

plugin

```
Core Plugin, Ethernet Plugin
```

Show vlan

Example 5-114 displays the VLAN assignments for all of the ports on the switch.

Example 5-114 Nexus5548core_1 show vlan output

VLAN	Name	Status	Ports
1	default	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/11, Eth1/12 Eth1/13, Eth1/14, Eth1/15 Eth1/16, Eth1/18, Eth1/20 Eth1/21, Eth1/22, Eth1/23 Eth1/24, Eth1/25, Eth1/26 Eth1/27, Eth1/28, Eth1/29 Eth1/30, Eth1/31, Eth1/32
4092	DATA_VLAN	active	Po5, Po6, Po100, Eth1/7, Eth1/8 Eth1/9, Eth1/10, Eth1/17 Eth1/19

Show interface status

Example 5-115 shows the full interface table, listing port status, speed, and so on, for the Nexus5548core_1 switch.

Example 5-115 Nexus5548core_1 show interface status output

Port	Name	Status	Vlan	Duplex	Speed	Type
Eth1/1	--	sfpAbsent	1	full	10G	--
Eth1/2	--	sfpAbsent	1	full	10G	--
Eth1/3	--	sfpAbsent	1	full	10G	--
Eth1/4	--	sfpAbsent	1	full	10G	--
Eth1/5	--	sfpAbsent	1	full	10G	--
Eth1/6	--	sfpAbsent	1	full	10G	--
Eth1/7	Po5 to G8264tor_1	connected	trunk	full	10G	10Gbase-(un
Eth1/8	Po5 to G8264tor_1	connected	trunk	full	10G	10Gbase-(un
Eth1/9	Po6 to G8264tor_2	connected	trunk	full	10G	10Gbase-(un
Eth1/10	Po6 to G8264tor_2	connected	trunk	full	10G	10Gbase-(un
Eth1/11	--	sfpAbsent	1	full	10G	--
Eth1/12	--	sfpAbsent	1	full	10G	--
Eth1/13	--	sfpAbsent	1	full	10G	--
Eth1/14	--	sfpAbsent	1	full	10G	--
Eth1/15	--	sfpAbsent	1	full	10G	--
Eth1/16	--	sfpAbsent	1	full	10G	--
Eth1/17	Po100 to Nexus5548	connected	trunk	full	10G	10Gbase-(un
Eth1/18	--	sfpAbsent	1	full	10G	--
Eth1/19	Po100 to Nexus5548	connected	trunk	full	10G	10Gbase-(un
Eth1/20	--	sfpAbsent	1	full	10G	--
Eth1/21	--	disabled	1	full	10G	10Gbase-(un
Eth1/22	--	sfpAbsent	1	full	10G	--
Eth1/23	--	sfpAbsent	1	full	10G	--
Eth1/24	--	sfpAbsent	1	full	10G	--
Eth1/25	--	sfpAbsent	1	full	10G	--
Eth1/26	--	sfpAbsent	1	full	10G	--
Eth1/27	--	sfpAbsent	1	full	10G	--
Eth1/28	--	sfpAbsent	1	full	10G	--

Eth1/29	--	sfpAbsent 1	full	10G	--
Eth1/30	--	sfpAbsent 1	full	10G	--
Eth1/31	--	sfpAbsent 1	full	10G	--
Eth1/32	--	sfpAbsent 1	full	10G	--
Po5	--	connected trunk	full	10G	--
Po6	--	connected trunk	full	10G	--
Po100	Switch-to-Switch 1	connected trunk	full	10G	--
mgmt0	--	connected routed	full	1000	--

Show lldp neighbors

Example 5-116 lists the LLDP information and verifies the physical connectivity.

Example 5-116 Nexus5548core_1 show lldp neighbors output

Capability codes:

(R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device

(W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

Device ID	Local Intf	Hold-time	Capability	Port ID
G8264TOR-1	Eth1/7	120	BR	18
G8264TOR-1	Eth1/8	120	BR	20
G8264TOR-2	Eth1/9	120	BR	22
G8264TOR-2	Eth1/10	120	BR	24
Nexus5548core_2	Eth1/17	120	B	Eth1/17
Nexus5548core_2	Eth1/19	120	B	Eth1/19
Total entries displayed: 6				

Show spanning-tree on Nexus5548core_1

Example 5-117 shows the spanning-tree output on the Nexus5548core_1 switch. As indicated in the output, Nexus5548core_1 is in a designated forwarding state from a spanning-tree perspective on all three physical interfaces.

Example 5-117 Nexus5548core_1 show spanning-tree output

VLAN4092

Spanning tree enabled protocol rstp

Root ID Priority 12284

Address 547f.ee2d.3641

This bridge is the root

Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 12284 (priority 8192 sys-id-ext 4092)

Address 547f.ee2d.3641

Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po5	Desg	FWD	1	128.4100	P2p
Po6	Desg	FWD	1	128.4101	P2p
Po100	Desg	FWD	1	128.4195	Network P2p

Show spanning-tree on Nexus5548core_2

Example 5-118 shows the spanning-tree output on the Nexus5548core_2 switch. As indicated in the output, Nexus5548core_2 is in a designated forwarding state from a spanning-tree perspective on all three physical interfaces.

Example 5-118 Nexus5548core_2 show spanning-tree output

VLAN4092					
Spanning tree enabled protocol rstp					
Root ID	Priority	12284			
	Address	547f.ee2d.3641			
	Cost	1			
	Port	4195 (port-channel100)			
	Hello Time	2 sec	Max Age	20 sec	Forward Delay 15 sec
Bridge ID	Priority	20476 (priority 16384 sys-id-ext 4092)			
	Address	0005.73bc.02bc			
	Hello Time	2 sec	Max Age	20 sec	Forward Delay 15 sec
Interface	Role	Sts	Cost	Prio.Nbr	Type

Po5	Desg	FWD	1	128.4100	P2p
Po6	Desg	FWD	1	128.4101	P2p
Po100	Root	FWD	1	128.4195	Network P2p

Ping output for equipment on VLAN 4092

To verify connectivity, issue **ping** commands to devices on VLAN 4092 (Data VLAN) as shown in Example 5-119.

Example 5-119 Ping verification for equipment on VLAN 4092

```
Nexus5548core_1# ping 10.1.4.243
PING 10.1.4.243 (10.1.4.243): 56 data bytes
64 bytes from 10.1.4.243: icmp_seq=0 ttl=254 time=1.838 ms
64 bytes from 10.1.4.243: icmp_seq=1 ttl=254 time=0.509 ms
64 bytes from 10.1.4.243: icmp_seq=2 ttl=254 time=0.672 ms
64 bytes from 10.1.4.243: icmp_seq=3 ttl=254 time=4.894 ms
64 bytes from 10.1.4.243: icmp_seq=4 ttl=254 time=18.616 ms

--- 10.1.4.243 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.509/5.305/18.616 ms

Nexus5548core_1# ping 10.1.4.238
PING 10.1.4.238 (10.1.4.238): 56 data bytes
64 bytes from 10.1.4.238: icmp_seq=0 ttl=254 time=2.078 ms
64 bytes from 10.1.4.238: icmp_seq=1 ttl=254 time=0.7 ms
64 bytes from 10.1.4.238: icmp_seq=2 ttl=254 time=1.139 ms
64 bytes from 10.1.4.238: icmp_seq=3 ttl=254 time=7.184 ms
64 bytes from 10.1.4.238: icmp_seq=4 ttl=254 time=9.611 ms

--- 10.1.4.238 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.7/4.142/9.611 ms

Nexus5548core_1# ping 10.1.4.10
```

```

PING 10.1.4.10 (10.1.4.10): 56 data bytes
64 bytes from 10.1.4.10: icmp_seq=0 ttl=63 time=0.547 ms
64 bytes from 10.1.4.10: icmp_seq=1 ttl=63 time=0.967 ms
64 bytes from 10.1.4.10: icmp_seq=2 ttl=63 time=0.777 ms
64 bytes from 10.1.4.10: icmp_seq=3 ttl=63 time=7.307 ms
64 bytes from 10.1.4.10: icmp_seq=4 ttl=63 time=9.598 ms

--- 10.1.4.10 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.547/3.839/9.598 ms

```

5.4.9 Full configuration files

This section shows the configuration on all of the devices in the Network Topology diagram.

EN4093flex-1

Example 5-120 lists the configuration for the EN4093flex-1 switch.

Example 5-120 EN4093-1 switch configuration file

```

version "7.3.1"
switch-type "IBM Flex System Fabric EN4093 10Gb Scalable Switch"
!
!

snmp-server name "en4093flex_1"
!
!
hostname "en4093flex_1"
!
!
interface port INTA1
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port INTB1
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT4
    name "ISL hlthchk"
    pvid 4000
    exit
!
interface port EXT7
    name "ISL"
    tagging
    pvid 4094
    exit
!

```

```

interface port EXT8
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT9
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT10
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT15
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT16
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT17
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT18
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT19
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT20
    name "Link to g8264tor_2"

```

```

        tagging
        tag-pvid
        pvid 4092
        exit
    !
interface port EXT21
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT22
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
vlan 1
    member INTA2-INTA14,INTB2-INTB14,EXT1-EXT3,EXT5-EXT6
    no member INTA1,INTB1,EXT4,EXT7-EXT10,EXT15-EXT22
!
vlan 4000
    enable
    name "ISL hlthchk"
    member EXT4
!
vlan 4092
    enable
    name "DATA"
    member INTA1,INTB1,EXT7-EXT10,EXT15-EXT22
!
vlan 4094
    enable
    name "ISL"
    member EXT7-EXT10
!
!
spanning-tree stp 125 vlan 4000
!
spanning-tree stp 126 vlan 4092
!
no spanning-tree stp 127 enable
spanning-tree stp 127 vlan 4094
!
!
interface port EXT7
    lacp mode active
    lacp key 1000
!
interface port EXT8
    lacp mode active
    lacp key 1000
!

```

```

interface port EXT9
    lacp mode active
    lacp key 1000
!
interface port EXT10
    lacp mode active
    lacp key 1000
!
interface port EXT15
    lacp mode active
    lacp key 2000
!
interface port EXT16
    lacp mode active
    lacp key 2000
!
interface port EXT17
    lacp mode active
    lacp key 2000
!
interface port EXT18
    lacp mode active
    lacp key 2000
!
interface port EXT19
    lacp mode active
    lacp key 2000
!
interface port EXT20
    lacp mode active
    lacp key 2000
!
interface port EXT21
    lacp mode active
    lacp key 2000
!
interface port EXT22
    lacp mode active
    lacp key 2000
!
failover enable
failover trigger 1 mmon monitor admin-key 2000
failover trigger 1 mmon control member INTA1-INTB14
failover trigger 1 enable
!
!
!
vlag enable
vlag tier-id 1
vlag isl vlan 4094
vlag hlthchk peer-ip 1.1.1.2
vlag isl adminkey 1000
vlag adminkey 2000 enable
!
!

```

```

!
!
!
!
!
!
!
lldp enable
!
interface ip 40
    ip address 1.1.1.1 255.255.255.0
    vlan 4000
    enable
    exit
!
interface ip 92
    ip address 10.1.4.238 255.255.255.0
    vlan 4092
    enable
    exit
!
!
!
!
!
ntp enable
ntp ipv6 primary-server fe80::211:25ff:fec3:9b69 MGT
ntp interval 15
ntp authenticate
ntp primary-key 8811
!
ntp message-digest-key 8811 md5-ekey
1e389d20083088209635f6e3cb802bd2b52a41c0125c9904874d06d2a3af9d16341b4054daa0d14523
ca25ad2e9ec7d8ef2248b85c18a59a2436918a0ee41cea
!
ntp trusted-key 8811
!
end

```

EN4093flex_2

Example 5-121 lists the configuration for the EN4093flex_2 switch.

Example 5-121 EN4093flex_2 switch configuration

```

version "7.3.1"
switch-type "IBM Flex System Fabric EN4093 10Gb Scalable Switch"
!
!

snmp-server name "en4093flex_2"
!
!
hostname "en4093flex_2"
!
!

```



```

interface port INTA1
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port INTB1
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT4
    name "ISL hlthchk"
    pvid 4000
    exit
!
interface port EXT7
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT8
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT9
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT10
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT15
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT16
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!

```

```

interface port EXT17
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT18
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT19
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT20
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT21
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT22
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
vlan 1
    member INTA2-INTA14,INTB2-INTB14,EXT1-EXT3,EXT5-EXT6
    no member INTA1,INTB1,EXT4,EXT7-EXT10,EXT15-EXT22
!
vlan 4000
    enable
    name "ISL hlthchk"
    member EXT4
!
vlan 4092
    enable
    name "DATA"
    member INTA1,INTB1,EXT7-EXT10,EXT15-EXT22

```

```

!
vlan 4094
    enable
    name "ISL"
    member EXT7-EXT10
!
!
spanning-tree stp 125 vlan 4000
!
spanning-tree stp 126 vlan 4092
!
no spanning-tree stp 127 enable
spanning-tree stp 127 vlan 4094
!
!
no logging console
!
interface port EXT7
    lacp mode active
    lacp key 1000
!
interface port EXT8
    lacp mode active
    lacp key 1000
!
interface port EXT9
    lacp mode active
    lacp key 1000
!
interface port EXT10
    lacp mode active
    lacp key 1000
!
interface port EXT15
    lacp mode active
    lacp key 2000
!
interface port EXT16
    lacp mode active
    lacp key 2000
!
interface port EXT17
    lacp mode active
    lacp key 2000
!
interface port EXT18
    lacp mode active
    lacp key 2000
!
interface port EXT19
    lacp mode active
    lacp key 2000
!
interface port EXT20
    lacp mode active

```

```

    lacp key 2000
!
interface port EXT21
    lacp mode active
    lacp key 2000
!
interface port EXT22
    lacp mode active
    lacp key 2000
!
failover enable
failover trigger 1 mmon monitor admin-key 2000
failover trigger 1 mmon control member INTA1-INTB14
failover trigger 1 enable
!
!
!
vlag enable
vlag tier-id 1
vlag isl vlan 4094
vlag hlthchk peer-ip 1.1.1.1
vlag isl adminkey 1000
vlag adminkey 2000 enable
!
!
!
!
!
!
!
!
!
lldp enable
!
interface ip 40
    ip address 1.1.1.2 255.255.255.0
    vlan 4000
    enable
    exit
!
interface ip 92
    ip address 10.1.4.239 255.255.255.0
    vlan 4092
    enable
    exit
!
!
!
!
!
ntp enable
ntp ipv6 primary-server fe80::211:25ff:fec3:9b69 MGT
ntp interval 15
ntp authenticate
ntp primary-key 8811

```

```
!  
ntp message-digest-key 8811 md5-ekey  
ef9d8bb6cf808aa2b6b6e2f70c3029501c9b293eb41d60e5ebbd0fbbd72171ed3c867d24b9976e2052  
771345e26681dc63a675b9033673c9923707f9d0f1c078  
!  
ntp trusted-key 8811  
!  
end
```

G8264tor_1

Example 5-122 lists the configuration for the G8264tor_1 switch.

Example 5-122 G8264tor_1 switch configuration

```
version "7.4.1"  
switch-type "IBM Networking Operating System RackSwitch G8264"  
!  
!  
ssh enable  
!  
  
!  
!  
no system dhcp  
no system default-ip mgt  
hostname "G8264TOR-1"  
!  
!  
interface port 1  
    name "ISL"  
    tagging  
    pvid 4094  
    exit  
!  
interface port 2  
    name "ISL"  
    tagging  
    pvid 4094  
    exit  
!  
interface port 3  
    name "ISL"  
    tagging  
    pvid 4094  
    exit  
!  
interface port 4  
    name "ISL"  
    tagging  
    pvid 4094  
    exit  
!  
interface port 5  
    name "ISL"  
    tagging
```

```

        pvid 4094
        exit
    !
interface port 6
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 7
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 8
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 9
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 10
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 11
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 12
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 13
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 14
    name "ISL"
    tagging
    pvid 4094

```

```

        exit
    !
interface port 15
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 16
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 18
    name "Po5 to Nexus5548Core_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 20
    name "Po5 to Nexus5548Core_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 22
    name "Po6 to Nexus5548Core_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 24
    name "Po6 to Nexus5548Core_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 25
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 26
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit

```

```

!
interface port 27
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 28
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 37
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 38
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 39
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 40
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 64
    name "ISL hlthchk"
    pvid 4000
    exit
!
vlan 1
    member 17-63
    no member 1-16,64
!
vlan 4000
    enable
    name "ISL hlthchk"

```



```

        member 64
    !
vlan 4092
    enable
    name "DATA"
    member 1-16,18,20,22,24-28,37-40
!
vlan 4094
    enable
    name "ISL"
    member 1-16
!
!
portchannel 5 port 18
portchannel 5 port 20
portchannel 5 enable
!
portchannel 6 port 22
portchannel 6 port 24
portchannel 6 enable
!
!
spanning-tree stp 125 vlan 4000
!
spanning-tree stp 126 vlan 4092
!
no spanning-tree stp 127 enable
spanning-tree stp 127 vlan 4094
!
!
interface port 1
    lacp mode active
    lacp key 1000
!
interface port 2
    lacp mode active
    lacp key 1000
!
interface port 3
    lacp mode active
    lacp key 1000
!
interface port 4
    lacp mode active
    lacp key 1000
!
interface port 5
    lacp mode active
    lacp key 1000
!
interface port 6
    lacp mode active
    lacp key 1000
!
interface port 7

```

```

        lacp mode active
        lacp key 1000
    !
interface port 8
    lacp mode active
    lacp key 1000
    !
interface port 9
    lacp mode active
    lacp key 1000
    !
interface port 10
    lacp mode active
    lacp key 1000
    !
interface port 11
    lacp mode active
    lacp key 1000
    !
interface port 12
    lacp mode active
    lacp key 1000
    !
interface port 13
    lacp mode active
    lacp key 1000
    !
interface port 14
    lacp mode active
    lacp key 1000
    !
interface port 15
    lacp mode active
    lacp key 1000
    !
interface port 16
    lacp mode active
    lacp key 1000
    !
interface port 25
    lacp mode active
    lacp key 2002
    !
interface port 26
    lacp mode active
    lacp key 2002
    !
interface port 27
    lacp mode active
    lacp key 2002
    !
interface port 28
    lacp mode active
    lacp key 2002
    !

```

```

interface port 37
    lacp mode active
    lacp key 2002
!
interface port 38
    lacp mode active
    lacp key 2002
!
interface port 39
    lacp mode active
    lacp key 2002
!
interface port 40
    lacp mode active
    lacp key 2002
!
!
!
vlag enable
vlag tier-id 2
vlag isl vlan 4094
vlag hlthchk peer-ip 1.1.1.2
vlag isl adminkey 1000
vlag adminkey 2002 enable
!
!
!
!
!
!
!
!
!
!
interface ip 40
    ip address 1.1.1.1 255.255.255.0
    vlan 4000
    enable
    exit
!
interface ip 92
    ip address 10.1.4.243 255.255.255.0
    vlan 4092
    enable
    exit
!
interface ip 128
    ip address 172.25.101.243
    enable
    exit
!
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
!
!

```

```
!  
!  
!  
!  
end
```

G8264tor_2

Example 5-123 lists the configuration for the G8264tor_2 switch.

Example 5-123 G8264tor_2 switch configuration

```
version "7.4.1"  
switch-type "IBM Networking Operating System RackSwitch G8264"  
!  
!  
ssh enable  
!  
  
!  
!  
no system dhcp  
no system default-ip mgt  
hostname "G8264TOR-2"  
!  
!  
interface port 1  
    name "ISL"  
    tagging  
    exit  
!  
interface port 2  
    name "ISL"  
    tagging  
    exit  
!  
interface port 3  
    name "ISL"  
    tagging  
    exit  
!  
interface port 4  
    name "ISL"  
    tagging  
    exit  
!  
interface port 5  
    name "ISL"  
    tagging  
    exit  
!  
interface port 6  
    name "ISL"  
    tagging  
    exit  
!
```

```

interface port 7
    name "ISL"
    tagging
    exit
!
interface port 8
    name "ISL"
    tagging
    exit
!
interface port 9
    name "ISL"
    tagging
    exit
!
interface port 10
    name "ISL"
    tagging
    exit
!
interface port 11
    name "ISL"
    tagging
    exit
!
interface port 12
    name "ISL"
    tagging
    exit
!
interface port 13
    name "ISL"
    tagging
    exit
!
interface port 14
    name "ISL"
    tagging
    exit
!
interface port 15
    name "ISL"
    tagging
    exit
!
interface port 16
    name "ISL"
    tagging
    exit
!
interface port 18
    name "Po5 to Nexus5548Core_2"
    tagging
    tag-pvid
    pvid 4092

```

```

        exit
    !
interface port 20
    name "Po5 to Nexus5548Core_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 22
    name "Po6 to Nexus5548Core_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 24
    name "Po6 to Nexus5548Core_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 25
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 26
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 27
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 28
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 37
    name "Link to EN4093-2"
    tagging
    tag-pvid

```

```

    pvid 4092
    exit
!
interface port 38
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 39
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 40
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 64
    name "ISL hlthchk"
    pvid 4000
    exit
!
vlan 1
    member 1-63
    no member 64
!
vlan 4000
    enable
    name "ISL hlthchk"
    member 64
!
vlan 4092
    enable
    name "DATA"
    member 1-16,18,20,22,24-28,37-40
!
vlan 4094
    enable
    name "ISL"
    member 1-16
!
!
portchannel 5 port 18
portchannel 5 port 20
portchannel 5 enable
!
portchannel 6 port 22
portchannel 6 port 24

```

```

portchannel 6 enable
!
!
spanning-tree stp 125 vlan 4000
!
spanning-tree stp 126 vlan 4092
!
no spanning-tree stp 127 enable
spanning-tree stp 127 vlan 4094
!
!
interface port 1
    lacp mode active
    lacp key 1000
!
interface port 2
    lacp mode active
    lacp key 1000
!
interface port 3
    lacp mode active
    lacp key 1000
!
interface port 4
    lacp mode active
    lacp key 1000
!
interface port 5
    lacp mode active
    lacp key 1000
!
interface port 6
    lacp mode active
    lacp key 1000
!
interface port 7
    lacp mode active
    lacp key 1000
!
interface port 8
    lacp mode active
    lacp key 1000
!
interface port 9
    lacp mode active
    lacp key 1000
!
interface port 10
    lacp mode active
    lacp key 1000
!
interface port 11
    lacp mode active
    lacp key 1000
!

```



```

interface port 12
    lacp mode active
    lacp key 1000
!
interface port 13
    lacp mode active
    lacp key 1000
!
interface port 14
    lacp mode active
    lacp key 1000
!
interface port 15
    lacp mode active
    lacp key 1000
!
interface port 16
    lacp mode active
    lacp key 1000
!
interface port 25
    lacp mode active
    lacp key 2002
!
interface port 26
    lacp mode active
    lacp key 2002
!
interface port 27
    lacp mode active
    lacp key 2002
!
interface port 28
    lacp mode active
    lacp key 2002
!
interface port 37
    lacp mode active
    lacp key 2002
!
interface port 38
    lacp mode active
    lacp key 2002
!
interface port 39
    lacp mode active
    lacp key 2002
!
interface port 40
    lacp mode active
    lacp key 2002
!
!
!
vlag enable

```

```

vlag tier-id 2
vlag isl vlan 4094
vlag hlthchk peer-ip 1.1.1.1
vlag isl adminkey 1000
vlag adminkey 2002 enable
!
!
!
!
!
!
!
!
!
!
interface ip 40
    ip address 1.1.1.2 255.255.255.0
    vlan 4000
    enable
    exit
!
interface ip 92
    ip address 10.1.4.244 255.255.255.0
    vlan 4092
    enable
    exit
!
interface ip 128
    ip address 172.25.101.244
    enable
    exit
!
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
!
!
!
!
!
!
end

```

Nexus5548core_1 switch

Example 5-124 lists the configuration of the Nexus5548core_1 switch.

Example 5-124 Nexus5548core_1 switch configuration

```

!Command: show running-config
!Time: Tue Oct 16 22:57:10 2012

version 5.2(1)N1(1b)
logging level feature-mgr 0
hostname Nexus5548core_1

```

```

feature telnet
cfs ipv4 distribute
cfs eth distribute
feature interface-vlan
feature lacp
feature lldp

username admin password 5 $1$huQeFTJf$dYim2oGvqYAGk3THH5KP.0 role network-admin

banner motd #Nexus 5000 Switch
#

no ip domain-lookup
class-map type qos class-fcoe
class-map type queuing class-fcoe
    match qos-group 1
class-map type queuing class-all-flood
    match qos-group 2
class-map type queuing class-ip-multicast
    match qos-group 2
class-map type network-qos class-fcoe
    match qos-group 1
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
snmp-server user admin network-admin auth md5 0x50d80b5959ad2a911a11fcaa8453db8a
priv 0x50d80b5959ad2a911a11fcaa8453db8a localizedkey

vrf context management
    ip route 0.0.0.0/0 172.25.1.1
vlan 1
vlan 4092
    name DATA_VLAN
spanning-tree vlan 4092 priority 8192
port-profile default max-ports 512

interface Vlan1

interface Vlan4092
    no shutdown
    ip address 10.1.4.249/24

interface port-channel5
    switchport mode trunk
    switchport trunk native vlan 4092
    switchport trunk allowed vlan 4092
    speed auto

interface port-channel6
    switchport mode trunk
    switchport trunk native vlan 4092
    switchport trunk allowed vlan 4092
    speed auto

```

```

interface port-channel100
  description Switch-to-Switch link
  switchport mode trunk
  switchport trunk allowed vlan 4092
  spanning-tree port type network

interface Ethernet1/1

interface Ethernet1/2

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7
  description Po5 to G8264tor_1
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  channel-group 5

interface Ethernet1/8
  description Po5 to G8264tor_1
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  channel-group 5

interface Ethernet1/9
  description Po6 to G8264tor_2
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  channel-group 6

interface Ethernet1/10
  description Po6 to G8264tor_2
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  channel-group 6

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

```

```
interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17
  description Po100 to Nexus5548core_2
  switchport mode trunk
  switchport trunk allowed vlan 4092
  speed auto
  channel-group 100 mode active

interface Ethernet1/18

interface Ethernet1/19
  description Po100 to Nexus5548core_2
  switchport mode trunk
  switchport trunk allowed vlan 4092
  speed auto
  channel-group 100 mode active

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface mgmt0
  ip address 172.25.101.249/16
  cli alias name wr copy run start
  line console
  line vty
  boot kickstart bootflash:/n5000-uk9-kickstart.5.2.1.N1.1b.bin
  boot system bootflash:/n5000-uk9.5.2.1.N1.1b.bin
```

Nexus5548core_2 switch

Example 5-125 lists the configuration of the Nexus5548core_2 switch.

Example 5-125 Nexus5548core_2 switch configuration

```
!Command: show running-config
!Time: Tue Oct 16 22:19:15 2012

version 5.2(1)N1(1b)
logging level feature-mgr 0
hostname Nexus5548core_2

feature telnet
cfs ipv4 distribute
cfs eth distribute
feature interface-vlan
feature lacp
feature lldp

username admin password 5 $1$W5m0kb.B$kFgCTs1WQy/ElfbozmrDt/ role network-admin

banner motd #Nexus 5000 Switch
#

no ip domain-lookup
class-map type qos class-fcoe
class-map type queuing class-fcoe
    match qos-group 1
class-map type queuing class-all-flood
    match qos-group 2
class-map type queuing class-ip-multicast
    match qos-group 2
class-map type network-qos class-fcoe
    match qos-group 1
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
snmp-server user adminnetwork-admin auth md5 0xf6e8ccc23aa981dc5c6c28cfa16eb886
priv 0xf6e8ccc23aa981dc5c6c28cfa16eb886 localizedkey

vrf context management
    ip route 0.0.0.0/0 172.25.1.1
vrf context VPCKeepAlive
vlan 1
vlan 4092
    name DATA_VLAN
spanning-tree vlan 4092 priority 16384
port-profile default max-ports 512

interface Vlan1

interface Vlan4092
    no shutdown
```

```

ip address 10.1.4.200/24

interface port-channel5
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  speed auto

interface port-channel6
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  speed auto

interface port-channel100
  description Switch-to-Switch link
  switchport mode trunk
  switchport trunk allowed vlan 4092
  spanning-tree port type network

interface Ethernet1/1

interface Ethernet1/2

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7
  description Po5 to G8264tor_2
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  channel-group 5

interface Ethernet1/8
  description Po5 to G8264tor_2
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  channel-group 5

interface Ethernet1/9
  description Po6 to G8264tor_1
  switchport mode trunk
  switchport trunk native vlan 4092
  switchport trunk allowed vlan 4092
  channel-group 6

interface Ethernet1/10
  description Po6 to G8264tor_1

```

```

switchport mode trunk
switchport trunk native vlan 4092
switchport trunk allowed vlan 4092
channel-group 6

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17
description Po100 to Nexus5548core_1
switchport mode trunk
switchport trunk allowed vlan 4092
speed auto
channel-group 100 mode active

interface Ethernet1/18

interface Ethernet1/19
description Po100 to Nexus5548core_1
switchport mode trunk
switchport trunk allowed vlan 4092
speed auto
channel-group 100 mode active

interface Ethernet1/20

interface Ethernet1/21
shutdown

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

```



```
interface Ethernet1/31

interface Ethernet1/32

interface Ethernet2/1

interface Ethernet2/2

interface Ethernet2/3

interface Ethernet2/4

interface Ethernet2/5

interface Ethernet2/6

interface Ethernet2/7

interface Ethernet2/8

interface Ethernet2/9

interface Ethernet2/10

interface Ethernet2/11

interface Ethernet2/12

interface Ethernet2/13

interface Ethernet2/14

interface Ethernet2/15

interface Ethernet2/16

interface mgmt0
  ip address 172.25.101.200/16

interface loopback1
  ip address 192.168.1.1/24
cli alias name wr copy run start
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.2.1.N1.1b.bin
boot system bootflash:/n5000-uk9.5.2.1.N1.1b.bin
```

5.5 Fully redundant with Open Shortest Path First (OSPF)

This section details the implementation of a fully redundant configuration that uses the Layer-3 routing protocol OSPF.

5.5.1 Topology and requirements

This implementation scenario uses the Layer-3 routing protocol OSPF to provide network connectivity to the G8264 switches. Although this design is different from all the presented Layer-2 Implementation scenarios, the goal of providing a fully redundant infrastructure to the compute nodes still applies. If you have upstream Cisco equipment and prefer to limit the exposure of Layer-2 to your core or aggregation layer, you can implement OSPF instead, but there are some caveats.

This approach has the following advantages:

- ▶ Limited Layer-2 exposure to network infrastructure equipment, limiting the ability of a mis-configuration resulting in a broadcast storm, ARP flooding, or other negative consequence of Layer-2
- ▶ OSPF builds adjacency matrixes and adjusts automatically to down equipment or links
- ▶ Placing Layer 3 IP routing on a switch closer to the servers allows for cross-subnet traffic at that level, freeing up the upstream router to handle just in-bound and out-bound traffic
- ▶ Because IBM System networking switches use ASICs for forwarding layer 3 packets, cross-subnet traffic can be routed within the switch at wirespeed layer 2 performance rates

This approach has these disadvantages:

- ▶ Less flexibility in exposing compute nodes to VLANs that can exist on other switches, either physically or geographically separated
- ▶ Applications that specifically require Layer-2 adjacency for functionality, such as virtual machine based mobility between hypervisors, do not function between differing chassis without Layer-2 adjacency
- ▶ IPv4 subnet address allocation cannot be completely efficient from an address use perspective

Components used

The following components are used in the example configuration:

- ▶ Cisco Nexus 5548UP (Qty. 2)
- ▶ IBM G8264 RackSwitch (Qty. 2)
- ▶ IBM Flex System Fabric EN4093/R 10Gb Scalable Switch (Qty. 2)

5.5.2 Network diagram and physical setup

Figure 5-6 show the network topology diagram for the fully redundant topology using OSPF.

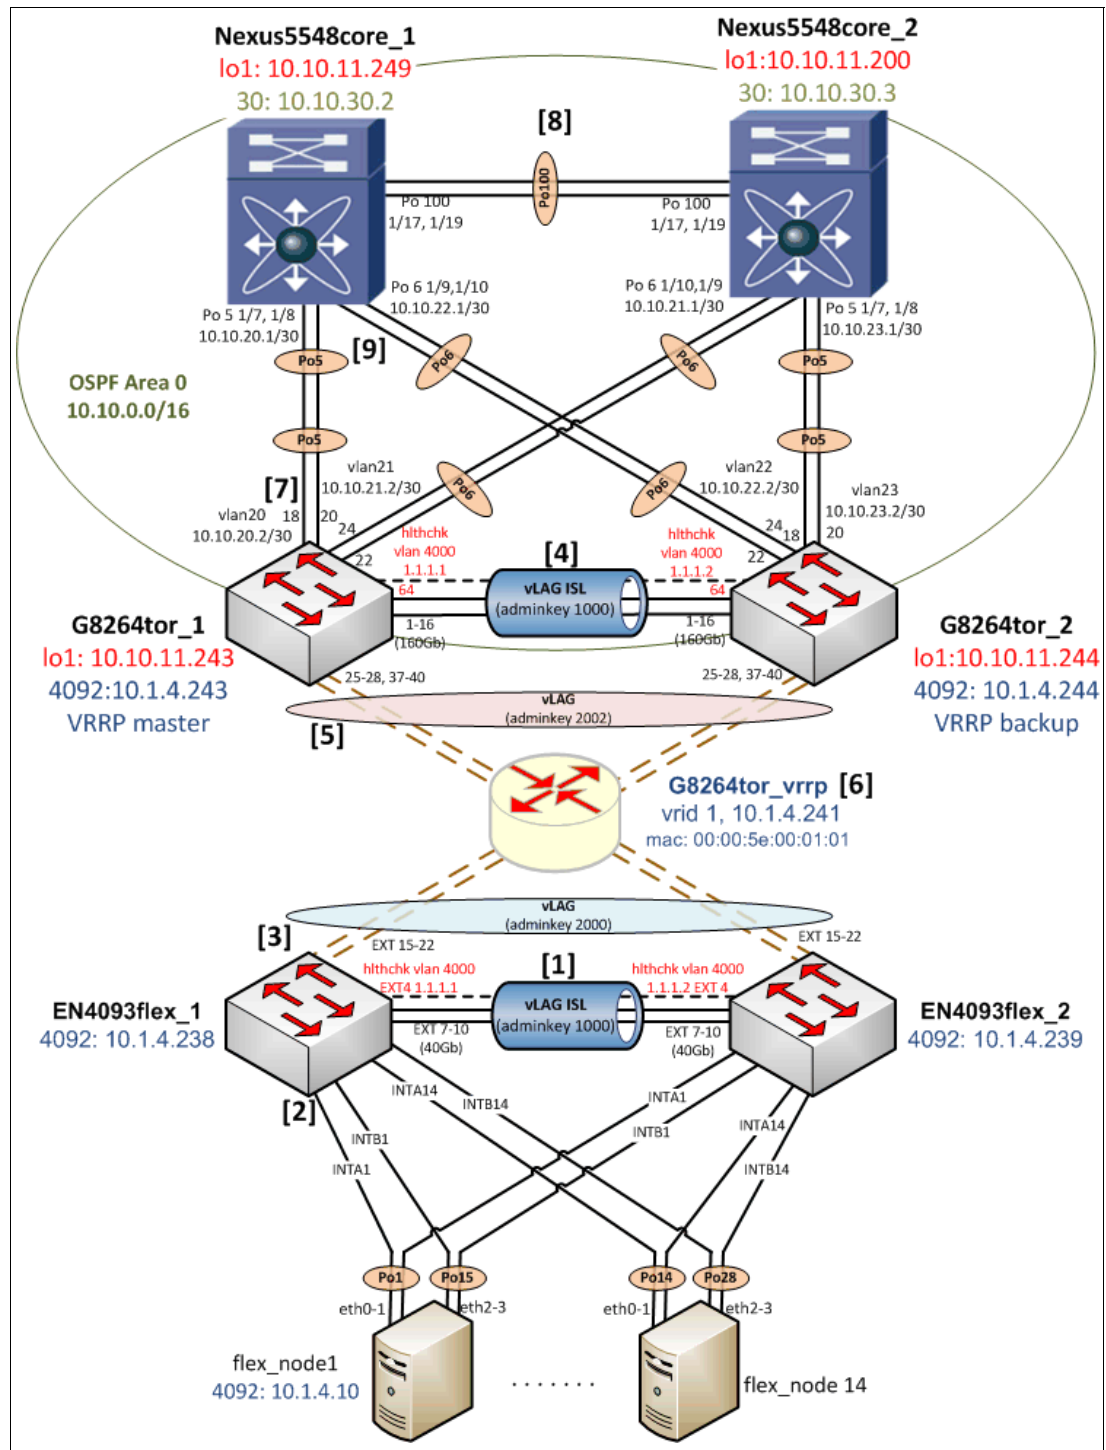


Figure 5-6 Network topology diagram for fully redundant topology using OSPF

Verify the physical cabling between the EN4093/R switches and G8264s. The example environment uses four IBM QSFP+ DAC Break Out Cables from the EN4093/R switches to

the upstream G8264s. This requires that the EN4093/R switches be licensed for these particular features so that the ports can be used.

- ▶ Four 1m IBM QSFP+-to-QSFP+ Cables were used to form the 160 Gb ISL between the G8264 switches.
- ▶ 10Gb SFP+ DAC cables were used for all other connections in the diagram.

5.5.3 EN4093flex_1 configuration

Begin the implementation of this scenario on the IBM Flex System Fabric EN4093/R switches, working up the diagram in Figure 5-6 on page 189. Each step provides the commands necessary and are labeled according to the numbering schema in the diagram.

General configuration

1. Create the ISL hlthchk, ISL data, and Data VLANs as shown in Example 5-126. Give them descriptive names, assign them to spanning-tree groups, and enable them. You can elect to allow the switch itself to create STP instances for you. The examples shows manually creating them instead.

Example 5-126 Create ISL hlthchk, Data, and ISL VLANs on EN4093flex_1

```
configure terminal
vlan 4000
    enable
    name "ISL hlthchk"
    stg 125
    exit
vlan 4092
    enable
    name "DATA"
    stg 126
    exit
vlan 4094
    enable
    name "ISL"
    stg 127
    exit
```

2. Assign IP addresses for both the ISL Healthcheck and Data VLANs as shown in Example 5-127. This allows you to verify connectivity between the various pieces of equipment when verifying the configuration. In this example, interface ip 40 represents the vLAG Health Check IP address, and interface ip 92 represents an address on the Data VLAN that uses the prefix 10.1.4. The last octet is borrowed from the network diagram's Management address to aid in the identification of which piece of equipment you are verifying connectivity to.

Example 5-127 Create IP interfaces and assigning VLANs and IP addresses on EN4093flex_1

```
configure terminal
interface ip 40
    ip address 1.1.1.1 255.255.255.0
    vlan 4000
    enable
    exit
interface ip 92
    ip address 10.1.4.238 255.255.255.0
```

```
vlan 4092
enable
exit
```

Configuring ISL between EN4093flex switches (step 1)

3. Configure the ISL in Example 5-128 between the EN4093flex switches by configuring them to have a default (untagged) VLAN of 4094, Use an LACP key of 1000 to bundle the ports together in an aggregation, with 802.1q tagging enabled so that L2 VLAN traffic can traverse the ISL. Carry Data VLAN 4092 over these links.

Example 5-128 Initial ISL configuration on EN4093flex_1

```
configure terminal
interface port ext7-ext10
    pvid 4094
    tagging
    exit
vlan 4092
    member ext7-ext10
    exit
interface port ext7-ext10
    lacp key 1000
    lacp mode active
    exit
```

4. Create the dedicated health check VLAN and physical interface shown in Example 5-129 for heartbeats between the EN4093/R switches. This example uses EXT4 as a dedicated interface and VLAN 4000 as the health check for the ISL.

Example 5-129 Creating vLAG hlthchk VLAN and interface on EN4093flex_1

```
configure terminal
vlan 4000
    name "ISL hlthchk"
    enable
    exit
interface port ext4
    pvid 4000
    exit
```

5. Disable STP between the EN4093/R switches and activate a vLAG between them so that they appear as a single entity to upstream and downstream infrastructure as shown in Example 5-130. Reference the LACP key that was configured in the previous step.

Example 5-130 Disable STP and activate ISL vLAG on EN4093flex_1

```
configure terminal
no spanning-tree stp 127 enable
vlag tier-id 1
vlag isl vlan 4094
vlag isl adminkey 1000
vlag hlthchk peer-ip 1.1.1.2
vlag enable
```

Configuring downstream internal node ports (step 2)

6. Configure the downstream node interfaces shown in Example 5-131 to have a default (untagged) VLAN of 4092 (data VLAN), with 802.1q tagging enabled. Add the ability for all member ports to be on VLAN 4092.

Example 5-131 Downstream internal node port configuration on EN4093flex_1

```
configure terminal
interface port inta1-intb14
    pvid 4092
    tagging
    spanning-tree edge
    exit
vlan 4092
    member inta1-intb14
exit
```

7. For redundancy, create two port-channels on each of the 14 nodes. Each port channel aggregates two ports, one from each EN4093flex switch. Port channels 1-14 match the “A” internally labelled ports, and port channels 15-28 match the “B” ports as shown in Example 5-132.

Example 5-132 Node-facing port channel creation and vLAG activation on EN4093flex_1

```
configure terminal
portchannel 1 port inta1
portchannel 1 enable
vlag portchannel 1 enable
portchannel 15 port intb1
portchannel 15 enable
vlag portchannel 15 enable
portchannel 2 port inta2
portchannel 2 enable
vlag portchannel 2 enable
portchannel 16 port intb2
portchannel 16 enable
vlag portchannel 16 enable
portchannel 3 port inta3
portchannel 3 enable
vlag portchannel 3 enable
portchannel 17 port intb3
portchannel 17 enable
vlag portchannel 17 enable
portchannel 4 port inta4
portchannel 4 enable
vlag portchannel 4 enable
portchannel 18 port intb4
portchannel 18 enable
vlag portchannel 18 enable
portchannel 5 port inta5
portchannel 5 enable
vlag portchannel 5 enable
portchannel 19 port intb5
portchannel 19 enable
vlag portchannel 19 enable
portchannel 6 port inta6
```

```
portchannel 6 enable
vlag portchannel 6 enable
portchannel 20 port intb6
portchannel 20 enable
vlag portchannel 20 enable
portchannel 7 port inta7
portchannel 7 enable
vlag portchannel 7 enable
portchannel 21 port intb7
portchannel 21 enable
vlag portchannel 21 enable
portchannel 8 port inta8
portchannel 8 enable
vlag portchannel 8 enable
portchannel 22 port intb8
portchannel 22 enable
vlag portchannel 22 enable
portchannel 9 port inta9
portchannel 9 enable
vlag portchannel 9 enable
portchannel 23 port intb9
portchannel 23 enable
vlag portchannel 23 enable
portchannel 10 port inta10
portchannel 10 enable
vlag portchannel 10 enable
portchannel 24 port intb10
portchannel 24 enable
vlag portchannel 24 enable
portchannel 11 port inta11
portchannel 11 enable
vlag portchannel 11 enable
portchannel 25 port intb11
portchannel 25 enable
vlag portchannel 25 enable
portchannel 12 port inta12
portchannel 12 enable
vlag portchannel 12 enable
portchannel 26 port intb12
portchannel 26 enable
vlag portchannel 26 enable
portchannel 13 port inta13
portchannel 13 enable
vlag portchannel 13 enable
portchannel 27 port intb13
portchannel 27 enable
vlag portchannel 27 enable
portchannel 14 port inta14
portchannel 14 enable
vlag portchannel 14 enable
portchannel 28 port intb14
portchannel 28 enable
vlag portchannel 28 enable
```

Configuring upstream G8264tor facing ports and layer2 failover (step 3)

8. Set up the upstream G8264tor facing ports in Example 5-133 with a default (untagged) VLAN of 4092 (data VLAN). Tag the PVID, and use an LACP key of 2000 to bundle the ports together in an aggregation.

Example 5-133 Upstream G8264tor facing port configuration on EN4093flex_1

```
configure terminal
interface port ext15-ext22
    pvid 4092
    tagging
    tag-pvid
    exit
vlan 4092
    member ext15-ext22
    exit
interface port ext15-ext22
    lacp key 2000
    lacp mode active
    exit
```

9. Activate the vLAG feature for the upstream EN4093/R ports so that the G8264s see the EN4093s as a single, virtualized entity as shown in Example 5-134. Use adminkey 2000, which represents the LACP key that bundles ports EXT15-22 together as one.

Example 5-134 Activating the upstream vLAG on G8264tor facing ports on EN4093flex_1

```
configure terminal
vlag adminkey 2000 enable
```

10. Enable Layer-2 failover in Example 5-135, which shuts down the links to the compute nodes if the uplinks for the EN4093/R switch fail. Doing so ensures that the downstream node is aware of the upstream failure. It can then fail traffic over to the other NIC in the node, which in this case is connected to the other EN4093/R switch in the Enterprise Chassis, ensuring that redundancy is maintained.

Example 5-135 Enabling L2 failover for the compute nodes on EN4093flex_1

```
configure terminal
failover trigger 1 mmon monitor admin-key 2000
failover trigger 1 mmon control member INTA1-INTB14
failover trigger 1 enable
failover enable
```

Now repeat this configuration for EN4093flex_2 on the other I/O module. The only difference between the EN4093flex_1 switch and the EN4093flex_2 switch is the vLAG health check peer address and the Data, and ISL h1thchk VLAN IP addresses. To verify EN4093flex switch configuration, run the **show** commands outlined in 5.5.8, “Verification and show command output” on page 208.

5.5.4 G8264tor_1 configuration

Next, configure the switch named G8264tor_1. Although the G8264 switches are mostly similar from a configuration standpoint, differences exist that warrant more explanation.

General configuration

1. Begin by creating the Point-to-Point L3 VLANs (20 and 21), ISL Healthcheck, ISL data, and Data VLANs as shown in Example 5-136. Give them descriptive names, assign them to spanning-tree groups, and enable them.

Example 5-136 Creating vlan 20, vlan21, ISL hlthchk, Data, and ISL vlans on G8264tor_1

```
configure terminal
vlan 20
    enable
    name "VLAN 20"
    stg 20
vlan 21
    enable
    name "VLAN 21"
    stg 21
vlan 4000
    enable
    name "ISL hlthchk"
    stg 125
    exit
vlan 4092
    enable
    name "Data"
    stg 126
    exit
vlan 4094
    enable
    name "ISL"
    stg 127
    exit
```

2. Assign IP addresses for the ISL Healthcheck, Data, and management VLANs as shown in Example 5-137. interface ip 128 represents the management IP address that is referenced in the Network Topology diagram, and IP gateway 4 is the upstream router interface for the 172 management network. Loopback 1 is created to use as the router-id when building the eventual OSPF adjacencies.

Example 5-137 Creating IP interfaces and assigning VLANs and IP addresses on G8264tor_1

```
configure terminal
interface ip 40
    ip address 1.1.1.1 255.255.255.0
    vlan 4000
    enable
    exit
interface ip 92
    ip address 10.1.4.243 255.255.255.0
    vlan 4092
    enable
    exit
interface ip 128
    ip address 172.25.101.243 255.255.0.0
    enable
    exit
interface loopback 1
```

```
ip address 10.10.11.243 255.255.255.255
enable
exit
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
```

3. Assign IP addresses for VLANs 20 and 21. In this implementation scenario, these VLANs represent the Point-to-Point Layer-3 links between the G8264s and the upstream Nexus equipment used to build the OSPF adjacencies. The example consistently uses /30 networks with the “.1” address on the upstream Nexus pair, and the “.2” address on the G8264s as illustrated in Example 5-138.

Example 5-138 Creating IP interfaces and assigning VLANs and IP addresses for vlan20 and vlan21 on G8264tor_1

```
configure terminal
interface ip 20
  ip address 10.10.20.2 255.255.255.252
  vlan 20
  enable
  exit
interface ip 21
  ip address 10.10.21.2 255.255.255.252
  vlan 21
  enable
  exit
```

Configuring ISL between G8264tor switches (step 4)

4. Configure the ISL between the G8264tor switches as shown in Example 5-139. Make the default (untagged) VLAN 4094, LACP key of 1000 to bundle the ports together in an aggregation, with 802.1q tagging enabled. This configuration allows L2 VLAN traffic to traverse the ISL. Carry Data VLAN 4092 over these links.

Example 5-139 Initial ISL configuration on G8264tor_1

```
configure terminal
interface port 1-16
  pvid 4094
  tagging
  exit
vlan 4092
  member 1-16
  exit
interface port 1-16
  lacp key 1000
  lacp mode active
  exit
```

5. Disable STP between the G8264 switches and activate a vLAG between them so that they appear as a single entity to upstream and downstream infrastructure as shown in Example 5-140. Reference the LACP key that was configured in the previous step.

Example 5-140 Disabling STP and activating ISL vLAG on G8253tor_1

```
configure terminal
no spanning-tree stp 127 enable
vlag tier-id 2
```

```
vlag isl vlan 4094
vlag isl adminkey 1000
vlag hlthchk peer-ip 1.1.1.2
vlag enable
```

Configuring downstream EN4093flex facing ports (step 5)

6. Configure the downstream EN4093flex facing ports in Example 5-141 to have a default (untagged) VLAN of 4092 (data VLAN), with 802.1q tagging enabled. Add the ability for all member ports to be on VLAN 4092.

Example 5-141 Downstream EN4093flex facing port configuration on G8264tor_1

```
configure terminal
interface port 25-28,37-40
    pvid 4092
    tagging
    tag-pvid
    exit
vlan 4092
    member 25-28,37-40
    exit
interface port 25-28,37-40
    lacp key 2002
    lacp mode active
    exit
```

7. Activate the vLAG for the downstream EN4093/R ports so that the 4093s see the G8264s as a single, virtualized entity as shown in Example 5-142. Use adminkey 2002, which represents the LACP key that bundles ports 25-28, and 37-40 together as one.

Example 5-142 Activating the downstream EN4093flex facing vLAG on G8264tor_1

```
configure terminal
vlag adminkey 2002 enable
```

Configuring Virtual Router Redundancy Protocol (VRRP) (step 6)

8. In this scenario, the G8264tor switches function at the interface between layer 2 and layer 3. System administrators typically code a single default gateway on host operating systems. By configuring VRRP on the G8264tor switches, you can create a shared virtual router gateway address 10.1.4.241 that is used by downstream hosts on the data VLAN (4092). G8264tor_1 functions as the primary gateway router. When it is unavailable, G8264tor_2 can seamlessly take over the functions of the 10.1.4.241 gateway router as shown in Example 5-143.

Example 5-143 VRRP configuration on G8264tor_1

```
configure terminal
router vrrp
    enable
!
virtual-router 1 virtual-router-id 1
virtual-router 1 interface 92
virtual-router 1 priority 120
virtual-router 1 address 10.1.4.241
virtual-router 1 enable
```

Configuring upstream Nexus5548core facing ports (step 7)

9. Configure the Layer-3 upstream Nexus5548core facing ports in Example 5-144 with a default (untagged) VLAN of 20 on ports 18 and 20, and a default (untagged) VLAN of 21 on ports 22 and 24.

Example 5-144 Upstream Nexus5548core facing port configuration on G8264tor_1

```
configure terminal
interface port 18,20
    name "Po5 to Nexus5548core_1"
    pvid 20
    exit
interface port 22,24
    name "Po5 to Nexus5548core_2"
    pvid 21
    exit
```

10. Activate link aggregation groups using static port-channeling as shown in Example 5-145. The example uses static port-channeling to illustrate that IBM System Networking equipment inter-operates with an upstream Cisco infrastructure with either LACP or static (no negotiation protocol) port-channeling.

Example 5-145 Creating port-channel 5 and 6 on G8264tor_1

```
configure terminal
portchannel 5 port 18
portchannel 5 port 20
portchannel 5 enable
!
portchannel 6 port 22
portchannel 6 port 24
portchannel 6 enable
```

OSPF configuration

11. Set the **router-id** parameter to the loopback 1 in Example 5-146. By default, OSPF uses the lowest configured IP address on the device in all OSPF neighbor advertisements. Generally, define a loopback “virtual” interface for the device and use this interface address in OSPF neighbor advertisements instead because this interface is not susceptible to physical link failures. The **router-id** parameter is used for OSPF neighbor advertisements.

Example 5-146 Setting the router-id on G8264tor_1

```
configure terminal
ip router-id 10.10.11.243
```

12. Create an OSPF instance and advertise OSPF routes through IP interfaces 20, 21, and 92 in Example 5-147 for area 0. OSPF will begin to build its associated adjacency matrixes after the completion of this step.

Example 5-147 Enabling OSPF process and associated interfaces on G8264tor_1

```
configure terminal
router ospf
    area 0 enable
    enable
!
```

```
interface ip 20
    ip ospf enable
!
interface ip 21
    ip ospf enable
!
interface ip 92
    ip ospf enable
```

5.5.5 G8264tor_2 configuration

Next, configure the switch named G8264tor_2.

General configuration

1. Create the Point-to-Point L3 VLANs (22 and 23), ISL Healthcheck, ISL data, and Data VLANs as shown in Example 5-148. Give them descriptive names, assign them to spanning-tree groups, and enable them.

Example 5-148 Create vlan 22, vlan23, ISL hlthchk, Data, and ISL VLANs on G8264tor_2

```
configure terminal
vlan 22
    enable
    name "VLAN 22"
    stg 22
vlan 23
    enable
    name "VLAN 23"
    stg 23
vlan 4000
    enable
    name "ISL hlthchk"
    stg 125
    exit
vlan 4092
    enable
    name "Data"
    stg 126
    exit
vlan 4094
    enable
    name "ISL"
    stg 127
    exit
```

2. Assign IP addresses for the ISL Healthcheck, Data, and management VLANs in Example 5-149. interface ip 128 represents the management IP address that is referenced in the Network Topology diagram, and IP gateway 4 is the upstream router interface for the 172 management network. Loopback 1 is created to use as the router-id when building the eventual OSPF adjacencies.

Example 5-149 Create IP interfaces, and assign VLANs and IP addresses on G8264tor_2

```
configure terminal
interface ip 40
```

```

        ip address 1.1.1.2 255.255.255.0
        vlan 4000
        enable
        exit
interface ip 92
        ip address 10.1.4.244 255.255.255.0
        vlan 4092
        enable
        exit
interface ip 128
        ip address 172.25.101.244 255.255.0.0
        enable
        exit
interface loopback 1
        ip address 10.10.11.244 255.255.255.255
        enable
        exit
ip gateway 4 address 172.25.1.1
ip gateway 4 enable

```

3. Assign IP addresses for VLANs 22 and 23. In this implementation scenario, these VLANs represent the Point-to-Point Layer-3 links between the G8264s and the upstream Nexus equipment that is used to build the OSPF adjacencies. This example consistently uses /30 networks with the “.1” address on the upstream Nexus pair, and the “.2” address on the G8264’s as illustrated in Example 5-150.

Example 5-150 Create IP interfaces and assign VLANs for vlan22 and vlan23 on G8264tor_2

```

configure terminal
interface ip 22
        ip address 10.10.22.2 255.255.255.252
        vlan 22
        enable
        exit
interface ip 23
        ip address 10.10.23.2 255.255.255.252
        vlan 23
        enable
        exit

```

Configuring ISL between G8264tor switches (step 4)

4. Configure the ISL between the G8264tor switches in Example 5-151. Make the default (untagged) VLAN 4094, LACP key of 1000 to bundle the ports together in an aggregation, with 802.1q tagging enabled so that L2 VLAN traffic can traverse the ISL. Carry Data VLAN 4092 over these links.

Example 5-151 Initial ISL configuration on G8264tor_2

```

configure terminal
interface port 1-16
        pvid 4094
        tagging
        exit
vlan 4092
        member 1-16
        exit

```

```
interface port 1-16
  lacp key 1000
  lacp mode active
exit
```

5. Disable STP between the G8264 switches and activate a vLAG between them so that they appear as a single entity to upstream and downstream infrastructure as shown in Example 5-152. Reference the LACP key that was configured in the previous step.

Example 5-152 Disable STP and activate ISL vLAG on G8264tor_2

```
configure terminal
no spanning-tree stp 127 enable
vlag tier-id 2
vlag isl vlan 4094
vlag isl adminkey 1000
vlag hlthchk peer-ip 1.1.1.1
vlag enable
```

Configuring downstream EN4093/R facing ports (step 5)

6. Configure the downstream EN4093flex facing ports as shown in Example 5-153 to have a default (untagged) VLAN of 4092 (data VLAN), with 802.1q tagging enabled. Add the ability for all member ports to be on VLAN 4092.

Example 5-153 Downstream EN4093flex facing port configuration on G8264tor_2

```
configure terminal
interface port 25-28,37-40
  pvid 4092
  tagging
  tag-pvid
exit
vlan 4092
  member 25-28,37-40
exit
interface port 25-28,37-40
  lacp key 2002
  lacp mode active
exit
```

7. Activate the vLAG for the downstream EN4093/R ports so that the 4093s see the G8264s as a single, virtualized entity as shown in Example 5-154. Use adminkey 2002, which represents the LACP key that bundles ports 25-28, and 37-40 together as one.

Example 5-154 Activate the downstream EN4093flex facing vLAG on G8264tor_2

```
configure terminal
vlag adminkey 2002 enable
```

Configuring VRRP (step 6)

8. Configure VRRP on G8264tor_2. G8264tor_2 functions as the secondary VRRP gateway router for the 10.1.4.241 IP address as shown in Example 5-155.

Example 5-155 VRRP configuration on G8264tor_2

```
configure terminal
router vrrp
  enable
!
  virtual-router 1 virtual-router-id 1
  virtual-router 1 interface 92
  virtual-router 1 priority 110
  virtual-router 1 address 10.1.4.241
  virtual-router 1 enable
```

Configuring upstream Nexus5548core facing ports (step 7)

9. Set up the Layer-3 upstream ports to the Nexus pair in Example 5-156 with a default (untagged) VLAN of 23 on ports 18 and 20, and a default (untagged) VLAN of 22 on ports 22 and 24.

Example 5-156 Upstream Nexus5548core facing layer3 configuration on G8264tor_2

```
configure terminal
interface port 18,20
  name "Po5 to Nexus5548core_1"
  pvid 23
  exit
interface port 22,24
  name "Po5 to Nexus5548core_2"
  pvid 22
  exit
```

10. Activate the link aggregation groups by using static port-channeling as shown in Example 5-157.

Example 5-157 Create port-channel interfaces 5 and 6 on G8264tor_2

```
configure terminal
portchannel 5 port 18
portchannel 5 port 20
portchannel 5 enable
!
portchannel 6 port 22
portchannel 6 port 24
portchannel 6 enable
```

OSPF configuration

11. Set the **router-id** parameter to be the loopback 1 IP address as shown in Example 5-158.

Example 5-158 Set router-id on G8264tor_2

```
configure terminal
ip router-id 10.10.11.244
```

12. Create an OSPF instance and enable IP interfaces 22, 23, and 92 for area 0 as shown in Example 5-159.

Example 5-159 Enabling OSPF process and associated interfaces on G8264tor_2

```
configure terminal
router ospf
  area 0 enable
  enable
!
interface ip 22
  ip ospf enable
!
interface ip 23
  ip ospf enable
!
interface ip 92
  ip ospf enable
```

5.5.6 Nexus5548core_1 switch configuration

Next, configure the Nexus5548core_1 switch.

General configuration

13. Be sure that the following features are enabled as shown in Example 5-160, particularly OSPF because that is the routing protocol that is used in this implementation scenario.

Example 5-160 Enable NX-OS features on Nexus5548core_1

```
configure terminal
feature ospf
feature interface-vlan
feature lacp
feature lldp
```

14. Create a server VLAN as shown in Example 5-161. An important difference with this scenario is the absence of vlan 4092 (data VLAN) on the Nexus5548core switches. Because layer 3 connectivity is used between Nexus5548core and G8264tor switches, the data VLAN does not extend up to the Nexus5548core switches. vlan 30 is created as a “Server” network that is used on the Nexus5548core switches only. This is used to demonstrate connectivity to compute nodes on vlan 4092.

Example 5-161 Vlan 30 configuration on Nexus5548core_1

```
configure terminal
vlan 30
  name Server
```

15. Create the SSVI for the server VLAN as shown in Example 5-162.

Example 5-162 IP address configuration for vlan30 on Nexus5548core_1

```
configure terminal
interface Vlan30
  ip address 10.10.30.2/24
  no shutdown
```

OSPF configuration

16. Create the loopback 1 interface and implement OSPF process 100. With the proliferation of up to 40 Gb Ethernet as of this writing, set the reference bandwidth that is used in cost calculation to 100 Gbps. Assign vlan30 to router ospf 100 area 0 as shown in Example 5-163.

Example 5-163 OSPF instance configuration on Nexus5548core_1

```
configure terminal
interface loopback1
  description OSPF router-id
  ip address 10.10.11.249/32
router ospf 100
  router-id 10.10.11.249
  log-adjacency-changes
  auto-cost reference-bandwidth 100 Gbps
interface Vlan30
  ip router ospf 100 area 0.0.0.0
```

VRRP configuration

17. Configure vlan 30 to participate in OSPF and configure VRRP to serve as a protection mechanism in case one of the Nexus switches lose network connectivity as shown in Example 5-164. This is NOT shown in the Network Topology diagram because it is assumed that these sorts of mechanisms are already present in your existing infrastructure, but are shown here to be thorough.

Example 5-164 VRRP configuration for vlan30 on Nexus5548core_1

```
configure terminal
vrrp 1
  priority 200
  address 10.10.30.1
  no shutdown
```

Configuring port channel between Nexus5548core switches (step 8)

18. Configure the physical interfaces that comprise the switch-to-switch link between Nexus5548core_1 and Nexus5548core_2 as shown in Example 5-165. Use port-channel100 and make it a Layer-2 link between the switches. Although you can use a vPC peer link between both Nexus core switches instead, this example shows a Layer-2 port-channel between them as a difference between this scenario and the virtualized chassis technology design.

Example 5-165 Switch-to-switch link physical and logical interface configuration on Nexus5548core_1

```
configure terminal
interface Ethernet1/17
  description Po100 to Nexus5548core_2
  switchport mode trunk
  switchport trunk allowed vlan 30
  channel-group 100 mode active
interface Ethernet1/19
  description Po100 to Nexus5548core_2
  switchport mode trunk
  switchport trunk allowed vlan 30
```

```
channel-group 100 mode active
interface port-channel100
description Switch-to-Switch Link
switchport mode trunk
switchport trunk allowed vlan 30
spanning-tree port type network
```

Configuring downstream G8264tor facing ports (step 9)

19. Finally, for Nexus5548core_1, configure the downstream physical and logical interfaces in Example 5-166. Bundle interfaces Ethernet1/7 and Ethernet1/8 in static aggregation Po5, and interfaces Ethernet1/9 and Ethernet1/10 in static aggregation Po6. Associate these port-channel interfaces with OSPF process 100, area 0.0.0.0.

Example 5-166 Downstream iG8264tor facing interface configuration on Nexus5548core_1

```
configure terminal
interface Ethernet1/7-8
description Po5 to G8264tor_1
no switchport
speed auto
channel-group 5 mode on
interface Ethernet1/9-10
description Po6 to G8264tor_2
no switchport
speed auto
channel-group 6 mode on
interface port-channel5
no switchport
ip address 10.10.20.1/30
ip router ospf 100 area 0.0.0.0
interface port-channel6
no switchport
ip address 10.10.22.1/30
ip router ospf 100 area 0.0.0.0
```

5.5.7 Nexus5548core_2 configuration

This section details configuring the Nexus5548core_2 switch.

General configuration

20. Be sure that the following features are enabled as shown in Example 5-167, particularly OSPF because that is the routing protocol that is used in this implementation scenario.

Example 5-167 Enable NX-OS features on Nexus5548core_2

```
configure terminal
feature ospf
feature interface-vlan
feature lacp
feature lldp
```

21. Create vlan 30 as shown in Example 5-168.

Example 5-168 Create vlan 30 on Nexus5548core_2

```
configure terminal
vlan 30
  name Server
```

22. Create the Switched Virtual Interface (SVI) for the server VLAN as shown in Example 5-169.

Example 5-169 Server vlan30 ip configuration on Nexus5548core_2

```
configure terminal
interface Vlan30
  ip address 10.10.30.3/24
  no shutdown
```

OSPF configuration

23. Create the loopback 1 interface and implement OSPF process 100. Set the reference bandwidth to 100 Gbps, and assign vlan30 to OSPF area 0 as shown in Example 5-170.

Example 5-170 OSPF configuration on Nexus5548core_2

```
configure terminal
interface loopback1
  description OSPF router-id
  ip address 10.10.11.200/32
router ospf 100
  router-id 10.10.11.200
  log-adjacency-changes
  auto-cost reference-bandwidth 100 Gbps
interface Vlan30
  ip router ospf 100 area 0.0.0.0
```

VRRP configuration

24. Configure VRRP to serve as a protection mechanism in case one of the Nexus switches lose network connectivity in Example 5-171. Notice the priority configured which means that the secondary core switch is the backup from a VRRP perspective.

Example 5-171 VRRP configuration on Nexus5548core_2

```
configure terminal
vrrp 1
  priority 150
  address 10.10.30.1
  no shutdown
```

Configuring port channel link between Nexus5548core switches (step 8)

25. Configure the physical interfaces that comprise the switch-to-switch link between the Nexus 5548-1 and 5548-2 switches as shown in Example 5-172. Use port-channel100 and make it a Layer-2 link between the switches.

Example 5-172 Switch-to-switch link physical and logical interface configuration on Nexus5548core_2

```
configure terminal
interface Ethernet1/17
  description Po100 to Nexus5548core_1
  switchport mode trunk
  switchport trunk allowed vlan 30
  channel-group 100 mode active
interface Ethernet1/19
  description Po100 to Nexus5548core_1
  switchport mode trunk
  switchport trunk allowed vlan 30
  channel-group 100 mode active
interface port-channel100
  description Switch-to-Switch link
  switchport mode trunk
  switchport trunk allowed vlan 30
  spanning-tree port type network
```

Configuring downstream G8264tor facing ports (step 9)

26. Finally, for Nexus5548core_2 configure the downstream G8264tor facing physical and logical interfaces as shown in Example 5-173. Bundle interfaces Ethernet1/7 and Ethernet1/8 in static aggregation Po5, and interfaces Ethernet1/9 and Ethernet1/10 in static aggregation Po6. Associate these port-channel interfaces with OSPF process 100, area 0.0.0.0.

Example 5-173 Downstream G8264tor facing interface configuration on Nexus5548core_2

```
configure terminal
interface Ethernet1/7-8
  description Po5 to G8264tor_2
  no switchport
  speed auto
  channel-group 5 mode on
interface Ethernet1/9-10
  description Po6 to G8264tor_1
  no switchport
  speed auto
  channel-group 6 mode on
interface port-channel5
  no switchport
  ip address 10.10.23.1/30
  ip router ospf 100 area 0.0.0.0
interface port-channel6
  no switchport
  ip address 10.10.21.1/30
  ip router ospf 100 area 0.0.0.0
```

5.5.8 Verification and show command output

The following section lists output from common show commands that can aid the network architect in the implementation of this scenario. Perform ping verification of the various IP addresses configured on the equipment for the Data VLAN to ensure that all of the devices can reach each other successfully.

As in the implementation section, the commands begin at the EN4093/R switches and work up the Network Topology diagram to the Cisco Nexus pair.

EN4093/R output

This section lists output from the switch with hostname EN4093flex_1. Similar or identical output exists for the switch with hostname EN4093flex_2.

Show version

The command output in Example 5-174 shows information about the switch and the associated code/firmware level.

Example 5-174 EN4093flex_1 show version output

```
System Information at 23:04:56 Fri Oct 12, 2012
Time zone: No timezone configured
Daylight Savings Time Status: Disabled

IBM Flex System Fabric EN4093 10Gb Scalable Switch

Switch has been up for 1 day, 2 hours, 1 minute and 21 seconds.
Last boot: 21:05:54 Thu Oct 11, 2012 (reset from Telnet/SSH)

MAC address: 6c:ae:8b:bf:6d:00    IP (If 40) address: 1.1.1.1
Internal Management Port MAC Address: 6c:ae:8b:bf:6d:ef
Internal Management Port IP Address (if 128): 172.25.101.238
External Management Port MAC Address: 6c:ae:8b:bf:6d:fe
External Management Port IP Address (if 127):
Software Version 7.3.1.0           (FLASH image1), active configuration.
```

```
Hardware Part Number      : 49Y4272
Hardware Revision         : 02
Serial Number             : Y250VT24M099
Manufacturing Date (WWYY) : 1712
PCBA Part Number         : BAC-00072-01
PCBA Revision            : 0
PCBA Number              : 00
Board Revision           : 02
PLD Firmware Version     : 1.5

Temperature Warning       : 32 C (Warn at 60 C/Recover at 55 C)
Temperature Shutdown     : 32 C (Shutdown at 65 C/Recover at 60 C)
Temperature Inlet        : 27 C
Temperature Exhaust      : 33 C
```

Power Consumption : 54.300 W (12.244 V, 4.435 A)

Switch is in I/O Module Bay 1

Show vlan

Example 5-175 shows output regarding VLAN assignment for all the various ports on the switch.

Example 5-175 EN4093flex_1 show vlan output

VLAN	Name	Status	MGT	Ports
1	Default VLAN	ena	dis	EXT1-EXT3 EXT5 EXT6
4000	ISL hlthchk	ena	dis	EXT4
4092	DATA	ena	dis	INTA1-INTB14 EXT7-EXT10 EXT15-EXT22
4094	ISL	ena	dis	EXT7-EXT10
4095	Mgmt VLAN	ena	ena	EXTM MGT1

Show interface status

Because there is only one compute node in the chassis (in slot 1), all the other internal ports are listed as “down” from a link perspective in the output shown in Example 5-176.

Example 5-176 EN4093flex_1 show interface status output

Alias	Port	Speed	Duplex	Flow Ctrl		Link	Name
				--TX--	--RX--		
INTA1	1	1000	full	no	no	up	INTA1
INTA2	2	1G/10G	full	yes	yes	down	INTA2
INTA3	3	1G/10G	full	yes	yes	down	INTA3
INTA4	4	1G/10G	full	yes	yes	down	INTA4
INTA5	5	1G/10G	full	yes	yes	down	INTA5
INTA6	6	1G/10G	full	yes	yes	down	INTA6
INTA7	7	1G/10G	full	yes	yes	down	INTA7
INTA8	8	1G/10G	full	yes	yes	down	INTA8
INTA9	9	1G/10G	full	yes	yes	down	INTA9
INTA10	10	1G/10G	full	yes	yes	down	INTA10
INTA11	11	1G/10G	full	yes	yes	down	INTA11
INTA12	12	1G/10G	full	yes	yes	down	INTA12
INTA13	13	1G/10G	full	yes	yes	down	INTA13
INTA14	14	1G/10G	full	yes	yes	down	INTA14
INTB1	15	1000	full	no	no	up	INTB1
INTB2	16	1G/10G	full	yes	yes	down	INTB2
INTB3	17	1G/10G	full	yes	yes	down	INTB3
INTB4	18	1G/10G	full	yes	yes	down	INTB4
INTB5	19	1G/10G	full	yes	yes	down	INTB5
INTB6	20	1G/10G	full	yes	yes	down	INTB6
INTB7	21	1G/10G	full	yes	yes	down	INTB7
INTB8	22	1G/10G	full	yes	yes	down	INTB8
INTB9	23	1G/10G	full	yes	yes	down	INTB9
INTB10	24	1G/10G	full	yes	yes	down	INTB10
INTB11	25	1G/10G	full	yes	yes	down	INTB11
INTB12	26	1G/10G	full	yes	yes	down	INTB12
INTB13	27	1G/10G	full	yes	yes	down	INTB13

INTB14	28	1G/10G	full	yes	yes	down	INTB14
EXT1	43	10000	full	no	no	up	EXT1
EXT2	44	10000	full	no	no	up	EXT2
EXT3	45	10000	full	no	no	up	EXT3
EXT4	46	10000	full	no	no	up	ISL h1thchk
EXT5	47	1G/10G	full	no	no	down	EXT5
EXT6	48	1G/10G	full	no	no	down	EXT6
EXT7	49	10000	full	no	no	up	ISL
EXT8	50	10000	full	no	no	up	ISL
EXT9	51	10000	full	no	no	up	ISL
EXT10	52	10000	full	no	no	up	ISL
EXT15	57	10000	full	no	no	up	Link to g8264tor_1
EXT16	58	10000	full	no	no	up	Link to g8264tor_1
EXT17	59	10000	full	no	no	up	Link to g8264tor_1
EXT18	60	10000	full	no	no	up	Link to g8264tor_1
EXT19	61	10000	full	no	no	up	Link to g8264tor_2
EXT20	62	10000	full	no	no	up	Link to g8264tor_2
EXT21	63	10000	full	no	no	up	Link to g8264tor_2
EXT22	64	10000	full	no	no	up	Link to g8264tor_2
EXTM	65	1000	half	yes	yes	down	EXTM
MGT1	66	1000	full	yes	yes	up	MGT1

Show lldp remote-device

The command output in Example 5-177 illustrates the physical topology and verifies that cables are plugged into the ports specified in both the Network Topology diagram, and the configuration specified in the appendix.

Example 5-177 EN4093flex_1 show lldp remote-device output

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
EXT16	3	08 17 f4 33 9d 00	25	G8264TOR-1
EXT15	4	08 17 f4 33 9d 00	26	G8264TOR-1
EXT18	5	08 17 f4 33 9d 00	27	G8264TOR-1
EXT17	6	08 17 f4 33 9d 00	28	G8264TOR-1
EXT21	7	08 17 f4 33 75 00	25	G8264TOR-2
EXT19	8	08 17 f4 33 75 00	26	G8264TOR-2
EXT22	9	08 17 f4 33 75 00	27	G8264TOR-2
EXT20	10	08 17 f4 33 75 00	28	G8264TOR-2
EXT4	12	6c ae 8b bf fe 00	46	en4093flex_2
EXT7	13	6c ae 8b bf fe 00	49	en4093flex_2
EXT8	14	6c ae 8b bf fe 00	50	en4093flex_2
EXT9	15	6c ae 8b bf fe 00	51	en4093flex_2
EXT10	16	6c ae 8b bf fe 00	52	en4093flex_2

Show vlag isl

Example 5-178 shows command output about the status of the ISL between the EN4093/R switches, and the ports that comprise the ISL itself.

Example 5-178 EN4093flex_1 show vlag isl output

ISL_ID	ISL_Vlan	ISL_Trunk	ISL_Members	Link_State	Trunk_State
65	4094	Adminkey 1000	EXT7	UP	UP
			EXT8	UP	UP
			EXT9	UP	UP
			EXT10	UP	UP

Show vlag information

The command output in Example 5-179 shows that the vLAG between the EN4093/R switches and G8264 switches is up and operational as referenced by the LACP admin key of 2000. The ISL between the EN4093/R switches is up as well.

EN4093flex_1 is acting as the admin and operational role of PRIMARY. For centralized vLAG functions, such as vLAG STP, one of the vLAG switches must control the protocol operations. To select the switch that controls the centralized vLAG function, perform role election. The switch with the primary role controls the centralized operation. Role election is non-preemptive. That is, if a primary already exists, another switch that is coming up remains as secondary even if it can become primary based on the role election logic.

Role election is determined by comparing the local vLAG system priority and local system MAC address. The switch with the smaller priority value becomes the vLAG primary switch. If priorities are the same, the switch with the smaller system MAC address becomes the vLAG primary switch. You can configure vLAG priority to anything between <0-65535>. Priority was left at the default value of 0 in all examples.

Example 5-179 EN4093flex_1 show vlag information output

```
vLAG Tier ID: 1
vLAG system MAC: 08:17:f4:c3:dd:00
Local MAC 6c:ae:8b:bf:6d:00 Priority 0 Admin Role PRIMARY (Operational Role
PRIMARY)
Peer MAC 6c:ae:8b:bf:fe:00 Priority 0
Health local 1.1.1.1 peer 1.1.1.2 State UP
ISL trunk id 65
ISL state Up
Startup Delay Interval: 120s (Finished)
```

```
vLAG 65: config with admin key 2000, associated trunk 66, state formed
```

Show vlag adminkey 2000

The output in Example 5-180 shows that the vLAG is formed and enabled using LACP reference key 2000.

Example 5-180 EN4093flex_1 show vlag adminkey 2000 output

```
vLAG is enabled on admin key 2000
Current LACP params for EXT15: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT16: active, Priority 32768, Admin Key 2000, Min-Links 1
```

Current LACP params for EXT17: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT18: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT19: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT20: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT21: active, Priority 32768, Admin Key 2000, Min-Links 1

Current LACP params for EXT22: active, Priority 32768, Admin Key 2000, Min-Links 1

Show lacp information state up

The command output in Example 5-181 shows which ports are participating in an LACP aggregation, and which reference keys are used on those specific interfaces.

Example 5-181 EN4093flex_1 show lacp information state up

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
EXT7	active	1000	1000	yes	32768	49	65	up	1
EXT8	active	1000	1000	yes	32768	49	65	up	1
EXT9	active	1000	1000	yes	32768	49	65	up	1
EXT10	active	1000	1000	yes	32768	49	65	up	1
EXT15	active	2000	2000	yes	32768	57	66	up	1
EXT16	active	2000	2000	yes	32768	57	66	up	1
EXT17	active	2000	2000	yes	32768	57	66	up	1
EXT18	active	2000	2000	yes	32768	57	66	up	1
EXT19	active	2000	2000	yes	32768	57	66	up	1
EXT20	active	2000	2000	yes	32768	57	66	up	1
EXT21	active	2000	2000	yes	32768	57	66	up	1
EXT22	active	2000	2000	yes	32768	57	66	up	1

Show failover trigger 1

The failover output in Example 5-182 shows which ports are monitored, and which ports are shut down if an issue is encountered. In this example, the upstream to G8264 links are monitored with LACP reference key 2000. The control ports are the downstream internal I/O module ports that are used by the Compute Nodes.

Example 5-182 EN4093flex_1 show failover output

```

Failover: On
VLAN Monitor: OFF

Trigger 1 Manual Monitor: Enabled
Trigger 1 limit: 0
Monitor State: Up
Member      Status
-----
adminkey 2000
EXT15      Operational
EXT16      Operational
EXT17      Operational
EXT18      Operational
EXT19      Operational

```

```

EXT20      Operational
EXT21      Operational
EXT22      Operational
Control State: Auto Controlled
Member      Status
-----
INTA1      Operational
INTA2      Operational
INTA3      Operational
INTA4      Operational
INTA5      Operational
INTA6      Operational
INTA7      Operational
INTA8      Operational
INTA9      Operational
INTA10     Operational
INTA11     Operational
INTA12     Operational
INTA13     Operational
INTA14     Operational
INTB1      Operational
INTB2      Operational
INTB3      Operational
INTB4      Operational
INTB5      Operational
INTB6      Operational
INTB7      Operational
INTB8      Operational
INTB9      Operational
INTB10     Operational
INTB11     Operational
INTB12     Operational
INTB13     Operational
INTB14     Operational

```

Trigger 2: Disabled

Trigger 3: Disabled

Trigger 4: Disabled

Trigger 5: Disabled

Trigger 6: Disabled

Trigger 7: Disabled

Trigger 8: Disabled

Show ARP

To verify VRRP configuration on the upstream G8264tor switches, run the **show arp** command on the EN4093flex switches. You can then see that the VRRP ip gateway address 10.1.4.241 is present in the ARP table. The MAC -address used by this IP address is the standard VRRP

mac-address 00-00-5e-00-01-xx where xx is defined by the vrrp virtual-router-id that is defined as 01 as shown in Example 5-183.

Example 5-183 EN4093flex_1 show ARP output

```
en4093flex_1#show arp
Current ARP configuration:
rearp 5
No static ARP configured.

-----
Total number of arp entries : 6
IP address  Flags  MAC address  VLAN  Age  Port
-----
1.1.1.1 P 6c:ae:8b:bf:6d:00 4000
1.1.1.2 6c:ae:8b:bf:fe:00 4000 11 EXT4
10.1.4.238 P 6c:ae:8b:bf:6d:00 4092
10.1.4.241 00:00:5e:00:01:01 4092 1 TRK65
10.1.4.243 08:17:f4:33:9d:00 4092 287 TRK65
10.1.4.244 08:17:f4:33:75:00 4092 279 TRK65
```

Ping output for equipment on VLAN 4092

To verify connectivity, issue **ping** commands to devices on VLAN 4092 (Data VLAN) as shown in Example 5-184. IP address 10.4.1.10 represents a compute node with an operating system installed, flex_node1 on the Network Topology diagram.

Example 5-184 Ping verification for equipment on VLAN 4092

```
en4093flex_1#ping 10.1.4.10 data-port
Connecting via DATA port.
[host 10.1.4.10, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl 255,
tos 0]
10.1.4.10: #1 ok, RTT 1 msec.
10.1.4.10: #2 ok, RTT 0 msec.
10.1.4.10: #3 ok, RTT 1 msec.
10.1.4.10: #4 ok, RTT 0 msec.
10.1.4.10: #5 ok, RTT 0 msec.
Ping finished.

en4093flex_1#ping 10.1.4.239 data-port
Connecting via DATA port.
[host 10.1.4.239, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.239: #1 ok, RTT 4 msec.
10.1.4.239: #2 ok, RTT 1 msec.
10.1.4.239: #3 ok, RTT 2 msec.
10.1.4.239: #4 ok, RTT 3 msec.
10.1.4.239: #5 ok, RTT 1 msec.
Ping finished.

en4093flex_1#ping 10.1.4.243 data-port
Connecting via DATA port.
[host 10.1.4.243, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.243: #1 ok, RTT 1 msec.
10.1.4.243: #2 ok, RTT 1 msec.
10.1.4.243: #3 ok, RTT 2 msec.
```

```
10.1.4.243: #4 ok, RTT 8 msec.  
10.1.4.243: #5 ok, RTT 6 msec.  
Ping finished.
```

```
en4093flex_1#ping 10.1.4.244 data-port  
Connecting via DATA port.  
[host 10.1.4.244, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl  
255, tos 0]  
10.1.4.244: #1 ok, RTT 1 msec.  
10.1.4.244: #2 ok, RTT 2 msec.  
10.1.4.244: #3 ok, RTT 1 msec.  
10.1.4.244: #4 ok, RTT 2 msec.  
10.1.4.244: #5 ok, RTT 0 msec.  
Ping finished.
```

```
en4093flex_1#ping 10.1.4.249 data-port  
Connecting via DATA port.  
[host 10.1.4.241, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl  
255, tos 0]  
10.1.4.241: #1 ok, RTT 2 msec.  
10.1.4.241: #2 ok, RTT 1 msec.  
10.1.4.241: #3 ok, RTT 2 msec.  
10.1.4.241: #4 ok, RTT 1 msec.  
10.1.4.241: #5 ok, RTT 3 msec.  
Ping finished.
```

```
en4093flex_1#ping 10.1.4.200 data-port  
Connecting via DATA port.  
[host 10.1.4.241, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl  
255, tos 0]  
10.1.4.241: #1 ok, RTT 2 msec.  
10.1.4.241: #2 ok, RTT 2 msec.  
10.1.4.241: #3 ok, RTT 2 msec.  
10.1.4.241: #4 ok, RTT 1 msec.  
10.1.4.241: #5 ok, RTT 3 msec.  
Ping finished
```

G8264 output

This section lists output from the switch with hostname G8264tor_1, noting specific differences on G8264tor_2 when applicable.

Show version

Example 5-185 shows information about the switch and the associated code/firmware level.

Example 5-185 G8264tor_1 show version output

```
System Information at 21:55:21 Wed Oct 24, 2012  
Time zone: No timezone configured  
Daylight Savings Time Status: Disabled
```

```
IBM Networking Operating System RackSwitch G8264
```

```
Switch has been up for 0 days, 3 hours, 55 minutes and 35 seconds.  
Last boot: 18:01:02 Wed Oct 24, 2012 (reset from Telnet/SSH)
```

MAC address: 08:17:f4:33:9d:00 IP (If 20) address: 10.10.20.2
 Management Port MAC Address: 08:17:f4:33:9d:fe
 Management Port IP Address (if 128): 172.25.101.243
 Hardware Revision: 0
 Hardware Part No: BAC-00065-00
 Switch Serial No: US71120007
 Manufacturing date: 11/13

Software Version 7.4.1.0 (FLASH image1), active configuration.

Temperature Mother Top: 26 C
 Temperature Mother Bottom: 32 C
 Temperature Daughter Top: 26 C
 Temperature Daughter Bottom: 30 C

Warning at 75 C and Recover at 90 C

Fan 1 in Module 1: RPM= 8450 PWM= 15(5%) Front-To-Back
 Fan 2 in Module 1: RPM= 3967 PWM= 15(5%) Front-To-Back
 Fan 3 in Module 2: RPM= 8667 PWM= 15(5%) Front-To-Back
 Fan 4 in Module 2: RPM= 4094 PWM= 15(5%) Front-To-Back
 Fan 5 in Module 3: RPM= 7883 PWM= 15(5%) Front-To-Back
 Fan 6 in Module 3: RPM= 4173 PWM= 15(5%) Front-To-Back
 Fan 7 in Module 4: RPM= 8837 PWM= 15(5%) Front-To-Back
 Fan 8 in Module 4: RPM= 3994 PWM= 15(5%) Front-To-Back

System Fan Airflow: Front-To-Back

Power Supply 1: OK
 Power Supply 2: OK

Power Faults: ()
 Fan Faults: ()
 Service Faults: ()

Show vlan on G8264tor_1

Example 5-186 shows VLAN assignments for all of the ports on G8264tor_1.

Example 5-186 G8264tor_1 show vlan output

VLAN	Name	Status	Ports
1	Default VLAN	ena	17 19 21 23 25-63
20	VLAN 20	ena	18 20
21	VLAN 21	ena	22 24
4000	ISL h1thchk	ena	64
4092	DATA	ena	1-16 25-28 37-40
4094	ISL	ena	1-16
4095	Mgmt VLAN	ena	MGT

Show vlan on G8264tor_2

Example 5-187 shows VLAN assignments for all of the ports on G8264tor_2.

Example 5-187 G8264tor_2 show vlan output

VLAN	Name	Status	Ports
1	Default VLAN	ena	17 19 21 23 25-63
22	VLAN 22	ena	18 20
23	VLAN 23	ena	22 24
4000	ISL hlthchk	ena	64
4092	DATA	ena	1-16 25-28 37-40
4094	ISL	ena	1-16
4095	Mgmt VLAN	ena	MGT

Show interface status

Because there is only one compute node in the chassis (in slot 1), all the other internal ports are listed as “down” from a link perspective in the output in Example 5-188.

Example 5-188 G8264tor_1 show interface status output

Alias	Port	Speed	Duplex	Flow Ctrl	Link	Name
				--TX-----RX--		
1	1	10000	full	no no	up	ISL
2	2	10000	full	no no	up	ISL
3	3	10000	full	no no	up	ISL
4	4	10000	full	no no	up	ISL
5	5	10000	full	no no	up	ISL
6	6	10000	full	no no	up	ISL
7	7	10000	full	no no	up	ISL
8	8	10000	full	no no	up	ISL
9	9	10000	full	no no	up	ISL
10	10	10000	full	no no	up	ISL
11	11	10000	full	no no	up	ISL
12	12	10000	full	no no	up	ISL
13	13	10000	full	no no	up	ISL
14	14	10000	full	no no	up	ISL
15	15	10000	full	no no	up	ISL
16	16	10000	full	no no	up	ISL
17	17	1G/10G	full	no no	down	17
18	18	10000	full	no no	up	Po5 to
Nexus5548core_1						
19	19	1G/10G	full	no no	down	19
20	20	10000	full	no no	up	Po5 to
Nexus5548core_1						
21	21	1G/10G	full	no no	down	21
22	22	10000	full	no no	up	Po6 to
Nexus5548core_2						
23	23	1G/10G	full	no no	down	23
24	24	10000	full	no no	up	Po6 to
Nexus5548core_2						
25	25	10000	full	no no	up	Link to EN4093-1
26	26	10000	full	no no	up	Link to EN4093-1
27	27	10000	full	no no	up	Link to EN4093-1
28	28	10000	full	no no	up	Link to EN4093-1

29	29	1G/10G	full	no	no	down	29
30	30	1G/10G	full	no	no	down	30
31	31	1G/10G	full	no	no	down	31
32	32	1G/10G	full	no	no	down	32
33	33	1G/10G	full	no	no	down	33
34	34	1G/10G	full	no	no	down	34
35	35	1G/10G	full	no	no	down	35
36	36	1G/10G	full	no	no	down	36
37	37	10000	full	no	no	up	Link to EN4093-2
38	38	10000	full	no	no	up	Link to EN4093-2
39	39	10000	full	no	no	up	Link to EN4093-2
40	40	10000	full	no	no	up	Link to EN4093-2
41	41	1G/10G	full	no	no	down	41
42	42	1G/10G	full	no	no	down	42
43	43	1G/10G	full	no	no	down	43
44	44	1G/10G	full	no	no	down	44
45	45	1G/10G	full	no	no	down	45
46	46	1G/10G	full	no	no	down	46
47	47	1G/10G	full	no	no	down	47
48	48	1G/10G	full	no	no	down	48
49	49	1G/10G	full	no	no	down	49
50	50	1G/10G	full	no	no	down	50
51	51	1G/10G	full	no	no	down	51
52	52	1G/10G	full	no	no	down	52
53	53	1G/10G	full	no	no	down	53
54	54	1G/10G	full	no	no	down	54
55	55	1G/10G	full	no	no	down	55
56	56	1G/10G	full	no	no	down	56
57	57	1G/10G	full	no	no	down	57
58	58	1G/10G	full	no	no	down	58
59	59	1G/10G	full	no	no	down	59
60	60	1G/10G	full	no	no	down	60
61	61	1G/10G	full	no	no	down	61
62	62	1G/10G	full	no	no	down	62
63	63	1G/10G	full	no	no	down	63
64	64	10000	full	no	no	up	ISL h1thchk
MGT	65	1000	full	yes	yes	up	MGT

Show lldp remote-device on G8264tor_1

The command output in Example 5-189 shows the physical topology and verifies that cables are plugged into the ports specified in both the Network Topology diagram, and the configuration specified in the appendix.

Example 5-189 G8264tor_1 show lldp remote-device output

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
1	1	08 17 f4 33 75 00	1	G8264TOR-2
2	2	08 17 f4 33 75 00	2	G8264TOR-2
3	3	08 17 f4 33 75 00	3	G8264TOR-2
4	4	08 17 f4 33 75 00	4	G8264TOR-2
5	5	08 17 f4 33 75 00	5	G8264TOR-2
6	7	08 17 f4 33 75 00	6	G8264TOR-2

7	8	08 17 f4 33 75 00	7	G8264TOR-2
8	9	08 17 f4 33 75 00	8	G8264TOR-2
9	10	08 17 f4 33 75 00	9	G8264TOR-2
10	11	08 17 f4 33 75 00	10	G8264TOR-2
11	12	08 17 f4 33 75 00	11	G8264TOR-2
12	13	08 17 f4 33 75 00	12	G8264TOR-2
13	14	08 17 f4 33 75 00	13	G8264TOR-2
14	15	08 17 f4 33 75 00	14	G8264TOR-2
18	16	54 7f ee 2d 36 0e	Eth1/7	Nexus5548core_1
15	17	08 17 f4 33 75 00	15	G8264TOR-2
20	18	54 7f ee 2d 36 0f	Eth1/8	Nexus5548core_1
16	19	08 17 f4 33 75 00	16	G8264TOR-2
25	20	6c ae 8b bf 6d 00	58	en4093flex_1
26	21	6c ae 8b bf 6d 00	57	en4093flex_1
27	22	6c ae 8b bf 6d 00	60	en4093flex_1
28	23	6c ae 8b bf 6d 00	59	en4093flex_1
37	24	6c ae 8b bf fe 00	57	en4093flex_2
38	25	6c ae 8b bf fe 00	59	en4093flex_2
39	26	6c ae 8b bf fe 00	58	en4093flex_2
40	27	6c ae 8b bf fe 00	60	en4093flex_2
64	28	08 17 f4 33 75 00	64	G8264TOR-2
24	30	54 7f ee 72 bd 11	Eth1/10	Nexus5548core_2
22	31	54 7f ee 72 bd 10	Eth1/9	Nexus5548core_2

Show lldp remote-device on G8264tor_2

Example 5-190 command output shows LLDP-related information for the second G8264 switch.

Example 5-190 G8264tor_2 show lldp remote-device output

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
1	1	08 17 f4 33 9d 00	1	G8264TOR-1
2	2	08 17 f4 33 9d 00	2	G8264TOR-1
3	3	08 17 f4 33 9d 00	3	G8264TOR-1
4	4	08 17 f4 33 9d 00	4	G8264TOR-1
5	5	08 17 f4 33 9d 00	5	G8264TOR-1
6	6	08 17 f4 33 9d 00	6	G8264TOR-1
7	7	08 17 f4 33 9d 00	7	G8264TOR-1
8	9	08 17 f4 33 9d 00	8	G8264TOR-1
9	10	08 17 f4 33 9d 00	9	G8264TOR-1
10	11	08 17 f4 33 9d 00	10	G8264TOR-1
11	12	08 17 f4 33 9d 00	11	G8264TOR-1
12	13	08 17 f4 33 9d 00	12	G8264TOR-1
13	14	08 17 f4 33 9d 00	13	G8264TOR-1
14	15	08 17 f4 33 9d 00	14	G8264TOR-1
15	16	08 17 f4 33 9d 00	15	G8264TOR-1
16	17	08 17 f4 33 9d 00	16	G8264TOR-1
22	18	54 7f ee 2d 36 10	Eth1/9	Nexus5548core_1
24	19	54 7f ee 2d 36 11	Eth1/10	Nexus5548core_1
25	20	6c ae 8b bf 6d 00	63	en4093flex_1
26	21	6c ae 8b bf 6d 00	61	en4093flex_1
27	22	6c ae 8b bf 6d 00	64	en4093flex_1

28	23	6c ae 8b bf 6d 00	62	en4093flex_1
37	24	6c ae 8b bf fe 00	61	en4093flex_2
38	25	6c ae 8b bf fe 00	63	en4093flex_2
64	26	08 17 f4 33 9d 00	64	G8264TOR-1
39	27	6c ae 8b bf fe 00	62	en4093flex_2
40	28	6c ae 8b bf fe 00	64	en4093flex_2
18	29	54 7f ee 72 bd 0e	Eth1/7	Nexus5548core_2
20	30	54 7f ee 72 bd 0f	Eth1/8	Nexus5548core_2

Show vlag isl

The command output in Example 5-191 shows the status of the ISL between the G8264 switches, and the ports that comprise the ISL itself.

Example 5-191 G8264tor_1 show vlag isl output

ISL_ID	ISL_Vlan	ISL_Trunk	ISL_Members	Link_State	Trunk_State
67	4094	Adminkey 1000	1	UP	UP
			2	UP	UP
			3	UP	UP
			4	UP	UP
			5	UP	UP
			6	UP	UP
			7	UP	UP
			8	UP	UP
			9	UP	UP
			10	UP	UP
			11	UP	UP
			12	UP	UP
			13	UP	UP
			14	UP	UP
			15	UP	UP
			16	UP	UP

Show vlag information

The output in Example 5-192 shows that the downstream vLAG between the G8264 and EN4093/R switches is up and operational as referenced by the LACP admin key of 2002. The ISL between the G8264 switches is up too.

G8264tor_1 is acting as the admin and operational role of SECONDARY. For centralized vLAG functions, such as vLAG STP, one of the vLAG switches must control the protocol operations. To select the switch that controls the centralized vLAG function, perform role election. The switch with the primary role controls the centralized operation. Role election is non-preemptive. That is, if a primary already exists, another switch that is coming up remains as secondary even if it can become primary based on the role election logic.

Role election is determined by comparing the local vLAG system priority and local system MAC address. The switch with the smaller priority value becomes the vLAG primary switch. If the priorities are the same, the switch with the smaller system MAC address becomes the vLAG primary switch. You can configure vLAG priority to anything between <0-65535>. The priority was left at the default value of 0 in all examples.

Example 5-192 G8264tor_1 show vlag information output

```
vLAG Tier ID: 2
vLAG system MAC: 08:17:f4:c3:dd:01
```

```

Local MAC 08:17:f4:33:9d:00 Priority 0 Admin Role SECONDARY (Operational Role
SECONDARY)
Peer MAC 08:17:f4:33:75:00 Priority 0
Health local 1.1.1.1 peer 1.1.1.2 State UP
ISL trunk id 67
ISL state Up
Startup Delay Interval: 120s (Finished)

```

vLAG 66: config with admin key 2002, associated trunk 66, state formed

Show vlag adminkey 2002

The output in Example 5-193 shows that the downstream vLAG towards the EN4093/R switches is formed and enabled using LACP reference key 2002.

Example 5-193 G8264tor_1 show vlag adminkey 2002 output

```

vLAG is enabled on admin key 2002
Current LACP params for 25: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 26: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 27: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 28: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 37: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 38: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 39: active, Priority 32768, Admin Key 2002, Min-Links 1

Current LACP params for 40: active, Priority 32768, Admin Key 2002, Min-Links 1

```

Show lacp information state up

Example 5-194 shows which ports are participating in an LACP aggregation, and which reference keys are used on those specific interfaces.

Example 5-194 G8264tor_1 show lacp information state up

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
1	active	1000	1000	yes	32768	1	67	up	1
2	active	1000	1000	yes	32768	1	67	up	1
3	active	1000	1000	yes	32768	1	67	up	1
4	active	1000	1000	yes	32768	1	67	up	1
5	active	1000	1000	yes	32768	1	67	up	1
6	active	1000	1000	yes	32768	1	67	up	1
7	active	1000	1000	yes	32768	1	67	up	1
8	active	1000	1000	yes	32768	1	67	up	1
9	active	1000	1000	yes	32768	1	67	up	1
10	active	1000	1000	yes	32768	1	67	up	1
11	active	1000	1000	yes	32768	1	67	up	1
12	active	1000	1000	yes	32768	1	67	up	1
13	active	1000	1000	yes	32768	1	67	up	1
14	active	1000	1000	yes	32768	1	67	up	1

15	active	1000	1000	yes	32768	1	67	up	1
16	active	1000	1000	yes	32768	1	67	up	1
25	active	2002	2002	yes	32768	26	66	up	1
26	active	2002	2002	yes	32768	26	66	up	1
27	active	2002	2002	yes	32768	26	66	up	1
28	active	2002	2002	yes	32768	26	66	up	1
37	active	2002	2002	yes	32768	26	66	up	1
38	active	2002	2002	yes	32768	26	66	up	1
39	active	2002	2002	yes	32768	26	66	up	1
40	active	2002	2002	yes	32768	26	66	up	1

Show ip ospf neighbor on G8264tor_1

Example 5-195 lists output from the **show ip ospf neighbor** command, showing that OSPF is enabled and displaying associated neighbor information. Use this information to verify the Network Topology diagram.

Example 5-195 G8264tor_1 show ip ospf neighbor output

Intf	NeighborID	Prio	State	Address
----	-----	----	-----	-----
20	10.10.11.249	1	Full	10.10.20.1
21	10.10.11.200	1	Full	10.10.21.1
92	10.10.11.244	1	Full	10.1.4.244

Show ip ospf neighbor on G8264tor_2

Example 5-196 lists output from the **show ip ospf neighbor** command on the second G8264 switch.

Example 5-196 G8264tor_2 show ip ospf neighbor output

Intf	NeighborID	Prio	State	Address
----	-----	----	-----	-----
22	10.10.11.249	1	Full	10.10.22.1
23	10.10.11.200	1	Full	10.10.23.1
92	10.10.11.243	1	Full	10.1.4.243

Show ip ospf routes on G8264tor_1

Example 5-197 lists output from the **show ip ospf routes** command, showing learned routes identified by using the neighboring interfaces.

Example 5-197 G8264tor_1 show ip ospf route output

Codes: IA - OSPF inter area,
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
* - best

```

10.1.4.0/24 directly connected
10.10.20.0/30 directly connected
10.10.21.0/30 directly connected
* 10.10.22.0/30 via 10.1.4.244
* 10.10.23.0/30 via 10.1.4.244
* 10.10.30.0/24 via 10.10.20.1
* 10.10.30.0/24 via 10.10.21.1

```

Show ip ospf routes on G8264tor_2

Example 5-198 lists output from the **show ip ospf routes** command on the second G8264 switch.

Example 5-198 G8264tor_2 show ip ospf route output

```
Codes: IA - OSPF inter area,
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       * - best

10.1.4.0/24 directly connected
10.10.22.0/30 directly connected
10.10.23.0/30 directly connected
* 10.10.20.0/30 via 10.1.4.243
* 10.10.21.0/30 via 10.1.4.243
* 10.10.30.0/24 via 10.10.22.1
* 10.10.30.0/24 via 10.10.23.1
```

Show ip ospf interface on G8264tor_1

OSPF interface-related information is displayed in Example 5-199.

Example 5-199 G8264tor_1 show ip ospf interface output

```
Ip Address 10.10.20.2, Area 0.0.0.0, Admin Status UP
Router ID 10.10.11.243, State BackupDR, Priority 1
Designated Router (ID) 10.10.11.249, Ip Address 10.10.20.1
Backup Designated Router (ID) 10.10.11.243, Ip Address 10.10.20.2
Timer intervals, Hello 10, Dead 40, Wait 40, Retransmit 5, Transit delay 1
Neighbor count is 1   If Events 5, Authentication type none

-----

Ip Address 10.10.21.2, Area 0.0.0.0, Admin Status UP
Router ID 10.10.11.243, State DR, Priority 1
Designated Router (ID) 10.10.11.243, Ip Address 10.10.21.2
Backup Designated Router (ID) 10.10.11.200, Ip Address 10.10.21.1
Timer intervals, Hello 10, Dead 40, Wait 40, Retransmit 5, Transit delay 1
Neighbor count is 1   If Events 2, Authentication type none

-----

Ip Address 10.1.4.243, Area 0.0.0.0, Admin Status UP
Router ID 10.10.11.243, State DR, Priority 1
Designated Router (ID) 10.10.11.243, Ip Address 10.1.4.243
Backup Designated Router (ID) 10.10.11.244, Ip Address 10.1.4.244
Timer intervals, Hello 10, Dead 40, Wait 40, Retransmit 5, Transit delay 1
Neighbor count is 1   If Events 6, Authentication type none

-----
```

Show ip ospf interface for G8264tor_2

OSPF interface-related information is displayed in Example 5-200 for the second G8264 switch.

Example 5-200 G8264tor_2 show ip ospf interface output

```
Ip Address 10.10.22.2, Area 0.0.0.0, Admin Status UP
  Router ID 10.10.11.244, State BackupDR, Priority 1
  Designated Router (ID) 10.10.11.249, Ip Address 10.10.22.1
  Backup Designated Router (ID) 10.10.11.244, Ip Address 10.10.22.2
  Timer intervals, Hello 10, Dead 40, Wait 40, Retransmit 5, Transit delay 1
  Neighbor count is 1   If Events 5, Authentication type none
```

```
-----
Ip Address 10.10.23.2, Area 0.0.0.0, Admin Status UP
  Router ID 10.10.11.244, State DR, Priority 1
  Designated Router (ID) 10.10.11.244, Ip Address 10.10.23.2
  Backup Designated Router (ID) 10.10.11.200, Ip Address 10.10.23.1
  Timer intervals, Hello 10, Dead 40, Wait 40, Retransmit 5, Transit delay 1
  Neighbor count is 1   If Events 2, Authentication type none
```

```
-----
Ip Address 10.1.4.244, Area 0.0.0.0, Admin Status UP
  Router ID 10.10.11.244, State BackupDR, Priority 1
  Designated Router (ID) 10.10.11.243, Ip Address 10.1.4.243
  Backup Designated Router (ID) 10.10.11.244, Ip Address 10.1.4.244
  Timer intervals, Hello 10, Dead 40, Wait 40, Retransmit 5, Transit delay 1
  Neighbor count is 1   If Events 6, Authentication type none
```

Show ip vrrp information for G8264tor_1

The VRRP information in Example 5-201 confirms that G8264tor_1 is the master for Virtual IP (VIP) address 10.1.4.241.

Example 5-201 G8264tor_1 show ip vrrp information output

```
VRRP information:
  1: vrid 1, 10.1.4.241,      if 92, renter, prio 120, master
```

Show ip vrrp information for G8264tor_2

The VRRP information in Example 5-202 confirms that G8264tor-2 is the backup.

Example 5-202 G8264tor_2 show ip vrrp information output

```
VRRP information:
  1: vrid 1, 10.1.4.241,      if 92, renter, prio 110, backup
```

Ping output for equipment on VLAN 4092

To verify connectivity, issue **ping** commands to devices on VLAN 4092 (Data VLAN) in Example 5-203. IP address 10.4.1.10 represents a compute node with an operating system installed, *flex_node1* on the Network Topology diagram. IP address 10.10.30.1 represents the VIP on the Nexus pair simulating the Server network.

Example 5-203 Ping verification for equipment on VLAN 4092

```
G8264TOR-1#ping 10.10.30.1 data-port
Connecting via DATA port.
[host 10.10.30.1, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.10.30.1: #1 ok, RTT 1 msec.
10.10.30.1: #2 ok, RTT 0 msec.
10.10.30.1: #3 ok, RTT 0 msec.
10.10.30.1: #4 ok, RTT 0 msec.
10.10.30.1: #5 ok, RTT 1 msec.
Ping finished.
G8264TOR-1#
G8264TOR-1#ping 10.1.4.238 data-port
Connecting via DATA port.
[host 10.1.4.238, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.238: #1 ok, RTT 7 msec.
10.1.4.238: #2 ok, RTT 3 msec.
10.1.4.238: #3 ok, RTT 2 msec.
10.1.4.238: #4 ok, RTT 1 msec.
10.1.4.238: #5 ok, RTT 0 msec.
Ping finished.
G8264TOR-1#
G8264TOR-1#ping 10.1.4.239 data-port
Connecting via DATA port.
[host 10.1.4.239, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl
255, tos 0]
10.1.4.239: #1 ok, RTT 5 msec.
10.1.4.239: #2 ok, RTT 0 msec.
10.1.4.239: #3 ok, RTT 13 msec.
10.1.4.239: #4 ok, RTT 0 msec.
10.1.4.239: #5 ok, RTT 0 msec.
Ping finished.
G8264TOR-1#
G8264TOR-1#ping 10.1.4.10 data-port
Connecting via DATA port.
[host 10.1.4.10, max tries 5, delay 1000 msec, length 0, ping source N/S, ttl 255,
tos 0]
10.1.4.10: #1 ok, RTT 2 msec.
10.1.4.10: #2 ok, RTT 0 msec.
10.1.4.10: #3 ok, RTT 0 msec.
10.1.4.10: #4 ok, RTT 0 msec.
10.1.4.10: #5 ok, RTT 0 msec.
Ping finished.
```

Nexus output

This section lists output from the switch with hostname Nexus5548core_1. Similar or identical output exists for the switch with hostname Nexus5548core_2 unless otherwise noted.

Show version

Example 5-204 shows information about the switch and the associated code/firmware level.

Example 5-204 Nexus5548core_1 show version output

```
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Documents: http://www.cisco.com/en/US/products/ps9372/tsd_products_support_series_home.html
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software are covered under the GNU Public
License. A copy of the license is available at
http://www.gnu.org/licenses/gpl.html.
```

Software

```
BIOS:          version 3.5.0
loader:        version N/A
kickstart:     version 5.2(1)N1(1b)
system:        version 5.2(1)N1(1b)
power-seq:     Module 1: version v1.0
               Module 3: version v5.0
uC:            version v1.2.0.1
SFP uC:        Module 1: v1.0.0.0
BIOS compile time:      02/03/2011
kickstart image file is: bootflash:///n5000-uk9-kickstart.5.2.1.N1.1b.bin
kickstart compile time: 9/17/2012 11:00:00 [09/17/2012 18:38:53]
system image file is:   bootflash:///n5000-uk9.5.2.1.N1.1b.bin
system compile time:    9/17/2012 11:00:00 [09/17/2012 20:38:22]
```

Hardware

```
cisco Nexus5548 Chassis ("02 32X10GE/Modular Universal Platform Supervisor")
Intel(R) Xeon(R) CPU          with 8263848 kB of memory.
Processor Board ID F0C15424504
```

```
Device name: Nexus5548core_1
bootflash:   2007040 kB
```

Kernel uptime is 0 day(s), 22 hour(s), 32 minute(s), 3 second(s)

Last reset

```
Reason: Unknown
System version: 5.2(1)N1(1b)
Service:
```

plugin

```
Core Plugin, Ethernet Plugin
```

Show vlan brief

Example 5-205 displays the VLAN assignments for all of the ports on the switch.

Example 5-205 Nexus5548core_1 show vlan brief output

VLAN	Name	Status	Ports
1	default	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/11, Eth1/12 Eth1/13, Eth1/14, Eth1/15 Eth1/16, Eth1/18, Eth1/20 Eth1/21, Eth1/22, Eth1/23 Eth1/24, Eth1/25, Eth1/26 Eth1/27, Eth1/28, Eth1/29 Eth1/30, Eth1/31, Eth1/32
30	Server	active	Po100, Eth1/17, Eth1/19

Show interface status

Example 5-206 shows the full interface table, listing port status, speed, and so on, for the Nexus5548core_1 switch.

Example 5-206 Nexus5548core_1 show interface status output

Port	Name	Status	Vlan	Duplex	Speed	Type
Eth1/1	--	sfpAbsent	1	full	10G	--
Eth1/2	--	sfpAbsent	1	full	10G	--
Eth1/3	--	sfpAbsent	1	full	10G	--
Eth1/4	--	sfpAbsent	1	full	10G	--
Eth1/5	--	sfpAbsent	1	full	10G	--
Eth1/6	--	sfpAbsent	1	full	10G	--
Eth1/7	Po5 to G8264tor_1	connected	routed	full	10G	10Gbase-(un
Eth1/8	Po5 to G8264tor_1	connected	routed	full	10G	10Gbase-(un
Eth1/9	Po6 to G8264tor_2	connected	routed	full	10G	10Gbase-(un
Eth1/10	Po6 to G8264tor_2	connected	routed	full	10G	10Gbase-(un
Eth1/11	--	sfpAbsent	1	full	10G	--
Eth1/12	--	sfpAbsent	1	full	10G	--
Eth1/13	--	sfpAbsent	1	full	10G	--
Eth1/14	--	sfpAbsent	1	full	10G	--
Eth1/15	--	sfpAbsent	1	full	10G	--
Eth1/16	--	sfpAbsent	1	full	10G	--
Eth1/17	Po100 to Nexus5548	connected	trunk	full	10G	10Gbase-(un
Eth1/18	--	sfpAbsent	1	full	10G	--
Eth1/19	Po100 to Nexus5548	connected	trunk	full	10G	10Gbase-(un
Eth1/20	--	sfpAbsent	1	full	10G	--
Eth1/21	--	disabled	1	full	10G	10Gbase-(un
Eth1/22	--	sfpAbsent	1	full	10G	--
Eth1/23	--	sfpAbsent	1	full	10G	--
Eth1/24	--	sfpAbsent	1	full	10G	--
Eth1/25	--	sfpAbsent	1	full	10G	--
Eth1/26	--	sfpAbsent	1	full	10G	--
Eth1/27	--	sfpAbsent	1	full	10G	--
Eth1/28	--	sfpAbsent	1	full	10G	--
Eth1/29	--	sfpAbsent	1	full	10G	--
Eth1/30	--	sfpAbsent	1	full	10G	--

Eth1/31	--	sfpAbsent 1	full	10G	--
Eth1/32	--	sfpAbsent 1	full	10G	--
Po5	--	connected routed	full	10G	--
Po6	--	connected routed	full	10G	--
Po100	Switch-to-Switch L	connected trunk	full	10G	--
mgmt0	--	connected routed	full	1000	--
Lo1	OSPF router-id	connected routed	auto	auto	--

Show lldp neighbors on Nexus5548core_1

Example 5-207 lists the LLDP information and verifies physical connectivity.

Example 5-207 Nexus5548core_1 show LLDP neighbors output

Capability codes:

(R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device

(W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

Device ID	Local Intf	Hold-time	Capability	Port ID
G8264TOR-1	Eth1/7	120	BR	18
G8264TOR-1	Eth1/8	120	BR	20
G8264TOR-2	Eth1/9	120	BR	22
G8264TOR-2	Eth1/10	120	BR	24
Nexus5548core_2	Eth1/17	120	B	Eth1/17
Nexus5548core_2	Eth1/19	120	B	Eth1/19

Total entries displayed: 6

Show lldp neighbors on Nexus5548core_2

Example 5-208 lists the LLDP information and verifies physical connectivity.

Example 5-208 Nexus5548core_2 show LLDP neighbors output

Capability codes:

(R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device

(W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

Device ID	Local Intf	Hold-time	Capability	Port ID
G8264TOR-2	Eth1/7	120	BR	18
G8264TOR-2	Eth1/8	120	BR	20
G8264TOR-1	Eth1/9	120	BR	22
G8264TOR-1	Eth1/10	120	BR	24
Nexus5548core_1	Eth1/17	120	B	Eth1/17
Nexus5548core_1	Eth1/19	120	B	Eth1/19

Total entries displayed: 6

Show ip ospf interface on Nexus5548core_1

Example 5-209 shows the OSPF interface output on the Nexus5548core_1 switch.

Example 5-209 Nexus5548core_1 show ip ospf interface output

Vlan30 is up, line protocol is up

IP address 10.10.30.2/24, Process ID 100 VRF default, area 0.0.0.0

Enabled by interface configuration

State DR, Network type BROADCAST, cost 100

Index 3, Transmit delay 1 sec, Router Priority 1

Designated Router ID: 10.10.11.249, address: 10.10.30.2

Backup Designated Router ID: 10.10.11.200, address: 10.10.30.3

1 Neighbors, flooding to 1, adjacent with 1

```

Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello timer due in 00:00:01
No authentication
Number of opaque link LSAs: 0, checksum sum 0
port-channel5 is up, line protocol is up
  IP address 10.10.20.1/30, Process ID 100 VRF default, area 0.0.0.0
  Enabled by interface configuration
  State DR, Network type BROADCAST, cost 5
  Index 1, Transmit delay 1 sec, Router Priority 1
  Designated Router ID: 10.10.11.249, address: 10.10.20.1
  Backup Designated Router ID: 10.10.11.243, address: 10.10.20.2
  1 Neighbors, flooding to 1, adjacent with 1
  Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello timer due in 00:00:04
  No authentication
  Number of opaque link LSAs: 0, checksum sum 0
port-channel6 is up, line protocol is up
  IP address 10.10.22.1/30, Process ID 100 VRF default, area 0.0.0.0
  Enabled by interface configuration
  State DR, Network type BROADCAST, cost 5
  Index 2, Transmit delay 1 sec, Router Priority 1
  Designated Router ID: 10.10.11.249, address: 10.10.22.1
  Backup Designated Router ID: 10.10.11.244, address: 10.10.22.2
  1 Neighbors, flooding to 1, adjacent with 1
  Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello timer due in 00:00:02
  No authentication
  Number of opaque link LSAs: 0, checksum sum 0

```

Show ip ospf interface on Nexus5548core_2

Example 5-210 shows the OSPF interface output on the Nexus5548core_2 switch.

Example 5-210 Nexus5548core_2 show ip ospf interface output

```

port-channel5 is up, line protocol is up
  IP address 10.10.23.1/30, Process ID 100 VRF default, area 0.0.0.0
  Enabled by interface configuration
  State BDR, Network type BROADCAST, cost 5
  Index 1, Transmit delay 1 sec, Router Priority 1
  Designated Router ID: 10.10.11.244, address: 10.10.23.2
  Backup Designated Router ID: 10.10.11.200, address: 10.10.23.1
  1 Neighbors, flooding to 1, adjacent with 1
  Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello timer due in 00:00:03
  No authentication
  Number of opaque link LSAs: 0, checksum sum 0
port-channel6 is up, line protocol is up
  IP address 10.10.21.1/30, Process ID 100 VRF default, area 0.0.0.0
  Enabled by interface configuration
  State BDR, Network type BROADCAST, cost 5
  Index 2, Transmit delay 1 sec, Router Priority 1
  Designated Router ID: 10.10.11.243, address: 10.10.21.2
  Backup Designated Router ID: 10.10.11.200, address: 10.10.21.1
  1 Neighbors, flooding to 1, adjacent with 1
  Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5

```

```

    Hello timer due in 00:00:02
    No authentication
    Number of opaque link LSAs: 0, checksum sum 0
Vlan30 is up, line protocol is up
    IP address 10.10.30.3/24, Process ID 100 VRF default, area 0.0.0.0
    Enabled by interface configuration
    State BDR, Network type BROADCAST, cost 100
    Index 3, Transmit delay 1 sec, Router Priority 1
    Designated Router ID: 10.10.11.249, address: 10.10.30.2
    Backup Designated Router ID: 10.10.11.200, address: 10.10.30.3
    1 Neighbors, flooding to 1, adjacent with 1
    Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello timer due in 00:00:04
    No authentication
    Number of opaque link LSAs: 0, checksum sum 0

```

Show ip ospf neighbor on Nexus5548core_1

Example 5-211 displays the OSPF neighbor data from the perspective of the Nexus5548core_1 switch.

Example 5-211 Nexus5548core_1 show ip ospf neighbor output

```

OSPF Process ID 100 VRF default
Total number of neighbors: 3
Neighbor ID      Pri State           Up Time  Address      Interface
10.10.11.200     1 FULL/BDR       00:06:16 10.10.30.3   Vlan30
10.10.11.243     1 FULL/BDR       02:36:17 10.10.20.2   Po5
10.10.11.244     1 FULL/BDR       02:34:32 10.10.22.2   Po6

```

Show ip ospf neighbor on Nexus5548core_2

Example 5-212 displays the OSPF neighbor data from the perspective of the Nexus5548core_2 switch.

Example 5-212 Nexus5548core_2 show ip ospf neighbor output

```

OSPF Process ID 100 VRF default
Total number of neighbors: 3
Neighbor ID      Pri State           Up Time  Address      Interface
10.10.11.244     1 FULL/DR        01:43:06 10.10.23.2   Po5
10.10.11.243     1 FULL/DR        01:42:14 10.10.21.2   Po6
10.10.11.249     1 FULL/DR        00:06:19 10.10.30.2   Vlan30

```

Show ip route ospf for Nexus5548core_1

Example 5-213 lists routes that were learned by using OSPF for Nexus5548core_1.

Example 5-213 Nexus5548core_1 show ip route ospf output

```

IP Route Table for VRF "default"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]

10.1.4.0/24, ubest/mbest: 2/0
  *via 10.10.20.2, Po5, [110/6], 02:37:43, ospf-100, intra
  *via 10.10.22.2, Po6, [110/6], 02:36:08, ospf-100, intra

```

```
10.10.21.0/30, ubest/mbest: 1/0
    *via 10.10.20.2, Po5, [110/6], 02:16:35, ospf-100, intra
10.10.23.0/30, ubest/mbest: 1/0
    *via 10.10.22.2, Po6, [110/6], 02:16:35, ospf-100, intra
```

Show ip route ospf for Nexus5548core_2

Example 5-214 lists routes that were learned by using OSPF for Nexus5548core_2.

Example 5-214 Nexus5548core_2 show ip route ospf output

```
IP Route Table for VRF "default"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]

10.1.4.0/24, ubest/mbest: 2/0
    *via 10.10.21.2, Po6, [110/6], 01:43:43, ospf-100, intra
    *via 10.10.23.2, Po5, [110/6], 01:44:36, ospf-100, intra
10.10.20.0/30, ubest/mbest: 1/0
    *via 10.10.21.2, Po6, [110/6], 01:43:43, ospf-100, intra
10.10.22.0/30, ubest/mbest: 1/0
    *via 10.10.23.2, Po5, [110/6], 01:44:36, ospf-100, intra
```

Show vrrp detail for Nexus5548core_1

To simulate the Server VLAN, output for the running VRRP process on Nexus5548core_1 is listed in Example 5-215.

Example 5-215 Nexus5548core_1 show vrrp detail output

```
Vlan30 - Group 1 (IPv4)
  State is Master
  Virtual IP address is 10.10.30.1
  Priority 200, Configured 200
  Forwarding threshold(for VPC), lower: 1 upper: 200
  Advertisement interval 1
  Preemption enabled
  Virtual MAC address is 0000.5e00.0101
  Master router is Local
```

Show vrrp detail for Nexus5548core_2

Example 5-216 shows the output for the running VRRP process on Nexus5548core_2. Notice that the virtual MAC address is the same across both switches.

Example 5-216 Nexus5548core_2 show vrrp detail output

```
Vlan30 - Group 1 (IPv4)
  State is Backup
  Virtual IP address is 10.10.30.1
  Priority 150, Configured 150
  Forwarding threshold(for VPC), lower: 1 upper: 150
  Advertisement interval 1
  Preemption enabled
  Virtual MAC address is 0000.5e00.0101
  Master router is 10.10.30.2
```

Ping output for equipment on VLAN 4092

To verify connectivity, issue **ping** commands to devices on VLAN 4092 (Data VLAN) as shown in Example 5-217. Included is the compute node with an assigned IP address of 10.1.4.10.

Example 5-217 Ping verification for equipment on VLAN 4092

```
Nexus5548core_1# ping 10.1.4.243
PING 10.1.4.243 (10.1.4.243): 56 data bytes
64 bytes from 10.1.4.243: icmp_seq=0 ttl=253 time=0.786 ms
64 bytes from 10.1.4.243: icmp_seq=1 ttl=253 time=0.512 ms
64 bytes from 10.1.4.243: icmp_seq=2 ttl=253 time=1.375 ms
64 bytes from 10.1.4.243: icmp_seq=3 ttl=253 time=1.196 ms
64 bytes from 10.1.4.243: icmp_seq=4 ttl=253 time=7.616 ms

--- 10.1.4.243 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.512/2.296/7.616 ms
```

```
Nexus5548core_1# ping 10.1.4.244
PING 10.1.4.244 (10.1.4.244): 56 data bytes
64 bytes from 10.1.4.244: icmp_seq=0 ttl=254 time=52.424 ms
64 bytes from 10.1.4.244: icmp_seq=1 ttl=254 time=0.723 ms
64 bytes from 10.1.4.244: icmp_seq=2 ttl=254 time=0.608 ms
64 bytes from 10.1.4.244: icmp_seq=3 ttl=254 time=7.32 ms
64 bytes from 10.1.4.244: icmp_seq=4 ttl=254 time=9.398 ms

--- 10.1.4.244 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.608/14.094/52.424 ms
```

```
Nexus5548core_1# ping 10.1.4.238
PING 10.1.4.238 (10.1.4.238): 56 data bytes
64 bytes from 10.1.4.238: icmp_seq=0 ttl=253 time=0.879 ms
64 bytes from 10.1.4.238: icmp_seq=1 ttl=253 time=0.655 ms
64 bytes from 10.1.4.238: icmp_seq=2 ttl=253 time=0.745 ms
64 bytes from 10.1.4.238: icmp_seq=3 ttl=253 time=7.503 ms
64 bytes from 10.1.4.238: icmp_seq=4 ttl=253 time=9.591 ms

--- 10.1.4.238 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.655/3.874/9.591 ms
```

```
Nexus5548core_1# ping 10.1.4.239
PING 10.1.4.239 (10.1.4.239): 56 data bytes
64 bytes from 10.1.4.239: icmp_seq=0 ttl=253 time=1.262 ms
64 bytes from 10.1.4.239: icmp_seq=1 ttl=253 time=0.644 ms
64 bytes from 10.1.4.239: icmp_seq=2 ttl=253 time=1.398 ms
64 bytes from 10.1.4.239: icmp_seq=3 ttl=253 time=4.71 ms
64 bytes from 10.1.4.239: icmp_seq=4 ttl=253 time=9.125 ms

--- 10.1.4.239 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.644/3.427/9.125 ms
```

```

Nexus5548core_1# ping 10.1.4.10
PING 10.1.4.10 (10.1.4.10): 56 data bytes
64 bytes from 10.1.4.10: icmp_seq=0 ttl=61 time=0.892 ms
64 bytes from 10.1.4.10: icmp_seq=1 ttl=61 time=0.659 ms
64 bytes from 10.1.4.10: icmp_seq=2 ttl=61 time=0.744 ms
64 bytes from 10.1.4.10: icmp_seq=3 ttl=61 time=7.473 ms
64 bytes from 10.1.4.10: icmp_seq=4 ttl=61 time=9.592 ms

--- 10.1.4.10 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.659/3.871/9.592 ms

```

5.5.9 Full configuration files

This section displays the configuration on all of the devices in the Network Topology diagram.

EN4093flex-1

Example 5-218 lists the configuration for the EN4093flex-1 switch.

Example 5-218 EN4093-1 switch configuration file

```

version "7.3.1"
switch-type "IBM Flex System Fabric EN4093 10Gb Scalable Switch"
!
!

snmp-server name "en4093flex_1"
!
!
hostname "en4093flex_1"
!
!
interface port INTA1
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port INTB1
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT4
    name "ISL hlthchk"
    pvid 4000
    exit
!
interface port EXT7
    name "ISL"
    tagging

```

```

        pvid 4094
        exit
    !
interface port EXT8
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT9
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT10
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT15
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT16
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT17
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT18
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT19
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit

```



```

!
interface port EXT20
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT21
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT22
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
vlan 1
    member INTA2-INTA14,INTB2-INTB14,EXT1-EXT3,EXT5-EXT6
    no member INTA1,INTB1,EXT4,EXT7-EXT10,EXT15-EXT22
!
vlan 4000
    enable
    name "ISL hlthchk"
    member EXT4
!
vlan 4092
    enable
    name "DATA"
    member INTA1,INTB1,EXT7-EXT10,EXT15-EXT22
!
vlan 4094
    enable
    name "ISL"
    member EXT7-EXT10
!
!
spanning-tree stp 125 vlan 4000
!
spanning-tree stp 126 vlan 4092
!
no spanning-tree stp 127 enable
spanning-tree stp 127 vlan 4094
!
!
interface port EXT7
    lacp mode active
    lacp key 1000
!
interface port EXT8

```

```

        lacp mode active
        lacp key 1000
    !
interface port EXT9
    lacp mode active
    lacp key 1000
    !
interface port EXT10
    lacp mode active
    lacp key 1000
    !
interface port EXT15
    lacp mode active
    lacp key 2000
    !
interface port EXT16
    lacp mode active
    lacp key 2000
    !
interface port EXT17
    lacp mode active
    lacp key 2000
    !
interface port EXT18
    lacp mode active
    lacp key 2000
    !
interface port EXT19
    lacp mode active
    lacp key 2000
    !
interface port EXT20
    lacp mode active
    lacp key 2000
    !
interface port EXT21
    lacp mode active
    lacp key 2000
    !
interface port EXT22
    lacp mode active
    lacp key 2000
    !
failover enable
failover trigger 1 mmon monitor admin-key 2000
failover trigger 1 mmon control member INTA1-INTB14
failover trigger 1 enable
    !
    !
    !
vlag enable
vlag tier-id 1
vlag isl vlan 4094
vlag hlthchk peer-ip 1.1.1.2
vlag isl adminkey 1000

```

```

vlag adminkey 2000 enable
!
!
!
!
!
!
!
!
!
lldp enable
!
interface ip 40
    ip address 1.1.1.1 255.255.255.0
    vlan 4000
    enable
    exit
!
interface ip 92
    ip address 10.1.4.238 255.255.255.0
    vlan 4092
    enable
    exit
!
!
!
!
!
ntp enable
ntp ipv6 primary-server fe80::211:25ff:fec3:9b69 MGT
ntp interval 15
ntp authenticate
ntp primary-key 8811
!
ntp message-digest-key 8811 md5-ekey
1e389d20083088209635f6e3cb802bd2b52a41c0125c9904874d06d2a3af9d16341b4054daa0d14523
ca25ad2e9ec7d8ef2248b85c18a59a2436918a0ee41cea
!
ntp trusted-key 8811
!
end

```

EN4093flex_2

Example 5-219 lists the configuration for the EN4093flex_2 switch.

Example 5-219 EN4093flex_2 switch configuration

```

version "7.3.1"
switch-type "IBM Flex System Fabric EN4093 10Gb Scalable Switch"
!
!

snmp-server name "en4093flex_2"
!
!

```

```

hostname "en4093flex_2"
!
!
interface port INTA1
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port INTB1
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT4
    name "ISL hlthchk"
    pvid 4000
    exit
!
interface port EXT7
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT8
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT9
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT10
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port EXT15
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port EXT16
    name "Link to g8264tor_1"
    tagging
    tag-pvid

```

```

        pvid 4092
        exit
    !
interface port EXT17
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port EXT18
    name "Link to g8264tor_1"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port EXT19
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port EXT20
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port EXT21
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port EXT22
    name "Link to g8264tor_2"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
vlan 1
    member INTA2-INTA14,INTB2-INTB14,EXT1-EXT3,EXT5-EXT6
    no member INTA1,INTB1,EXT4,EXT7-EXT10,EXT15-EXT22
    !
vlan 4000
    enable
    name "ISL hlthchk"
    member EXT4
    !
vlan 4092

```

```

        enable
        name "DATA"
        member INTA1,INTB1,EXT7-EXT10,EXT15-EXT22
    !
vlan 4094
    enable
    name "ISL"
    member EXT7-EXT10
    !
    !
spanning-tree stp 125 vlan 4000
    !
spanning-tree stp 126 vlan 4092
    !
no spanning-tree stp 127 enable
spanning-tree stp 127 vlan 4094
    !
    !
no logging console
    !
interface port EXT7
    lacp mode active
    lacp key 1000
    !
interface port EXT8
    lacp mode active
    lacp key 1000
    !
interface port EXT9
    lacp mode active
    lacp key 1000
    !
interface port EXT10
    lacp mode active
    lacp key 1000
    !
interface port EXT15
    lacp mode active
    lacp key 2000
    !
interface port EXT16
    lacp mode active
    lacp key 2000
    !
interface port EXT17
    lacp mode active
    lacp key 2000
    !
interface port EXT18
    lacp mode active
    lacp key 2000
    !
interface port EXT19
    lacp mode active
    lacp key 2000

```

```

!
interface port EXT20
    lacp mode active
    lacp key 2000
!
interface port EXT21
    lacp mode active
    lacp key 2000
!
interface port EXT22
    lacp mode active
    lacp key 2000
!
failover enable
failover trigger 1 mmon monitor admin-key 2000
failover trigger 1 mmon control member INTA1-INTB14
failover trigger 1 enable
!
!
!
vlag enable
vlag tier-id 1
vlag isl vlan 4094
vlag hlthchk peer-ip 1.1.1.1
vlag isl adminkey 1000
vlag adminkey 2000 enable
!
!
!
!
!
!
!
!
!
lldp enable
!
interface ip 40
    ip address 1.1.1.2 255.255.255.0
    vlan 4000
    enable
    exit
!
interface ip 92
    ip address 10.1.4.239 255.255.255.0
    vlan 4092
    enable
    exit
!
!
!
!
!
ntp enable
ntp ipv6 primary-server fe80::211:25ff:fec3:9b69 MGT

```

```

ntp interval 15
ntp authenticate
ntp primary-key 8811
!
ntp message-digest-key 8811 md5-ekey
ef9d8bb6cf808aa2b6b6e2f70c3029501c9b293eb41d60e5ebbd0fbbd72171ed3c867d24b9976e2052
771345e26681dc63a675b9033673c9923707f9d0f1c078
!
ntp trusted-key 8811
!
end

```

G8264tor_1

Example 5-220 lists the configuration for the G8264tor_1 switch.

Example 5-220 G8264tor_1 switch configuration

```

version "7.4.1"
switch-type "IBM Networking Operating System RackSwitch G8264"
!
!
ssh enable
!

!
!
no system dhcp
no system default-ip
hostname "G8264TOR-1"
!
!
interface port 1
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 2
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 3
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 4
    name "ISL"
    tagging
    pvid 4094
    exit
!

```



```

interface port 5
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 6
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 7
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 8
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 9
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 10
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 11
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 12
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 13
    name "ISL"
    tagging
    pvid 4094
    exit
!
interface port 14

```

```

        name "ISL"
        tagging
        pvid 4094
        exit
    !
interface port 15
    name "ISL"
    tagging
    pvid 4094
    exit
    !
interface port 16
    name "ISL"
    tagging
    pvid 4094
    exit
    !
interface port 18
    name "Po5 to Nexus5548core_1"
    pvid 20
    exit
    !
interface port 20
    name "Po5 to Nexus5548core_1"
    pvid 20
    exit
    !
interface port 22
    name "Po6 to Nexus5548core_2"
    pvid 21
    exit
    !
interface port 24
    name "Po6 to Nexus5548core_2"
    pvid 21
    exit
    !
interface port 25
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 26
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
    !
interface port 27
    name "Link to EN4093-1"
    tagging
    tag-pvid

```

```

    pvid 4092
    exit
!
interface port 28
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 37
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 38
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 39
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 40
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 64
    name "ISL hlthchk"
    pvid 4000
    exit
!
vlan 1
    member 17,19,21,23,25-63
    no member 1-16,18,20,22,24,64
!
vlan 20
    enable
    name "VLAN 20"
    member 18,20
!
vlan 21
    enable
    name "VLAN 21"

```

```

        member 22,24
    !
vlan 4000
    enable
    name "ISL hlthchk"
    member 64
!
vlan 4092
    enable
    name "DATA"
    member 1-16,25-28,37-40
!
vlan 4094
    enable
    name "ISL"
    member 1-16
!
!
portchannel 5 port 18
portchannel 5 port 20
portchannel 5 enable
!
portchannel 6 port 22
portchannel 6 port 24
portchannel 6 enable
!
!
spanning-tree stp 20 vlan 20
!
spanning-tree stp 21 vlan 21
!
spanning-tree stp 125 vlan 4000
!
spanning-tree stp 126 vlan 4092
!
no spanning-tree stp 127 enable
spanning-tree stp 127 vlan 4094
!
!
no logging console
!
interface port 1
    lacp mode active
    lacp key 1000
!
interface port 2
    lacp mode active
    lacp key 1000
!
interface port 3
    lacp mode active
    lacp key 1000
!
interface port 4
    lacp mode active

```

```

        lacp key 1000
    !
interface port 5
    lacp mode active
    lacp key 1000
    !
interface port 6
    lacp mode active
    lacp key 1000
    !
interface port 7
    lacp mode active
    lacp key 1000
    !
interface port 8
    lacp mode active
    lacp key 1000
    !
interface port 9
    lacp mode active
    lacp key 1000
    !
interface port 10
    lacp mode active
    lacp key 1000
    !
interface port 11
    lacp mode active
    lacp key 1000
    !
interface port 12
    lacp mode active
    lacp key 1000
    !
interface port 13
    lacp mode active
    lacp key 1000
    !
interface port 14
    lacp mode active
    lacp key 1000
    !
interface port 15
    lacp mode active
    lacp key 1000
    !
interface port 16
    lacp mode active
    lacp key 1000
    !
interface port 25
    lacp mode active
    lacp key 2002
    !
interface port 26

```

```

        lacp mode active
        lacp key 2002
    !
interface port 27
    lacp mode active
    lacp key 2002
    !
interface port 28
    lacp mode active
    lacp key 2002
    !
interface port 37
    lacp mode active
    lacp key 2002
    !
interface port 38
    lacp mode active
    lacp key 2002
    !
interface port 39
    lacp mode active
    lacp key 2002
    !
interface port 40
    lacp mode active
    lacp key 2002
    !
    !
    !
vlag enable
vlag tier-id 2
vlag isl vlan 4094
vlag hlthchk peer-ip 1.1.1.2
vlag isl adminkey 1000
vlag adminkey 2002 enable
    !
    !
    !
    !
    !
    !
    !
    !
    !
ip router-id 10.10.11.243
    !
interface ip 20
    ip address 10.10.20.2 255.255.255.252
    vlan 20
    enable
    exit
    !
interface ip 21
    ip address 10.10.21.2 255.255.255.252

```

```

    vlan 21
    enable
    exit
!
interface ip 40
    ip address 1.1.1.1 255.255.255.0
    vlan 4000
    enable
    exit
!
interface ip 92
    ip address 10.1.4.243 255.255.255.0
    vlan 4092
    enable
    exit
!
interface ip 128
    ip address 172.25.101.243
    enable
    exit
!
interface loopback 1
    ip address 10.10.11.243 255.255.255.255
    enable
    exit
!
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
!
!
!
!
!
!
router vrrp
    enable
!
    virtual-router 1 virtual-router-id 1
    virtual-router 1 interface 92
    virtual-router 1 priority 120
    virtual-router 1 address 10.1.4.241
    virtual-router 1 enable
!
router ospf
    enable
!
    area 0 enable
!
interface ip 20
    ip ospf enable
!
interface ip 21
    ip ospf enable
!
interface ip 92

```

```
    ip ospf enable
!
ntp enable
ntp primary-server 172.25.101.237 MGT
!
end
```

G8264tor_2

Example 5-221 lists the configuration for the G8264tor_2 switch.

Example 5-221 G8264tor_2 switch configuration

```
version "7.4.1"
switch-type "IBM Networking Operating System RackSwitch G8264"
!
!
ssh enable
!

!
!
no system dhcp
no system default-ip
hostname "G8264TOR-2"
!
!
interface port 1
    name "ISL"
    tagging
    exit
!
interface port 2
    name "ISL"
    tagging
    exit
!
interface port 3
    name "ISL"
    tagging
    exit
!
interface port 4
    name "ISL"
    tagging
    exit
!
interface port 5
    name "ISL"
    tagging
    exit
!
interface port 6
    name "ISL"
    tagging
    exit
```



```

!
interface port 7
    name "ISL"
    tagging
    exit
!
interface port 8
    name "ISL"
    tagging
    exit
!
interface port 9
    name "ISL"
    tagging
    exit
!
interface port 10
    name "ISL"
    tagging
    exit
!
interface port 11
    name "ISL"
    tagging
    exit
!
interface port 12
    name "ISL"
    tagging
    exit
!
interface port 13
    name "ISL"
    tagging
    exit
!
interface port 14
    name "ISL"
    tagging
    exit
!
interface port 15
    name "ISL"
    tagging
    exit
!
interface port 16
    name "ISL"
    tagging
    exit
!
interface port 18
    name "Po5 to Nexus5548core_2"
    pvid 23
    exit

```

```

!
interface port 20
    name "Po5 to Nexus5548core_2"
    pvid 23
    exit
!
interface port 22
    name "Po6 to Nexus5548core_1"
    pvid 22
    exit
!
interface port 24
    name "Po6 to Nexus5548core_1"
    pvid 22
    exit
!
interface port 25
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 26
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 27
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 28
    name "Link to EN4093-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 37
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 38
    name "Link to EN4093-2"
    tagging
    tag-pvid

```

```

    pvid 4092
    exit
!
interface port 39
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 40
    name "Link to EN4093-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 64
    name "ISL hlthchk"
    pvid 4000
    exit
!
vlan 1
    member 1-17,19,21,23,25-63
    no member 18,20,22,24,64
!
vlan 22
    enable
    name "VLAN 22"
    member 22,24
!
vlan 23
    enable
    name "VLAN 23"
    member 18,20
!
vlan 4000
    enable
    name "ISL hlthchk"
    member 64
!
vlan 4092
    enable
    name "DATA"
    member 1-16,25-28,37-40
!
vlan 4094
    enable
    name "ISL"
    member 1-16
!
!
portchannel 5 port 18
portchannel 5 port 20
portchannel 5 enable

```

```

!
portchannel 6 port 22
portchannel 6 port 24
portchannel 6 enable
!
!
spanning-tree stp 22 vlan 22
!
spanning-tree stp 23 vlan 23
!
spanning-tree stp 125 vlan 4000
!
spanning-tree stp 126 vlan 4092
!
no spanning-tree stp 127 enable
spanning-tree stp 127 vlan 4094
!
!
interface port 1
    lacp mode active
    lacp key 1000
!
interface port 2
    lacp mode active
    lacp key 1000
!
interface port 3
    lacp mode active
    lacp key 1000
!
interface port 4
    lacp mode active
    lacp key 1000
!
interface port 5
    lacp mode active
    lacp key 1000
!
interface port 6
    lacp mode active
    lacp key 1000
!
interface port 7
    lacp mode active
    lacp key 1000
!
interface port 8
    lacp mode active
    lacp key 1000
!
interface port 9
    lacp mode active
    lacp key 1000
!
interface port 10

```

```

        lacp mode active
        lacp key 1000
    !
interface port 11
    lacp mode active
    lacp key 1000
    !
interface port 12
    lacp mode active
    lacp key 1000
    !
interface port 13
    lacp mode active
    lacp key 1000
    !
interface port 14
    lacp mode active
    lacp key 1000
    !
interface port 15
    lacp mode active
    lacp key 1000
    !
interface port 16
    lacp mode active
    lacp key 1000
    !
interface port 25
    lacp mode active
    lacp key 2002
    !
interface port 26
    lacp mode active
    lacp key 2002
    !
interface port 27
    lacp mode active
    lacp key 2002
    !
interface port 28
    lacp mode active
    lacp key 2002
    !
interface port 37
    lacp mode active
    lacp key 2002
    !
interface port 38
    lacp mode active
    lacp key 2002
    !
interface port 39
    lacp mode active
    lacp key 2002
    !

```

```

interface port 40
    lacp mode active
    lacp key 2002
!
!
!
vlag enable
vlag tier-id 2
vlag isl vlan 4094
vlag hlthchk peer-ip 1.1.1.1
vlag isl adminkey 1000
vlag adminkey 2002 enable
!
!
!
!
!
!
!
!
!
!
ip router-id 10.10.11.244
!
interface ip 22
    ip address 10.10.22.2 255.255.255.252
    vlan 22
    enable
    exit
!
interface ip 23
    ip address 10.10.23.2 255.255.255.252
    vlan 23
    enable
    exit
!
interface ip 40
    ip address 1.1.1.2 255.255.255.0
    vlan 4000
    enable
    exit
!
interface ip 92
    ip address 10.1.4.244 255.255.255.0
    vlan 4092
    enable
    exit
!
interface ip 128
    ip address 172.25.101.244
    enable
    exit
!
interface loopback 1
    ip address 10.10.11.244 255.255.255.255

```

```

        enable
    exit
!
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
!
!
!
!
!
!
router vrrp
    enable
!
    virtual-router 1 virtual-router-id 1
    virtual-router 1 interface 92
    virtual-router 1 priority 110
    virtual-router 1 address 10.1.4.241
    virtual-router 1 enable
!
router ospf
    enable
!
    area 0 enable
!
interface ip 22
    ip ospf enable
!
interface ip 23
    ip ospf enable
!
interface ip 92
    ip ospf enable
!
ntp enable
ntp primary-server 172.25.101.237 MGT
!
end

```

Nexus5548core_1 switch

Example 5-222 lists the configuration of the Nexus5548core_1 switch.

Example 5-222 Nexus5548core_1 switch configuration

```

!Command: show startup-config
!Time: Wed Oct 24 21:39:06 2012
!Startup config saved at: Wed Oct 24 21:38:37 2012

version 5.2(1)N1(1b)
logging level feature-mgr 0
hostname Nexus5548core_1

feature telnet
feature vrrp

```

```

cfs ipv4 distribute
cfs eth distribute
feature ospf
feature interface-vlan
feature lacp
feature lldp

username admin password 5 $1$huQeFTJf$dYim2oGvqYAGk3THH5KP.0 role network-admin

banner motd #Nexus 5000 Switch
#

no ip domain-lookup
class-map type qos class-fcoe
class-map type queuing class-fcoe
    match qos-group 1
class-map type queuing class-all-flood
    match qos-group 2
class-map type queuing class-ip-multicast
    match qos-group 2
class-map type network-qos class-fcoe
    match qos-group 1
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
snmp-server user admin network-admin auth md5 0x50d80b5959ad2a911a11fcaa8453db8a
priv 0x50d80b5959ad2a911a11fcaa8453db8a localizedkey

vrf context management
    ip route 0.0.0.0/0 172.25.1.1
vrf context VPCKeepAlive
vlan 1
vlan 30
    name Server
spanning-tree vlan 1000 priority 24576
port-profile default max-ports 512

interface Vlan1

interface Vlan30
    no shutdown
    ip address 10.10.30.2/24
    ip router ospf 100 area 0.0.0.0
    vrrp 1
        priority 200
        address 10.10.30.1
    no shutdown

interface port-channel5
    no switchport
    ip address 10.10.20.1/30
    ip router ospf 100 area 0.0.0.0

```



```

interface port-channel6
  no switchport
  ip address 10.10.22.1/30
  ip router ospf 100 area 0.0.0.0

interface port-channel100
  description Switch-to-Switch Link
  switchport mode trunk
  switchport trunk allowed vlan 30
  spanning-tree port type network

interface Ethernet1/1

interface Ethernet1/2

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7
  description Po5 to G8264tor_1
  no switchport
  speed auto
  channel-group 5

interface Ethernet1/8
  description Po5 to G8264tor_1
  no switchport
  speed auto
  channel-group 5

interface Ethernet1/9
  description Po6 to G8264tor_2
  no switchport
  speed auto
  channel-group 6

interface Ethernet1/10
  description Po6 to G8264tor_2
  no switchport
  speed auto
  channel-group 6

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

```

```

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17
    description Po100 to Nexus5548core_2
    switchport mode trunk
    switchport trunk allowed vlan 30
    channel-group 100 mode active

interface Ethernet1/18

interface Ethernet1/19
    description Po100 to Nexus5548core_2
    switchport mode trunk
    switchport trunk allowed vlan 30
    channel-group 100 mode active

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface mgmt0
    ip address 172.25.101.249/16

interface loopback1
    description OSPF router-id
    ip address 10.10.11.249/32
    cli alias name wr copy run start
    line console
    line vty
    boot kickstart bootflash:/n5000-uk9-kickstart.5.2.1.N1.1b.bin
    boot system bootflash:/n5000-uk9.5.2.1.N1.1b.bin

```

```
router ospf 100
  router-id 10.10.11.249
  log-adjacency-changes
  auto-cost reference-bandwidth 100 Gbps
```

Nexus5548core_2 switch

Example 5-223 lists the configuration of the Nexus5548core_2 switch.

Example 5-223 Nexus5548core_2 switch configuration

```
!Command: show startup-config
!Time: Tue Aug  4 21:49:38 2009
!Startup config saved at: Tue Aug  4 21:49:16 2009

version 5.2(1)N1(1b)
logging level feature-mgr 0
hostname Nexus5548core_2

feature telnet
feature vrrp
cfs ipv4 distribute
cfs eth distribute
feature ospf
feature interface-vlan
feature lacp
feature lldp

username admin password 5 $1$huQeFTJf$dYim2oGvqYAGk3THH5KP.0 role network-admin
no password strength-check

banner motd #Nexus 5000 Switch#

no ip domain-lookup
class-map type qos class-fcoe
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
policy-map type control-plane copp-system-policy-customized
  class copp-system-class-default
    police cir 2048 kbps bc 6400000 bytes
snmp-server user admin network-admin auth md5 0x50d80b5959ad2a911a11fcaa8453db8a
priv 0x50d80b5959ad2a911a11fcaa8453db8a localizedkey

vrf context management
  ip route 0.0.0.0/0 172.25.1.1
```

```

vrf context VPCKeepAlive
vlan 1
vlan 30
    name Server
port-profile default max-ports 512

interface Vlan1

interface Vlan30
    no shutdown
    ip address 10.10.30.3/24
    ip router ospf 100 area 0.0.0.0
    vrrp 1
        priority 150
        address 10.10.30.1
        no shutdown

interface port-channel5
    no switchport
    ip address 10.10.23.1/30
    ip router ospf 100 area 0.0.0.0

interface port-channel6
    no switchport
    ip address 10.10.21.1/30
    ip router ospf 100 area 0.0.0.0

interface port-channel100
    description Switch-to-Switch link
    switchport mode trunk
    switchport trunk allowed vlan 30
    spanning-tree port type network

interface Ethernet1/1

interface Ethernet1/2

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7
    description Po5 to G8264tor_2
    no switchport
    channel-group 5

interface Ethernet1/8
    description Po5 to G8264tor_2
    no switchport
    channel-group 5

```

```

interface Ethernet1/9
  description Po6 to G8264tor_1
  no switchport
  channel-group 6

interface Ethernet1/10
  description Po6 to G8264tor_1
  no switchport
  channel-group 6

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17
  description Po100 to Nexus5548core_1
  switchport mode trunk
  switchport trunk allowed vlan 30
  channel-group 100 mode active

interface Ethernet1/18

interface Ethernet1/19
  description Po100 to Nexus5548core_1
  switchport mode trunk
  switchport trunk allowed vlan 30
  channel-group 100 mode active

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

```

```
interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface mgmt0
  ip address 172.25.101.200/16

interface loopback1
  description OSPF router-id
  ip address 10.10.11.200/32
cli alias name wr copy run start
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.2.1.N1.1b.bin
boot system bootflash:/n5000-uk9.5.2.1.N1.1b.bin
router ospf 100
  router-id 10.10.11.200
  log-adjacency-changes
  auto-cost reference-bandwidth 100 Gbps
```



Troubleshooting and maintenance

This chapter addresses the troubleshooting and maintenance steps on IBM PureFlex Systems switches, with emphasis on EN4093 switch.

This chapter includes the following sections:

- ▶ Troubleshooting
- ▶ Configuration management
- ▶ Firmware management
- ▶ Logging and reporting

6.1 Troubleshooting

This section introduces the basic troubleshooting tools and techniques. It addresses various troubleshooting steps, such as inspecting LEDs on the switch, troubleshooting network connectivity, port mirroring for capturing data traffic, and the use of serial connection.

6.1.1 Basic troubleshooting procedures

This section contains basic troubleshooting information to help resolve problems that might occur during the installation and operation of your EN4093 switch. Before getting started, download and use the EN4093 documentation, available on the IBM Flex System Fabric EN4093 10Gb Scalable Switch InfoCenter at:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.net.workdevices.doc/Io_module_compass.html

LEDs on EN4093

EN4093 switch contains the following LEDs for easy identification of switch and port status:

- System status LEDs (Figure 6-1)




Figure 6-1 System status LEDs: OK, Identify, and Error (left to right)

The system status LEDs (OK, Identify and Error) have the following meanings:

- OK (green)

When this LED is lit, it indicates that the switch is powered on.

When this LED is not lit, but the yellow Error LED is lit, it indicates a critical alert.

When both LEDs are off, it indicates that the switch is off.

- Identify (blue)

You can use this LED to identify the location of switch in chassis. Use CMM web interface to change the state of this LED:

- Click **Chassis Management** → **I/O Modules** in the CMM web GUI as shown in Figure 6-2.

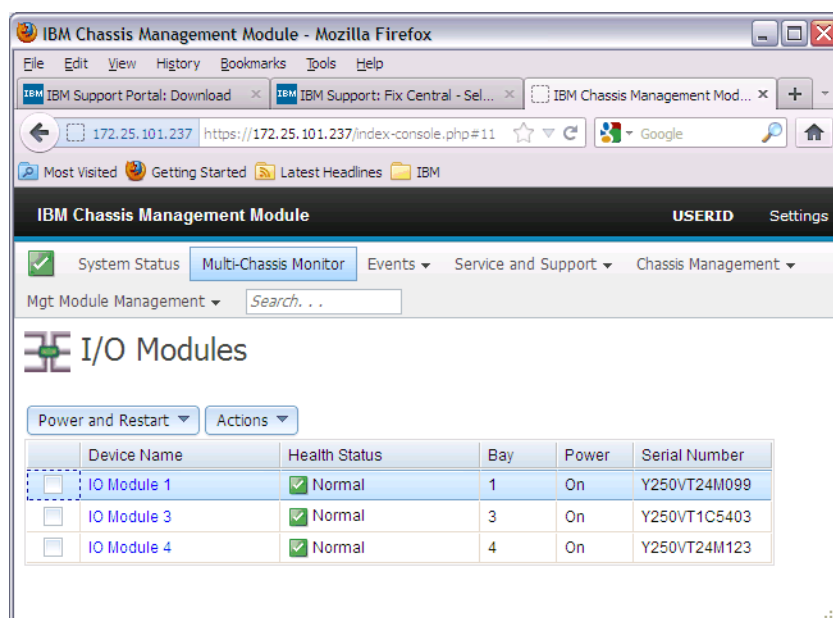


Figure 6-2 Selecting I/O module management

- Click the I/O module that you want to identify. In this case, click **IO Module 1**. This opens the window shown in Figure 6-3.

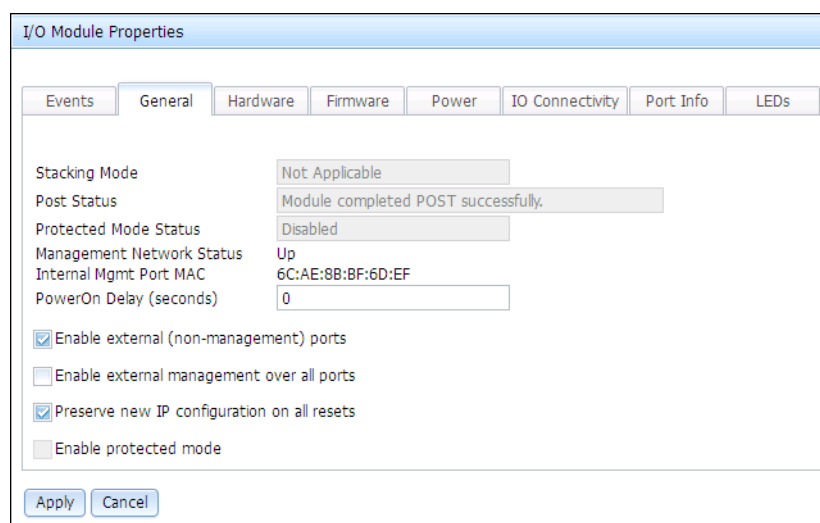


Figure 6-3 I/O module properties

- iii. Click **LEDs** tab to display the window that is shown in Figure 6-4.




Figure 6-4 Toggling the Identify LED state

You can now toggle the Identify LED state for easy identification of switch in the chassis.

- Error (yellow)

When this LED is lit, it indicates a critical alert or POST failure.

- SFP+ and QSFP+ module port LEDs (Figure 6-5 and Figure 6-6 on page 269).




Figure 6-5 SFP+ port LEDs

Figure 6-6 shows the LEDs for the QSFP+ port.




Figure 6-6 QSFP+ port LEDs

The Link and Tx/Rx LEDs have these functions:

- Link (green)

When this LED is lit, there is an active connection between the port and the connected device.

When the LED is not lit, there is no signal on the port, or the link is down.

- Tx/Rx (green)

When this LED is flashing, link activity is occurring on the port.

Port link LED does not light

Symptom: The port link LED does not light.

Solution 1: Check the port configuration. If the port is configured with a specific speed or duplex mode, check the other device to verify that it is set to the same configuration. If the switch port is set to autonegotiate, verify that the other device is also set to autonegotiate.

Solution 2: Check the cables that connect the port to the other device. Make sure that they are connected. Verify that you are using the correct cable type.

Switch does not boot

Symptom: All the switch LEDs stay on, and the command prompt does not appear on the console.

Solution: The switch firmware might be damaged. Use the console port to run a serial upgrade of the switch firmware. For more information, see 6.3.3, “Recovering from a failed firmware upgrade” on page 287.

6.1.2 Connectivity troubleshooting

This section contains basic information about how to troubleshoot the IP connectivity in a network built on IBM System Networking switches. IBM switches come with a set of simple tools that can be helpful for troubleshooting IP connectivity issues.

Ping

The **ping** command is a simple tool, based on a request-response mechanism, to verify connectivity to a remote network node. The **ping** command is based on ICMP. The request is an ICMP Echo packet, and the reply is an ICMP Echo Reply. Like a regular IP packet, an ICMP packet is forwarded based on the intermediate routers' routing table until it reaches the destination. After it reaches the destination, the ICMP Echo Reply packet is generated and forwarded back to the originating node.

Important: In IBM switches, **ping** sends an ICMP Echo packet on the management interface first. If you want to change that option, you must add the data-port keyword to a command as a parameter.

Example 6-1 shows the use of **ping** command to verify connectivity between the switch and IP address 172.25.101.237.

Example 6-1 Ping command example

```
en4093flex_1#ping 172.25.101.237
Connecting via MGT port.
[host 172.25.101.237, max tries 5, delay 1000 msec, length 0, ping

source N/S, ttl 255, tos 0]
172.25.101.237: #1 ok, RTT 1 msec.
172.25.101.237: #2 ok, RTT 2 msec.
172.25.101.237: #3 ok, RTT 2 msec.
172.25.101.237: #4 ok, RTT 1 msec.
172.25.101.237: #5 ok, RTT 2 msec.
Ping finished.
```

You can see in the output that all five ICMP Echo requests received the replies. There is also more information about the Round Trip Time (RTT), that is, the time it took for the switch to receive response.

Traceroute

You can use the **traceroute** command to not only verify connectivity to a remote network node, but to track the responses from intermediate nodes as well. This action is done by using the time to live (TTL) field in IP packets. The **traceroute** command sends a UDP packet to a port that is not likely to be used on a remote node with a TTL of 1. After the packet reaches the intermediate router, the TTL is decremented. The ICMP time-exceeded message is then sent back to the originating node, which increments the TTL to 2, and the process repeats. After the UDP packet reaches a destination host, an ICMP port-unreachable message is sent back to the sender. This action provides the sender with information about all intermediate routers on the way to the destination.

The command shown in Example 6-2 verifies which hops are on the way from switch to the system with IP address 10.0.100.1.

Example 6-2 Traceroute command example

```
ACC-2#traceroute 10.0.100.1 data-port
Connecting via DATA port.
[host 10.0.100.1, max-hops 32, delay 2048 msec]
 1  10.0.100.1      0 ms
Trace host responded.
```

From the output, you see that there is only one hop on the way from switch to destination. OSPF in this network, which selects this path as the shortest one.

For test purposes, shut down the direct link between the switch and target system and run **traceroute** again. The output is shown in Example 6-3.

Example 6-3 Traceroute command example without direct link

```
ACC-2#traceroute 10.0.100.1 data-port
Connecting via DATA port.
[host 10.0.100.1, max-hops 32, delay 2048 msec]
 1  10.0.104.1      0 ms
 2  10.0.100.1      1 ms
Trace host responded.
```

Now, to reach destination, the switch uses the 10.0.104.1 system as the intermediate router.

6.1.3 Port mirroring

You can use the IBM System Networking switches port mirroring feature to mirror (copy) the packets of a target port, and forward them to a monitoring port. Port mirroring functions for all Layer 2 and Layer 3 traffic on a port. This feature can be used as a troubleshooting tool or to enhance the security of your network.

For example, an intrusion detection system (IDS) server or other traffic sniffer device or analyzer can be connected to the monitoring port to detect intruders that attack the network.

IBM System Networking switches support a “many to one” mirroring model. As shown in Figure 6-7, selected traffic for ports 1 and 2 is being monitored by port 3. In the example, both ingress traffic and egress traffic on port 2 are copied and forwarded to the monitor. However, port 1 mirroring is configured so that only ingress traffic is copied and forwarded to the monitor. A device that is attached to port 3 can capture and analyze the resulting mirrored traffic.

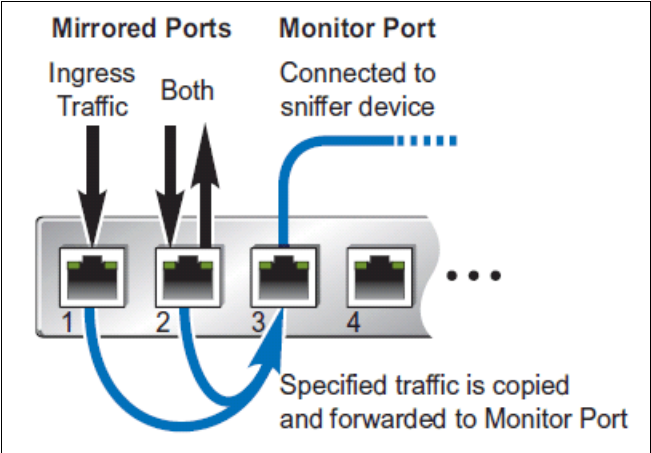


Figure 6-7 Mirroring ports

The monitored packets in the EN4093 have the following composition, based on the configuration of the ports:

- ▶ Packets that are mirrored at port egress are mirrored before VLAN tag processing. They can have a different PVID than packets that egress the port toward their actual network destination.
- ▶ Packets that are mirrored at port ingress are not modified.

Example 6-4 shows the ISCLI commands to enable port mirroring and to mirror ingress and egress traffic on ports EXT1 - EXT4 to monitoring port EXT6.

Example 6-4 Port mirroring ISCLI commands

```
en4093flex_1(config)#port-mirroring enable
en4093flex_1(config)#port-mirroring monitor-port EXT6 mirroring-port EXT1-EXT4
both
```

You can check the port mirroring configuration with ISCLI command **show port-mirroring**. As shown in Example 6-5, both ingress and egress traffic on ports EXT1 - EXT4 is mirrored to monitoring port EXT6.

Example 6-5 Port mirroring configuration verification

```
en4093flex_1(config)#show port-mirroring
Port Mirroring is enabled

Monitoring port  Mirrored ports
INTA1            none
INTA2            none
INTA3            none
...
Lines deleted for clarity
...
```

```
EXT5          none
EXT6          (EXT1,both) (EXT2,both) (EXT3,both) (EXT4,both)
EXT7          none
...
Lines deleted for clarity
...
```

6.1.4 Serial cable troubleshooting procedures

When all else fails, you can use the serial cable that is delivered with EN4093 to connect to the switch and investigate the problem. A terminal emulation utility must run on management system (such as Windows Hyperterminal or PuTTY). Use the following serial connection parameters:

- ▶ Speed: 9600 bps
- ▶ Data Bits: 8
- ▶ Stop Bits: 1
- ▶ Parity: None
- ▶ Flow Control: None

When the serial session is established, you must reboot the EN4093 switch to start the Boot Management Menu with recovery options. In the CMM web GUI, you can either power-cycle the affected EN4093 switch, or restart it.

When you see the memory test run in terminal window, press Shift+B to display the menu with recovery options. Example 6-6 shows the Boot Management Menu.

Example 6-6 Boot Management Menu

```
Resetting the System ...
Memory Test .....
Boot Management Menu
    1 - Change booting image
    2 - Change configuration block
    3 - Boot in recovery mode (tftp and xmodem download of images to recover
switch)
    4 - Xmodem download (for boot image only - use recovery mode for
application images)
    5 - Reboot
    6 - Exit
Please choose your menu option:
```

Using the Boot Management Menu, you can perform the following tasks:

- ▶ Change the active boot image from image1 to image2 or vice versa. For more information, see “Changing the boot image by using the serial interface” on page 283.
- ▶ Change the active configuration block. You can select between active, backup, and factory default configuration blocks. This option can be used to restore the EN4093 switch to factory defaults. For more information, see “Resetting with no terminal access to the switch” on page 281.
- ▶ Download new firmware to the switch. This option can be helpful if you must recover the switch after a failed firmware upgrade. For more information, see 6.3.3, “Recovering from a failed firmware upgrade” on page 287.

6.2 Configuration management

This section describes how to manage configuration files, and how to save and restore a configuration in the switch.

6.2.1 Configuration files

The switch stores its configuration in two files:

- ▶ `startup-config` is the configuration that the switch uses when it is reloaded.
- ▶ `running-config` is the configuration that reflects all the changes you made from the CLI. It is stored in memory, and is lost after the reload of the switch.

6.2.2 Configuration blocks

The switch stores its configuration in one of two configuration blocks:

- ▶ `active-config` is stored in the active configuration block.
- ▶ `backup-config` is stored in the backup configuration block.

When you save the running configuration (`copy running-config startup-config`), the new configuration is placed into the active configuration block. The previous configuration is copied into the backup configuration block.

In addition, there is also a factory configuration block. This block holds the factory default configuration, which you can use to restore the switch to factory defaults if needed.

This setup has the flexibility that you need to manage the configuration of the switch and run a configuration rollback.

Use the following command to select configuration block the switch will load on next reboot:

```
Switch# boot configuration-block {active|backup|factory}
```

6.2.3 Managing configuration files

This section describes the different ways of managing the configuration files.

Managing the configuration using ISCLI

You can manage the configuration files by using these commands:

- ▶ Run the following command to display the current configuration file:

```
Switch#show running-config
```

- ▶ Run the following command to copy the current (running) configuration from switch memory to the `startup-config` partition:

```
Switch#copy running-config startup-config
```

The following command also copies running configuration to the startup configuration:

```
Switch#write memory
```

- ▶ Run the following command to copy the current (running) configuration from switch memory to the `backup-config` block:

```
Switch#copy running-config backup-config
```


- ▶ Run the following command to back up the current configuration to a file on an FTP/TFTP server:

```
Switch#copy running-config {ftp|tftp}
```

- ▶ Run the following command to restore the current configuration from an FTP/TFTP server:

```
Switch#copy {ftp|tftp} running-config
```

Managing the configuration through SNMP

This section describes how to use MIB calls to work with switch configuration files.

You can use a standard SNMP tool to perform the actions, using the MIBs listed in Table 6-1. For more information about how to set up your switch to use SNMP, see 6.4.2, “SNMP” on page 292.

Table 6-1 SNMP MIBs for managing switch configuration and firmware

MIB name	MIB OID
agTransferServer	1.3.6.1.4.1872.2.5.1.1.7.1.0
agTransferImage	1.3.6.1.4.1872.2.5.1.1.7.2.0
agTransferImageFileName	1.3.6.1.4.1872.2.5.1.1.7.3.0
agTransferCfgFileName	1.3.6.1.4.1872.2.5.1.1.7.4.0
agTransferDumpFileName	1.3.6.1.4.1872.2.5.1.1.7.5.0
agTransferAction	1.3.6.1.4.1872.2.5.1.1.7.6.0
agTransferLastActionStatus	1.3.6.1.4.1872.2.5.1.1.7.7.0
agTransferUserName	1.3.6.1.4.1872.2.5.1.1.7.9.0
agTransferPassword	1.3.6.1.4.1.1872.2.5.1.1.7.10.0
agTransferTSDumpFileName	1.3.6.1.4.1.1872.2.5.1.1.7.11.0

The following configuration-related SNMP actions can be performed by using the MIBs listed in Table 6-1:

- ▶ Load a previously saved switch configuration from an FTP/TFTP server.
- ▶ Save the switch configuration to an FTP/TFTP server.

You can also use the SNMP MIBs in Table 6-1 to perform other functions, such as upgrading the switch firmware and saving the switch dump to an FTP/TFTP server.

Loading a saved configuration

To load a saved switch configuration with the name `MyRunningConfig.cfg` into the switch, complete the following steps. This example shows a TFTP server at IPv4 address 172.25.101.200 (although IPv6 is also supported) where the previously saved configuration is available for download.

1. Set the FTP/TFTP server address where the switch configuration file is located:

```
Set agTransferServer.0 "172.25.101.200"
```

2. Set the name of the configuration file:

```
Set agTransferCfgFileName.0 "MyRunningConfig.cfg"
```

3. If you are using an FTP server, enter a user name:

```
Set agTransferUserName.0 "MyName"
```

4. If you are using an FTP server, enter a password:

```
Set agTransferPassword.0 "MyPassword"
```

5. Initiate the transfer. To restore a running configuration, use transfer action 3:

```
Set agTransferAction.0 "3"
```

Saving the configuration

To save the switch configuration to an FTP/TFTP server, complete the following steps. This example shows an FTP/TFTP server at IPv4 address 172.25.101.200, although IPv6 is also supported.

1. Set the FTP/TFTP server address where the configuration file is saved:

```
Set agTransferServer.0 "172.25.101.200"
```

2. Set the name of the configuration file:

```
Set agTransferCfgFileName.0 "MyRunningConfig.cfg"
```

3. If you are using an FTP server, enter a user name:

```
Set agTransferUserName.0 "MyName"
```

4. If you are using an FTP server, enter a password:

```
Set agTransferPassword.0 "MyPassword"
```

5. Initiate the transfer. To save a running configuration file, use transfer action 4.

```
Set agTransferAction.0 "4"
```

Other tasks: Saving a switch dump

SNMP MIBs are not only useful to save and load switch configuration. You can also perform other tasks, such as saving a switch dump. To save a switch dump to an FTP/TFTP server, complete the following steps. This example shows an FTP/TFTP server at 172.25.101.200, although IPv6 is also supported.

1. Set the FTP/TFTP server address where the configuration is saved:

```
Set agTransferServer.0 "172.25.101.200"
```

2. Set the name of the dump file:

```
Set agTransferDumpFileName.0 "MyDumpFile.dmp"
```

3. If you are using an FTP server, enter a user name:

```
Set agTransferUserName.0 "MyName"
```

4. If you are using an FTP server, enter a password:

```
Set agTransferPassword.0 "MyPassword"
```

5. Initiate the transfer. To save a dump file, use transfer action 5.

```
Set agTransferAction.0 "5"
```

6.2.4 Resetting to factory defaults

You might need to reset the switch to factory defaults in certain situations. For example, when you redeploy the switch for use in a different scenario, or when you troubleshoot a configuration issue. To reset the switch to factory defaults, you must perform one of the following procedures.

Resetting EN4093 to factory defaults by using CMM

Complete these steps to reset EN4093 to factory defaults by using CMM:

1. Point your web browser to CMM IP address, and log in as shown in Figure 6-8.

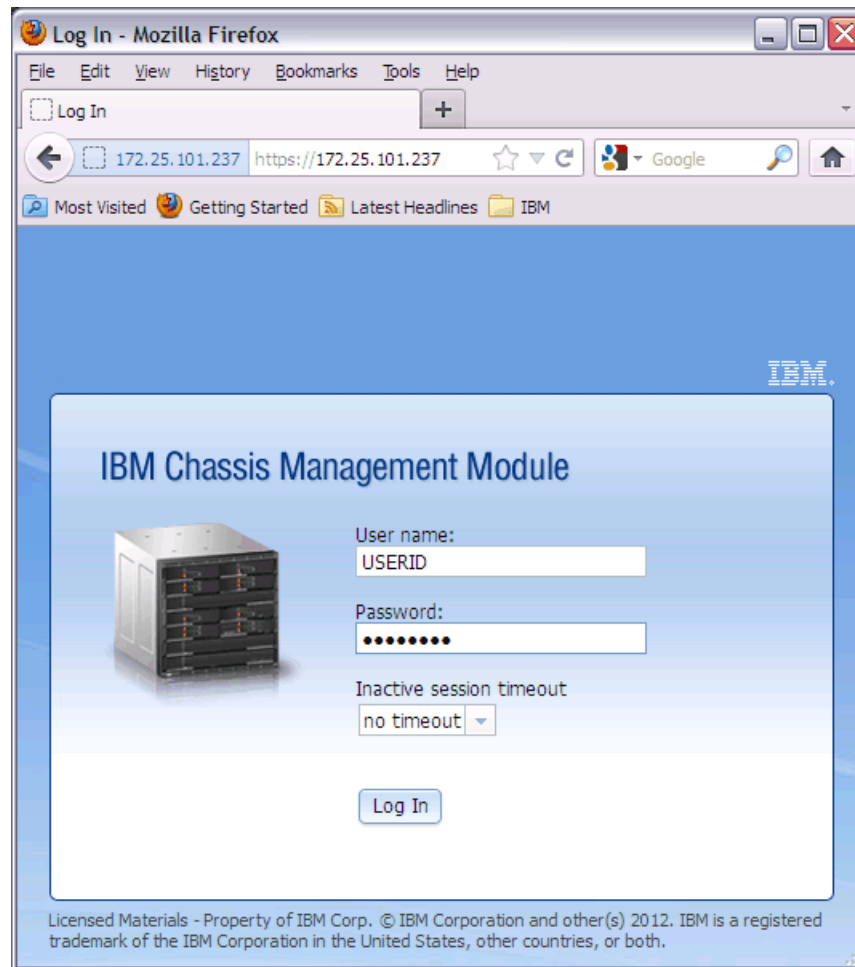


Figure 6-8 Logging in to CMM

After successful login, CMM GUI displays as shown in Figure 6-9.

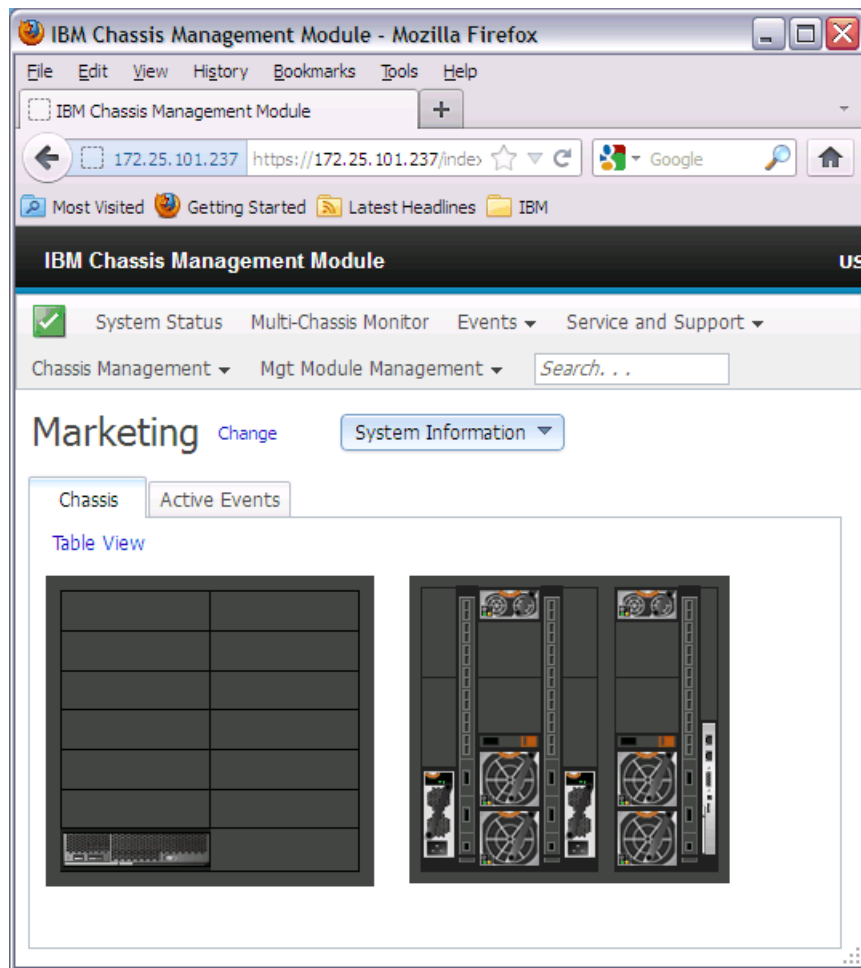


Figure 6-9 CMM GUI

2. Select **Chassis Management** → **I/O Modules** as shown in Figure 6-10.




Figure 6-10 Selecting I/O Modules management

3. Select the I/O module that you want to reset to factory defaults, and click **Actions** → **Restore Factory Defaults** as shown in Figure 6-11.




Figure 6-11 I/O Module 1 restore factory defaults

Resetting with terminal access to the switch

If you have terminal access to the switch and want to reset the switch to factory defaults, use the ISCLI command **boot configuration-block factory** and reload the switch as shown in Example 6-7.

Example 6-7 Resetting to factory defaults using ISCLI

```
EN4093flex_2(config)#boot configuration-block factory
Next boot will use factory default config block instead of active.
```

```
EN4093flex_2(config)#reload
```

Reset will use software "image2" and the factory default config block.

>> Note that this will RESTART the Spanning Tree,

>> which will likely cause an interruption in network service.

Confirm reload (y/n) ? y

The switch reloads with the factory default configuration.

Resetting with no terminal access to the switch

If you want to reset the switch to factory defaults and have no terminal access, you can use the serial console port. Complete the following steps:

1. Connect the management system to the serial port on the switch. Run a terminal emulation utility (such as Windows Hyperterminal or PuTTY) and use the following communication parameters to establish a session:
 - Speed: 9600 bps
 - Data Bits: 8
 - Stop Bits: 1
 - Parity: None
 - Flow Control: None
2. Restart the switch by powering it off and back on, or by restarting it in CMM web interface.
3. Interrupt the boot process and enter the Boot Management menu from the serial console port. When the system shows Memory Test, press Shift+B. The Boot Management Menu opens as shown in Example 6-8.

Example 6-8 Boot Management Menu

```
Boot Management Menu
  1 - Change booting image
  2 - Change configuration block
  3 - Boot in recovery mode (tftp and xmodem download of images to
recover switch)
  4 - Xmodem download (for boot image only - use recovery mode for
application images)
  5 - Reboot
  6 - Exit
Please choose your menu option:
```

4. Enter 2 to change the configuration block (Example 6-9).

Example 6-9 Changing the configuration block

```
Please choose your menu option: 2

Unknown current config block 255
Enter configuration block: a, b or f (active, backup or factory):
```

5. As displayed in Example 6-10, enter f to use the factory defaults configuration block.

Example 6-10 Using the factory defaults configuration block

```
Enter configuration block: a, b or f (active, backup or factory): f
```

6. The initial menu is displayed again. Enter 6 to exit and reset the switch with the default configuration as shown in Example 6-11.

Example 6-11 Exiting from the Boot Management Menu

```
Boot Management Menu
  1 - Change booting image
  2 - Change configuration block
  3 - Boot in recovery mode (tftp and xmodem download of images to
recover switch)
  4 - Xmodem download (for boot image only - use recovery mode for
application images)
```

```
5 - Reboot
6 - Exit
Please choose your menu option: 6
```

The switch resets to the factory default configuration.

Important: If you set the configuration block to factory, do not forget to change it back to active configuration by running the following command:

```
Switch(config)#boot configuration-block active
```

6.2.5 Password recovery

To perform password recovery, you must set the switch to the factory default by using one of the procedures that are described in 6.2.4, “Resetting to factory defaults” on page 276.

After you reset the switch, run the following command:

```
Switch#copy active-config running-config
```

After the command finishes running, the switch is in enable mode without a password. Change the password by running **password** in configuration mode:

```
Switch(config)#password
```

6.3 Firmware management

The switch firmware is the executable code that runs on the switch. The device comes preinstalled with a certain firmware level. As new firmware versions are released, upgrade the code that runs on your switch. You can find the latest version of firmware that is supported for your switch on the IBM Fix Central website at:

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

6.3.1 Firmware images

IBM switches can store up to two different IBM NOS (OS) images (called image1 and image2) and special boot image (called boot). When you load new firmware, make sure that you upgrade both the OS and boot image.

Run the ISCLI command **show boot** to see what images are installed. The output is shown in Example 6-12.

Example 6-12 Showing the current version of boot and OS images on the switch

```
EN4093flex_2#show boot
Currently set to boot software image1, active config block.
NetBoot: disabled, NetBoot tftp server: , NetBoot cfgfile:
Current CLI mode set to IBMNOS-CLI with selectable prompt enabled.
Current FLASH software:
  image1: version 7.2.2.2, downloaded 14:55:26 Mon Jun 18, 2012
  image2: version 7.3.1, downloaded 22:55:05 Mon Oct 1, 2012
  boot kernel: version 7.3.1
Currently scheduled reboot time: none
```

In Example 6-12 on page 282, you can see that the system has two OS images:

- ▶ image1: Version 7.2.2.2
- ▶ image2: Version 7.3.1

The boot image version is 7.3.1. However, the switch is set to boot from OS image1, which is at version 7.2.2.2. Make sure that the switch uses the same version for boot image and OS image. To boot from OS image2, run the command **boot image image2**, as shown in Example 6-13.

Example 6-13 Changing to boot from image2

```
EN4093flex_2(config)#boot image image2
Next boot will use switch software image2 instead of image1.
```

Changing the boot image by using the serial interface

You can use the serial connection and Boot Management Menu to change the boot image. Complete these steps:

1. Connect serial cable to the switch serial management port and the management system.
2. Start the terminal emulation utility on the management system.
3. Use the following set of parameters to establish terminal emulation session:
 - Speed: 9600 bps
 - Data Bits: 8
 - Stop Bits: 1
 - Parity: None
 - Flow Control: None
4. When the system shows Memory Test, press Shift+B. The Boot Management Menu is displayed as shown in Example 6-14.

Example 6-14 Boot Management Menu

```
Boot Management Menu
  1 - Change booting image
  2 - Change configuration block
  3 - Boot in recovery mode (tftp and xmodem download of images to
recover switch)
  4 - Xmodem download (for boot image only - use recovery mode for
application images)
  5 - Reboot
  6 - Exit
Please choose your menu option: 1

Current boot image is 1. Enter image to boot: 1 or 2: 2
Booting from image 2
```

5. Select menu option 1 to change boot image from image1 to image2.

6.3.2 Upgrading the firmware with ISCLI

This section shows how to upgrade the firmware of Flex System embedded switch EN4093. The latest firmware version at the time of writing was 7.3.1.0. This code level is available on IBM Fix Central and on the following link:

<http://www.ibm.com/support/entry/portal/docdisplay?lnodocid=migr-5090394>

To upgrade the firmware, complete these steps:

1. First, download the code update package (either from IBM Fix Central or from the link above) and unpack it. The update package contains two image files:

- Boot image file *GbScSE-10G-7.3.1.0_Boot.img*
- OS image file *GbScSE-10G-7.3.1.0_OS.img*

For convenience, rename these files as follows:

- Boot image file *7310boot.img*
- OS image file *7310os.img*

2. Put the two files onto an FTP or SFTP server. This example uses the CMM built-in TFTP server. Figure 6-12 shows the two files on CMM TFTP server.

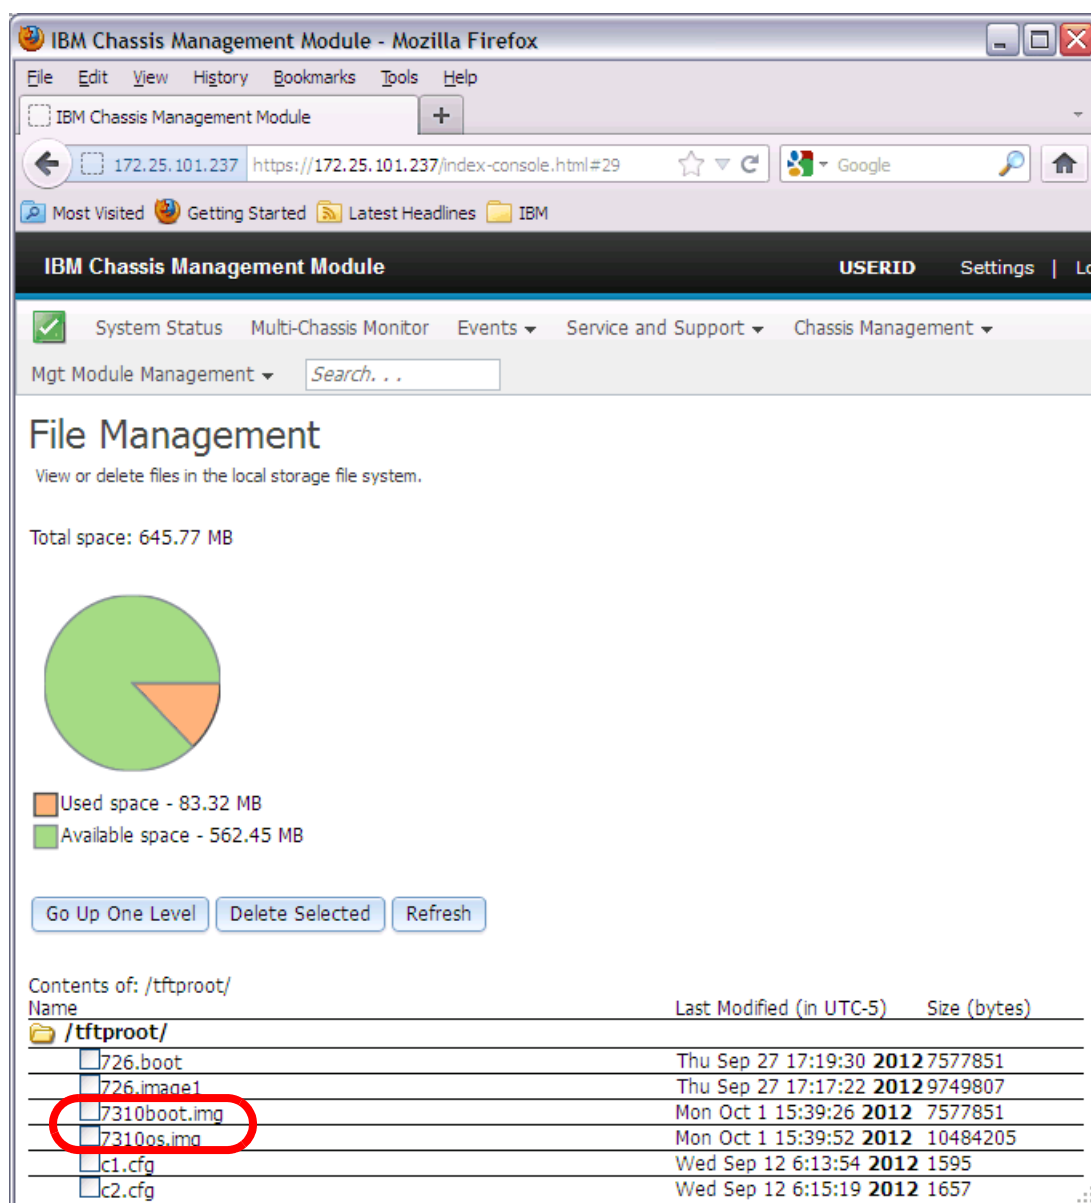


Figure 6-12 Firmware v7.3.1.0 image files on the CMM TFTP server

3. Download the image files to EN4093. First, log in to EN4093 as administrator, as shown in Example 6-15. When prompted to select CLI mode, choose **isccli**.

Example 6-15 Logging in to EN4093

```
login as: admin
Using keyboard-interactive authentication.
Enter password:
```

```
IBM Flex System Fabric EN4093 10Gb Scalable Switch.
```

```
Select Command Line Interface mode (ibmnos-cli/isccli): isccli
System Information at 14:41:22 Mon Oct 1, 2012
Time zone: America/US/Pacific
Daylight Savings Time Status: Disabled
```

```
IBM Flex System Fabric EN4093 10Gb Scalable Switch
```

```
Switch has been up for 2 days, 23 hours, 22 minutes and 43 seconds.
Last boot: 15:20:45 Fri Sep 28, 2012 (reset from Telnet/SSH)
```

```
MAC address: 6c:ae:8b:bf:fe:00   IP (If 10) address: 10.10.10.239
Internal Management Port MAC Address: 6c:ae:8b:bf:fe:ef
Internal Management Port IP Address (if 128): 172.25.101.239
External Management Port MAC Address: 6c:ae:8b:bf:fe:fe
External Management Port IP Address (if 127):
Software Version 7.2.2.2 (FLASH image1), active configuration.
```

```
Hardware Part Number      : 49Y4272
Hardware Revision         : 02
Serial Number             : Y250VT24M123
Manufacturing Date (WWYY) : 1712
PCBA Part Number          : BAC-00072-01
PCBA Revision             : 0
PCBA Number               : 00
Board Revision            : 02
PLD Firmware Version      : 1.5
```

```
Temperature Warning       : 29 C (Warn at 60 C/Recover at 55 C)
Temperature Shutdown      : 30 C (Shutdown at 65 C/Recover at 60 C)
Temperature Inlet         : 24 C
Temperature Exhaust       : 30 C
```

```
Power Consumption        : 43.530 W (12.184 V, 3.572 A)
```

```
Switch is in I/O Module Bay 4
```

4. Enable privileged EXEC mode (command **enable**) and download the boot image file. As shown in Example 6-16, use the command **copy tftp boot-image** to download the boot image file.

Example 6-16 Enabling privileged EXEC mode and downloading boot image

```
EN4093flex_2>enable
```

```
Enable privilege granted.
EN4093flex_2#copy tftp boot-image
```

```

Port type ["DATA"/"MGT"/"EXTM"]: MGT
Address or name of remote host: 172.25.101.237
Source file name: 7310boot.img

boot kernel currently contains Software Version 7.2.2.2
New download will replace boot kernel with file "7310boot.img"
  from FTP/TFTP server 172.25.101.237.
Connecting via MGT port.
Confirm download operation (y/n) ? y
Starting download...

File appears valid
Download in progress
.....
.....
.....
.....
.....
Boot image (FS, 7577851 bytes) download complete.
Writing to flash...This can take up to 90 seconds. Please wait
FS Sector now contains Software Version 7.3.1
Boot image (Kernel, 7577851 bytes) download complete.
Writing to flash...This can take up to 90 seconds. Please wait
Kernel Sector now contains Software Version 7.3.1
Boot image (DFT, 7577851 bytes) download complete.
Writing to flash...This can take up to 90 seconds. Please wait
DFT Sector now contains Software Version 7.3.1
Boot image (Boot, 7577851 bytes) download complete.
Writing to flash...This can take up to 90 seconds. Please wait
Boot Sector now contains Software Version 7.3.1

```

5. Download the OS image file into image2 and set switch to boot from image2 with the command **copy tftp image2** as displayed in Example 6-17.

Example 6-17 Downloading the OS image file

```

EN4093flex_2#copy tftp image2
Port type ["DATA"/"MGT"/"EXTM"]: MGT
Address or name of remote host: 172.25.101.237
Source file name: 7310os.img

image2 currently contains Software Version 7.2.2.2
  that was downloaded at 6:57:31 Mon Jun 18, 2012.
New download will replace image2 with file "7310os.img"
  from FTP/TFTP server 172.25.101.237.
Connecting via MGT port.
Confirm download operation (y/n) ? y
Starting download...

File appears valid
Download in progress
.....
.....
.....
.....
.....

```

```

.....
Image download complete (10484205 bytes)
Writing to flash...This takes about 10 seconds. Please wait
Write complete (10484205 bytes), now verifying FLASH...
Verification of new image2 in FLASH successful.
image2 now contains Software Version 7.3.1
Switch is currently set to boot software image1.
Do you want to change that to the new image2? [y/n]
Oct 1 14:55:05 EN4093flex_2 INFO    mgmt: image2 downloaded from host

172.25.101.237, file '7310os.img', software version 7.3.1
y
Next boot will use new software image2.

```

6. Reboot the switch to activate the new code as shown in Example 6-18.

Example 6-18 Reboot the switch

```

EN4093flex_2#reload

Reset will use software "image2" and the active config block.
>> Note that this will RESTART the Spanning Tree,
>> which will likely cause an interruption in network service.
Confirm reload (y/n) ? y

```

7. When the switch reloads, use command **show boot** to verify that the new firmware 7.3.1.0 is installed and running as shown in Example 6-19.

Example 6-19 New firmware verification

```

EN4093flex_2#show boot
Currently set to boot software image2, active config block.
NetBoot: disabled, NetBoot tftp server: , NetBoot cfgfile:
Current CLI mode set to IBMNOS-CLI with selectable prompt enabled.
Current FLASH software:
  image1: version 7.2.2.2, downloaded 14:55:26 Mon Jun 18, 2012
  image2: version 7.3.1, downloaded 22:55:05 Mon Oct 1, 2012
  boot kernel: version 7.3.1
Currently scheduled reboot time: none

```

6.3.3 Recovering from a failed firmware upgrade

Although it is unlikely, the firmware upgrade process might fail. If this situation occurs, you can still recover the EN4093 switch. To do so, complete the following steps:

1. Connect a PC running a terminal emulation utility to the serial port of your switch while the switch is off. Then, access the switch as described in the User's Guide. Use the following communication parameters to establish terminal emulation session:
 - Speed: 9600 bps
 - Data Bits: 8
 - Stop Bits: 1
 - Parity: None
 - Flow Control: None

Important: The procedure that is described in this section might also be useful when you boot the switch, and the boot and OS versions are not equal.

2. Power on the switch. From your terminal window, press Shift + B while the memory tests are processing and dots are showing the progress. A menu opens as shown in Example 6-20.

Example 6-20 Boot management menu

```
Boot Management Menu
  1 - Change booting image
  2 - Change configuration block
  3 - Boot in recovery mode (tftp and xmodem download of images to
recovery switch)
  4 - Xmodem download (for boot image only - use recovery mode for
application images)
  5 - Reboot
  6 - Exit
Please choose your menu option:
```

3. Select 4 for Xmodem download of boot image. Change the serial connection speed as follows:

Switch baudrate to 115200 bps and press ENTER ...

Change the settings of your terminal to meet the 115200 bps requirement and press Enter.
4. The system switches to download accept mode. You see a series of C characters on the panel that prompt you when the switch is ready. Start an Xmodem terminal to push the boot code you want to restore into the switch.
5. Select the boot code for your system, and the switch starts the download. You should see a panel similar to Example 6-21.

Example 6-21 Xmodem boot image download

```
xyzModem - CRC mode, 62106(SOH)/0(STX)/0(CAN) packets, 3 retries

Extracting images ... Do *NOT* power cycle the switch.
**** RAMDISK ****
Un-Protected 33 sectors
Erasing Flash...
..... done
Erased 33 sectors
Writing to Flash...9....8....7....6....5....4....3....2....1....0done
Protected 33 sectors
**** KERNEL ****
Un-Protected 25 sectors
Erasing Flash...
..... done
Erased 25 sectors
Writing to Flash...9....8....7....6....5....4....3....2....1....done
Protected 25 sectors
**** DEVICE TREE ****
Un-Protected 1 sectors
Erasing Flash...
. done
```

```

Erased 1 sectors
Writing to Flash...9....8....7....6....5....4....3....2....1....done
Protected 1 sectors
**** BOOT CODE ****
Un-Protected 4 sectors
Erasing Flash...
.... done
Erased 4 sectors
Writing to Flash...9....8....7....6....5....4....3....2....1....done
Protected 4 sectors

```

6. When this process is finished, you are prompted to reconfigure your terminal to 9600 bps speed:

Change the baud rate back to 9600 bps, hit the <ESC> key

Change the speed of your serial connection, and then press Esc.

7. The Boot Management Menu opens again. Select option 3 now, and change the speed to 115000 bps when the following message appears to start pushing the OS image.

```
## Switch baudrate to 115200 bps and press ENTER ...
```

When speed is changed to 115200 bps, press Enter to continue download.

8. Select the OS image that you want to upload to the switch. The Xmodem client starts sending the image to the switch. When the upload is complete, you see a panel similar to the one in Example 6-22.

Example 6-22 OS image upgrade

```

xyzModem - CRC mode, 27186(SOH)/0(STX)/0(CAN) packets, 6 retries
Extracting images ... Do *NOT* power cycle the switch.
**** Switch OS ****
Please choose the Switch OS Image to upgrade [1|2|n] :

```

9. You are prompted to select the image space in the switch you want to upgrade. After you select the OS image bank, you see a panel similar to the one in Example 6-23.

Example 6-23 Upgrading the OS image

```

Switch OS Image 1 ...
Un-Protected 27 sectors
Erasing Flash..... done
Writing to Flash.....done
Protected 27 sectors

```

10. When this process is done, you are prompted to reconfigure your terminal to 9600 bps speed again:

Change the baud rate back to 9600 bps, hit the <ESC> key

Press **Esc** to show the Boot Management Menu, and choose option 6 to exit and boot the new image.

6.4 Logging and reporting

This section addresses the following topics:

- ▶ Managing and configuring system logs
- ▶ Configuring an SNMP agent and SNMP traps
- ▶ Remote monitoring
- ▶ sFlow

6.4.1 System logs

IBM Networking OS can provide valuable maintenance and troubleshooting information through a system log (syslog) that uses the following fields in log entries: Date, time, switch name, criticality level, and message.

You can view the latest system logs by running the **show logging messages** command as shown in Example 6-24.

Example 6-24 Example of syslog output

```
Oct 17 22:30:47 en4093flex_1 NOTICE mgmt: admin(admin) login from host
10.10.53.121
Oct 17 22:30:53 en4093flex_1 INFO mgmt: new configuration saved from ISCLI
Oct 17 22:32:27 en4093flex_1 INFO telnet/ssh-1: Current config successfully
tftp'd to 10.10.53.121:en4093flex_1-OSPF
Oct 17 22:32:29 en4093flex_1 NOTICE mgmt: admin(admin) connection closed from
Telnet/SSH
Oct 17 22:35:16 en4093flex_1 NOTICE ntp: System clock updated
Oct 17 22:49:06 en4093flex_1 NOTICE mgmt: USERID(Admin) login from BBI.
Oct 17 22:50:16 en4093flex_1 NOTICE ntp: System clock updated
Oct 17 23:25:08 en4093flex_1 NOTICE mgmt: USERID(Admin) logout from BBI.
Oct 17 23:35:23 en4093flex_1 NOTICE ntp: System clock updated
Oct 17 23:45:18 en4093flex_1 NOTICE mgmt: admin(admin) login from host
10.10.53.121
Oct 17 23:45:45 en4093flex_1 ALERT vlag: vLAG on portchannel 1 is up
Oct 17 23:45:46 en4093flex_1 ALERT vlag: vLAG on portchannel 15 is up
Oct 17 23:46:26 en4093flex_1 INFO cfgchg: Configured from SSHv2 by admin on
host 10.10.53.121
```

Each syslog message has a criticality level associated with it, included in text form as a prefix to the log message. One of eight different prefixes is used, depending on the condition that the administrator is being notified of:

- ▶ Level 0 - EMERG: Indicates that the system is unusable.
- ▶ Level 1 - ALERT: Indicates that action should be taken immediately.
- ▶ Level 2 - CRIT: Indicates critical conditions.
- ▶ Level 3 - ERR: Indicates error conditions or operations in error.
- ▶ Level 4 - WARNING: Indicates warning conditions.
- ▶ Level 5 - NOTICE: Indicates a normal but significant condition.
- ▶ Level 6 - INFO: Indicates an information message.
- ▶ Level 7 - DEBUG: Indicates a debug-level message.

Information logged

You can selectively choose what information is logged by Syslog. You have a number of options:

all	All
bgp	BGP
cfg	Configuration
cli	Command-line interface
console	Console
dcbx	DCB Capability Exchange
difftrak	Configuration difference tracking
failover	Failover
fcoe	Fibre Channel over Ethernet
hotlinks	Hot Links
ip	Internet protocol
ipv6	IPv6
lacp	Link Aggregation Control Protocol
link	System port link
lldp	LLDP
management	Management
mld	MLD
netconf	NETCONF Configuration Protocol
ntp	Network time protocol
ospf	OSPF
ospfv3	OSPFv3
rmon	Remote monitoring
server	Syslog server
spanning-tree-group	Spanning Tree Group
ssh	Secure Shell
system	System
vlag	Virtual Link Aggregation
vlan	VLAN
vm	Virtual machine
vnic	VNIC
vrrp	Virtual Router Redundancy Protocol
web	Web

Use the following ISCLI command syntax:

```
[no] logging log [<feature>]
```

For example, the following command enables syslog messages generation for SSH:

```
logging log ssh
```

The following command disables syslog messages generation for LACP:

```
no logging log lacp
```

The following command displays a list of features for which syslog messages are generated:

```
show logging
```

Logging destinations

You can set up to two destinations for reporting. A destination of 0.0.0.0 means logs are stored locally on the switch. Another instance of a log destination host can be a remote logging server. In this case, the logs are sent to the server through Syslog. For each of the

two destinations, you can define many parameters, including the severity of logs to be sent to that particular destination.

Example 6-25 shows a configuration set to log locally the messages with ALERT (Level 1) severity and to send all critical (severity CRIT, Level 2) events to 172.25.101.200.

Example 6-25 Example of Syslog configuration

```
en4093flex_1(config)#logging host 1 address 0.0.0.0
en4093flex_1(config)#logging host 1 severity 1
en4093flex_1(config)#logging host 2 address 172.25.101.200

Oct 18 0:54:32 en4093flex_1 NOTICE mgmt: second syslog host changed to
172.25.101.200 via MGT port
en4093flex_1(config)#logging host 2 severity 2
```

You can also use the **logging host** command to specify the interface used for logging. The command has these options:

- ▶ data-port
- ▶ extm-port
- ▶ mgt-port

For example, to send the logs to a second destination from a data port, run the command shown in Example 6-26.

Example 6-26 Changing the logging interface

```
en4093flex_1(config)#logging host 2 data-port

Oct 18 0:57:13 en4093flex_1 NOTICE mgmt: second syslog host changed to 0.0.0.0
via Data port
```

Logging console

To make logging output visible on the console, run **logging console**. You can select the severity level of messages to be logged with the following syntax:

```
logging console severity <0-7>
```

6.4.2 SNMP

IBM Networking OS provides Simple Network Management Protocol (SNMP) version 1, version 2, and version 3 support for access through any network management software, such as IBM Systems Director. The default SNMP version support is SNMPv3 only.

Important: SNMP read and write functions are enabled by default. If SNMP is not needed for your network, disable these functions before you connect the switch to the network.

SNMP versions 1 and 2

To access the SNMP agent on the EN4093, configure the read and write community strings on the SNMP manager to match the community strings on the switch. The default read community string on the switch is public, and the default write community string is private.

The read and write community strings on the switch can be changed by running the following commands:

```
en4093flex_1(config)# snmp-server read-community <1-32 characters>
en4093flex_1(config)# snmp-server write-community <1-32 characters>
```

The SNMP manager must be able to reach the management interface or any of the IP interfaces on the switch.

For the SNMP manager to receive the SNMPv1 traps sent out by the SNMP agent on the switch, configure the trap host on the switch by running the following command:

```
en4093flex_1(config)# snmp-server trap-src-if <trap source IP interface>
en4093flex_1(config)# snmp-server host <IPv4 address> <trap host community string>
```

SNMP version 3

SNMP version 3 (SNMPv3) is an enhanced version of the Simple Network Management Protocol that was approved by the Internet Engineering Steering Group in March 2002. SNMPv3 contains more security and authentication features that provide data origin authentication, data integrity checks, timeliness indicators, and encryption to protect against threats such as masquerade, modification of information, message stream modification, and disclosure.

Using SNMPv3, your clients can query the MIBs securely.

Default configuration

IBM Networking OS has two SNMPv3 users by default. Both of the following users have access to all the MIBs supported by the switch:

- ▶ User 1 name is adminmd5 (password adminmd5). The authentication used is MD5.
- ▶ User 2 name is adminsha (password adminsha). The authentication used is SHA.

Up to 16 SNMP users can be configured on the switch. To modify an SNMP user, run the following command:

```
en4093flex_1(config)# snmp-server user <1-16> name <1-32 characters>
```

Users can be configured to use the authentication and privacy options. The EN4093 switch supports two authentication algorithms, MD5 and SHA, as specified in the following command:

```
en4093flex_1(config)# snmp-server user <1-16> authentication-protocol
{md5|sha} authentication-password
```

User configuration example

To configure a user, complete the following steps:

1. To configure a user with the name admin, the authentication type MD5, the authentication password of admin, and the privacy option DES with a privacy password of admin, run the commands shown in Example 6-27.

Example 6-27 SNMP v3 user configuration example

```
en4093flex_1(config)# snmp-server user 5 name admin
en4093flex_1(config)# snmp-server user 5 authentication-protocol md5
authentication-password
Changing authentication password; validation required:
Enter current admin password: <admin. password>
Enter new authentication password: <auth. password>
```

```
Re-enter new authentication password: <auth. password>
New authentication password accepted.
en4093flex_1(config)# snmp-server user 5 privacy-protocol des
privacy-password
Changing privacy password; validation required:
Enter current admin password: <admin. password>
Enter new privacy password: <privacy password>
Re-enter new privacy password: <privacy password>
New privacy password accepted.
```

2. Configure a user access group, along with the views the group can access, by running the commands shown in Example 6-28. Use the access table to configure the group's access level.

Example 6-28 SNMPv3 group and view configuration example

```
en4093flex_1(config)# snmp-server access 5 name admingrp
en4093flex_1(config)# snmp-server access 5 level authpriv
en4093flex_1(config)# snmp-server access 5 read-view iso
en4093flex_1(config)# snmp-server access 5 write-view iso
en4093flex_1(config)# snmp-server access 5 notify-view iso
```

Because the read view, write view, and notify view are all set to iso, the user type has access to all private and public MIBs.

3. Assign the user to the user group by running the commands shown in Example 6-29. Use the group table to link the user to a particular access group.

Example 6-29 SNMPv3 user assignment configuration

```
en4093flex_1(config)# snmp-server group 5 user-name admin
en4093flex_1(config)# snmp-server group 5 group-name admingrp
```

Configuring SNMP traps

This section describes the steps for configuring the SNMP traps.

SNMPv2 trap configuration

To configure the SNMPv2 trap, complete the following steps:

1. Configure a user with no authentication and password, as shown in Example 6-30.

Example 6-30 SNMP user configuration example

```
en4093flex_1(config)#snmp-server user 10 name v2trap
```

2. Configure an access group and group table entries for the user. Use the menu that is shown in Example 6-31 to specify which traps can be received by the user.

Example 6-31 SNMP group configuration

```
en4093flex_1(config)#snmp-server group 10 security snmpv2
en4093flex_1(config)#snmp-server group 10 user-name v2trap
en4093flex_1(config)#snmp-server group 10 group-name v2trap
en4093flex_1(config)#snmp-server access 10 name v2trap
en4093flex_1(config)#snmp-server access 10 security snmpv2
en4093flex_1(config)#snmp-server access 10 notify-view iso
```

3. Configure an entry in the notify table as shown in Example 6-32.

Example 6-32 SNMP notify entry configuration

```
en4093flex_1(config)#snmp-server notify 10 name v2trap
en4093flex_1(config)#snmp-server notify 10 tag v2trap
```

4. Specify the IPv4 address and other trap parameters in the targetAddr and targetParam tables. Use the commands that are shown in Example 6-33 to specify the user name that is associated with the targetParam table.

Example 6-33 SNMP trap destination and trap parameters configuration

```
en4093flex_1(config)#snmp-server target-address 10 name v2trap address
100.10.2.1
en4093flex_1(config)#snmp-server target-address 10 taglist v2trap
en4093flex_1(config)#snmp-server target-address 10 parameters-name v2param
en4093flex_1(config)#snmp-server target-parameters 10 name v2param
en4093flex_1(config)#snmp-server target-parameters 10 message snmpv2c
en4093flex_1(config)#snmp-server target-parameters 10 user-name v2trap
en4093flex_1(config)#snmp-server target-parameters 10 security snmpv2
```

5. Use the community table to specify which community string is used in the trap, as shown in Example 6-34.

Example 6-34 SNMP community configuration

```
en4093flex_1(config)#snmp-server community 10 index v2trap
en4093flex_1(config)#snmp-server community 10 user-name v2trap
```

SNMPv3 trap configuration

To configure a user for SNMPv3 traps, you can send the traps with both privacy and authentication, with authentication only, or without privacy or authentication.

You can configure these settings in the access table by running the following commands:

```
en4093flex_1(config)#snmp-server access <1-32> level
en4093flex_1(config)#snmp-server target-parameters <1-16>
```

Configure the user in the user table.

It is not necessary to configure the community table for SNMPv3 traps because the community string is not used by SNMPv3.

Example 6-35 shows how to configure a SNMPv3 user v3trap with authentication only:

Example 6-35 SNMPv3 trap configuration

```
en4093flex_1(config)#snmp-server user 11 name v3trap
en4093flex_1(config)#snmp-server user 11 authentication-protocol md5
authentication-password
Changing authentication password; validation required:
Enter current admin password: <admin. password>
Enter new authentication password: <auth. password>
Re-enter new authentication password: <auth. password>
New authentication password accepted.
en4093flex_1(config)#snmp-server access 11 notify-view iso
en4093flex_1(config)#snmp-server access 11 level authnopriv
en4093flex_1(config)#snmp-server group 11 user-name v3trap
```

```
en4093flex_1(config)#snmp-server group 11 tag v3trap
en4093flex_1(config)#snmp-server notify 11 name v3trap
en4093flex_1(config)#snmp-server notify 11 tag v3trap
en4093flex_1(config)#snmp-server target-address 11 name v3trap address
172.25.101.200
en4093flex_1(config)#snmp-server target-address 11 taglist v3trap
en4093flex_1(config)#snmp-server target-address 11 parameters-name v3param
en4093flex_1(config)#snmp-server target-parameters 11 name v3param
en4093flex_1(config)#snmp-server target-parameters 11 user-name v3trap
en4093flex_1(config)#snmp-server target-parameters 11 level authNoPriv
```

6.4.3 Remote Monitoring (RMON)

The IBM switches provide a Remote Monitoring (RMON) interface that allows network devices to exchange network monitoring data. RMON allows the switch to perform the following functions:

- ▶ Track events and trigger alarms when a threshold is reached.
- ▶ Notify administrators by issuing a syslog message or SNMP trap.

The RMON MIB provides an interface between the RMON agent on the switch and an RMON management application. The RMON MIB is described in RFC 1757:

<http://www.ietf.org/rfc/rfc1757.txt>

The RMON standard defines objects that are suitable for the management of Ethernet networks. The RMON agent continuously collects statistics and proactively monitors switch performance. You can use RMON to monitor traffic that flows through the switch.

The switch supports the following RMON Groups, as described in RFC 1757:

- ▶ Group 1: Statistics
- ▶ Group 2: History
- ▶ Group 3: Alarms
- ▶ Group 9: Events

RMON Group 1: Statistics

The switch supports collection of Ethernet statistics as outlined in the RMON statistics MIB, referring to etherStatsTable. You can configure RMON statistics on a per-port basis. RMON statistics are sampled every second, and new data overwrites any old data on a port.

Important: RMON port statistics must be enabled for the port before you can view them.

Here is an example configuration:

1. Enable RMON on a port. To enable RMON on a port, run **interface** and **rmon**:
 - en4093flex_1(config)# interface port 1
 - en4093flex_1(config-if)# rmon
2. To view the RMON statistics, run **interface**, run **rmon**, and run **show** to show the interface, as shown in Example 6-36.

Example 6-36 Viewing the RMON statistics

```
en4093flex_1(config)# interface port INTA1
en4093flex_1(config-if)# rmon
```

```
en4093flex_1(config-if)# show interface port INTA1 rmon-counters
```

```
-----  
RMON statistics for port INTA1:  
etherStatsDropEvents: NA  
etherStatsOctets: 7305626  
etherStatsPkts: 48686  
etherStatsBroadcastPkts: 4380  
etherStatsMulticastPkts: 6612  
etherStatsCRCAlignErrors: 0  
etherStatsUndersizePkts: 0  
etherStatsOversizePkts: 0  
etherStatsFragments: 2  
etherStatsJabbers: 0  
etherStatsCollisions: 0  
etherStatsPkts64Octets: 27445  
etherStatsPkts65to127Octets: 12253  
etherStatsPkts128to255Octets: 1046  
etherStatsPkts256to511Octets: 619  
etherStatsPkts512to1023Octets: 7283  
etherStatsPkts1024to15180Octets: 38  
-----
```

RMON Group 2: History

You can use the RMON History Group to sample and archive Ethernet statistics for a specific interface during a specific time interval. History sampling is done per port.

Important: RMON port statistics must be enabled for the port before an RMON History Group can monitor the port.

Data is stored in buckets, which store data gathered during discreet sampling intervals. At each configured interval, the History index takes a sample of the current Ethernet statistics, and places them into a bucket. History data buckets are in dynamic memory. When the switch is rebooted, the buckets are emptied.

Requested buckets are the number of buckets, or data slots, requested by the user for each History Group. Granted buckets are the number of buckets that are granted by the system, based on the amount of system memory available. The system grants a maximum of 50 buckets.

You can use an SNMP browser to view History samples.

The type of data that can be sampled must be of an Index object type, as described in RFC 1213 and RFC 1573:

<http://www.ietf.org/rfc/rfc1213.txt>
<http://www.ietf.org/rfc/rfc1573.txt>

The most common data type for the History sample is as follows:

1.3.6.1.2.1.2.2.1.1.<x>

The last digit (x) represents the number of the port to monitor.

6.4.4 Using sFlow to monitor traffic

IBM System Networking switches support sFlow technology for monitoring traffic in data networks. The switch includes an embedded sFlow agent that can be configured to provide continuous monitoring information of IPv4 traffic to a central sFlow analyzer.

The switch is responsible only for forwarding sFlow information. A separate sFlow analyzer is required elsewhere in the network to interpret sFlow data.

Use the following commands to enable and configure sFlow:

- ▶ Enable sFlow on the switch:
`sflow enable`
- ▶ Set sFlow analyzer IP address:
`sflow server <IP address>`
- ▶ Optionally, set the UDP port for sFlow analyzer (default is 6343):
`sflow port <1-65535>`
- ▶ Display sFlow configuration settings:
`show sflow`

sFlow statistical counters

IBM System Networking switch can be configured to send network statistics to an sFlow analyzer at regular intervals. For each port, a polling interval of 5 - 60 seconds can be configured, or 0 (the default) can be set to disable this feature.

Use the following command to set the sFlow port polling interval:

```
sflow polling <5-60>
```

When polling is enabled, at the end of each configured polling interval, the switch reports general port statistics and port Ethernet statistics.

sFlow network sampling

In addition to statistical counters, IBM System Networking switches can be configured to collect periodic samples of the traffic data received on each port. For each sample, 128 bytes are copied, UDP-encapsulated, and sent to the configured sFlow analyzer.

For each port, the sFlow sampling rate can be configured to occur every 256 - 65536 packets, or set to 0 (the default) to disable this feature. A sampling rate of 256 means that one sample is taken for approximately every 256 packets that are received on the port. The sampling rate is statistical, however. It is possible to have more or fewer samples sent to the analyzer for any specific group of packets (especially under low traffic conditions). The actual sample rate becomes most accurate over time, and under higher traffic flow.

Use the following command to set the sFlow port sampling rate:

```
sflow sampling <256-65536>
```


sFlow sampling has the following restrictions:

- ▶ **Sample rate:** The fastest sFlow sample rate is 1 out of every 256 packets.
- ▶ **ACLs:** sFlow sampling is done before ACLs are processed. For ports configured with both sFlow sampling and one or more ACLs, sampling occurs regardless of the action of the ACL.
- ▶ **Port mirroring:** sFlow sampling does not occur on mirrored traffic. If sFlow sampling is enabled on a port that is configured as a port monitor, the mirrored traffic is not sampled.

sFlow sampling: Although sFlow sampling is not generally a processor-intensive operation, configuring fast sampling rates (such as once every 256 packets) on ports under heavy traffic loads can cause switch processor utilization to reach maximum. Use larger rate values for ports that experience heavy traffic.



Cisco IOS to IBM isCLI Command Comparison

IBM switches offer two different command-line interfaces, IBMNOS-CLI, and IBM isCLI. The industry standard CLI is designed to be familiar to network professionals accustomed to Cisco's IOS CLI. This appendix shows a command comparison between Cisco's IOS and IBM isCLI. Common commands are illustrated here to aid you in implementing IBM System Networking products.

This appendix includes the following sections:

- ▶ General configuration
- ▶ Authentication
- ▶ BPDU Guard
- ▶ DHCP snooping
- ▶ Hostname and DNS server configuration
- ▶ Banner configuration
- ▶ Interface speed and duplex
- ▶ LLDP
- ▶ Management network configuration
- ▶ NTP
- ▶ OSPF configuration
- ▶ Port mirroring
- ▶ SNMP
- ▶ Spanning Tree Protocol (STP)
- ▶ SSH and Telnet
- ▶ Syslog
- ▶ Port aggregation (static)
- ▶ Port aggregation (LACP)
- ▶ VLAN tagging (802.1q)

General configuration

Many commands in IBM isCLI are the same or similar to Cisco IOS commands. See Table A-1.

Table A-1 Basic isCLI commands

Command	Purpose
switch>enable	Enter Privilege Exec mode
switch#configure terminal	Enter Configuration mode
switch(config)#exit	Exit configuration
switch#copy running-config startup-config	Save configuration with verification
switch#write	Save configuration without verification
switch#show running-config	Display current running configuration
switch#show vlan	Show configured VLANs and assigned ports
switch#show interface status	Show status of all ports

Authentication

This section lists commands used for both local and remote authentication.

Local authentication

The following commands are used for local authentication.

Cisco IOS

The mechanism for creating a new user and enable password are shown in Example A-1.

Example: A-1 Cisco IOS username configuration

```
configure terminal
username <username> secret <username_secret>
enable secret <enable_secret>
```

IBM isCLI

There are three usernames that are defined on the system as factory default (Table A-2).

Table A-2 IBM ISCLI factory default usernames

User	Factory default state
user	enabled, offline
oper	disable, offline
admin	always enabled, online

The admin account cannot be disabled. The procedure to change the admin password and to create new users is outlined in Example A-2.

Example: A-2 Changing the admin password and creating a new user account

```
switch(config)#access user administrator-password
Changing ADMINISTRATOR password; validation required:
Enter current local admin password: <old admin password>
Enter new admin password (max 128 characters): <new admin password>
Re-enter new admin password:<new admin password>
New admin password accepted.
access user user-password

switch(config)#access user <1-10> name <username>
switch(config)#access user <1-10> password
Changing privlusr password; validation required:
Enter current admin password:<admin password>
Enter new <username> password (max 128 characters):<password>
Re-enter new privlusr password:<password>
New <username> password accepted.

switch(config)#access user <1-10> level <user/operator/administrator>
switch(config)#access user <1-10> enable
```

Remote authentication

TACACS+ is a commonly used authentication protocol by network engineers.

Cisco IOS

Example A-3 shows a common implementation in Cisco IOS.

Example: A-3 Cisco External Authentication commands (Cisco IOS)

```
configure terminal
aaa authentication login default group tacacs+ local
aaa authentication enable default group tacacs+ enable
aaa authorization commands 15 default group tacacs+ if-authenticated
aaa accounting exec default start-stop group tacacs+
aaa accounting commands 15 default start-stop group tacacs+
aaa accounting network default start-stop group tacacs+
tacacs-server host <host_ip>
tacacs-server key <key>
```

IBM isCLI

Similar functionality can be implemented in IBM isCLI as shown in Example A-4.

Example: A-4 IBM External Authentication commands (IBM isCLI)

```
configure terminal
tacacs-server primary-host <host-ip> key <key>
tacacs-server enable-bypass
tacacs-server enable
```

BPDU Guard

BPDU Guard is often implemented to prevent general users from plugging in management network equipment into the overall network infrastructure.

Cisco IOS

Example A-5 shows a common BPDU Guard implementation in Cisco IOS.

Example: A-5 Cisco BPDU Guard commands

```
(config)#interface GigabitEthernet0/1
(config-if)#spanning-tree bpduguard enable
```

IBM isCLI

Similar functionality can be implemented in IBM isCLI as shown in Example A-6.

Example: A-6 IBM BPDU Guard commands

```
(config)#interface port 1
(config-if)#bpdu-guard
```

DHCP snooping

DHCP snooping is a DHCP security feature that provides network security by filtering untrusted DHCP messages and by building and maintaining a DHCP snooping binding database. The database is also referred to as a DHCP snooping binding table. The DHCP snooping binding table contains the MAC address, IP address, lease time, binding type, VLAN number, and port number that correspond to the local untrusted interface on the switch. By default, DHCP snooping is disabled on all VLANs. You can enable DHCP snooping on one or more VLANs. You must enable DHCP snooping globally. To use DHCP snooping, you must configure the DHCP server interface as trusted.

Cisco IOS

Configuration for Cisco IOS DHCP snooping is outlined in Example A-7

Example: A-7 Enabling DHCP Snooping in Cisco IOS, globally and per VLAN

```
(config)#ip dhcp snooping
(config)#ip dhcp snooping vlan <vlan>
(config)#interface <interface>
(config-if)#ip dhcp snooping trust
```

IBM isCLI

Configuration in IBM isCLI is the same as in Cisco IOS as shown in Example A-8.

Example: A-8 Enabling DHCP snooping in IBM isCLI

```
(config)#ip dhcp snooping
(config)#ip dhcp snooping vlan <vlan>
(config)#interface port <port>
(config-if)#ip dhcp snooping trust
```

Hostname and DNS server configuration

The following commands are used to configure the hostname and DNS server details.

Cisco IOS

Configure hostname and DNS server details for Cisco IOS as shown in Example A-9.

Example: A-9 Hostname and DNS server configuration for Cisco IOS

```
configure terminal
hostname <hostname>
ip name-server <dns-server>
ip domain-name <host.location.company.com>
```

IBM isCLI

Configure hostname and DNS server details for IBM isCLI as shown in Example A-10.

Example: A-10 Hostname and DNS server configuration for IBM isCLI

```
configure terminal
hostname <hostname>
ip dns primary-server <dns-server>
ip dns domain-name <host.location.company.com>
```

Banner configuration

Use these commands to configure a login or motd banner.

Cisco IOS

To create a login or motd banner in Cisco IOS, use the commands shown in Example A-11.

Example: A-11 Cisco IOS banner motd configuration

```
switch(config)#banner motd ^
Enter TEXT message. End with the character '^'
You have accessed a Cisco switch
contact the network admin if you require access details
```

```
etc.  
^  
switch(config)#
```

IBM isCLI

Configure a login notice in IBM isCLI with the commands shown in Example A-12.

Example: A-12 IBM isCLI system notice configuration, use the addline option to add a line

```
switch(config)#system notice  
Enter new login notice line by line (enter single '.' to end) :  
>>You have accessed an IBM System Neetworking switch  
>>contact the network admin if you require access details  
>>etc.  
>>.  
switch(config)#system notice addline <add some banner test>
```

Interface speed and duplex

Set interface speed and duplex by using the following commands.

Cisco IOS

Example A-13 shows how to set interface speed and duplex in Cisco IOS.

Example: A-13 Configuring interface speed and duplex in Cisco IOS

```
configure terminal  
interface FastEthernet0/12  
speed 100  
duplex full
```

IBM isCLI

Example A-14 shows how to configure duplex and port speed in IBM isCLI.

Example: A-14 Configuring interface speed and duplex in IBM isCLI

```
configure terminal  
interface port 12  
speed 100  
duplex full
```

LLDP

To enable LLDP globally and on a per interface basis, use the following commands.

Cisco IOS

Example A-15 shows how to configure LLDP in Cisco IOS.

Example: A-15 LLDP configuration in Cisco IOS

```
switch(config)#lldp run
switch(config)#interface GigabitEthernet1/0/23
switch(config)#lldp transmit
switch(config)#lldp receive
```

IBM isCLI

Example A-16 shows how to configure LLDP in IBM isCLI.

Example: A-16 LLDP configuration in IBM isCLI

```
switch(config)#lldp enable
switch(config)#interface port 1
switch(config-if)#no lldp admin-status
```

Management network configuration

IBM System Networking switches have dedicated network management ports that can be configured as shown below. Example A-17 shows how to configure network management port and gateway.

Example: A-17 Configuring the network management port and gateway for IBM isCLI

```
switch(config)#interface ip 128
switch(config-if)#ip address <ip address> <mask>
switch(config-if)#enable
switch(config)#ip gateway 4 address <ip address>
switch(config)#ip gateway 4 enable
```

NTP

To configure the timezone and NTP server, use the commands described in this section.

Cisco IOS

Example A-18 shows how to configure NTP and timezone information in Cisco IOS.

Example: A-18 Timezone and NTP server configuration in Cisco IOS

```
configure terminal
clock timezone EST -5
clock summer-time EDT recurring
ntp server <ntp server>
```

IBM isCLI

Example A-19 shows how to configure NTP and timezone information in IBM isCLI. To find out the local time-zone number, it may be easier to configure this through the IBM System Networking BBI GUI interface at least the first time.

Example: A-19 Timezone and NTP server configuration in IBM isCLI.

```
configure terminal
system timezone <time-zone number>
system daylight
ntp primary-server <ntp server>
ntp enable
```

OSPF configuration

Some of the more basic OSPF commands for Cisco IOS and IBM isCLI are described in this section.

Cisco IOS

Example A-20 shows basic OSPF configuration commands in Cisco IOS.

Example: A-20 Configuration commands for OSPF in Cisco IOS

```
switch(config)#interface loopback 0
switch(config-if)#ip address <ip address> 255.255.255.255
switch(config)#router ospf <id>
switch(config-router)#router-id <loopback ip address>
switch(config-router)#area <area> range <network number> <mask>
switch(config-router)#network <network number> <OSPF wildcard> area <area>
```

IBM isCLI

Example A-21 shows basic OSPF configuration commands in IBM isCLI.

Example: A-21 Configuration commands for OSPF in IBM isCLI

```
switch(config)#interface loopback 1
switch(config)#ip address <ip address> 255.255.255.255
switch(config)#enable
switch(config)#exit
switch(config)#ip router-id <loopback ip address>
switch(config)#router ospf
switch(config-router-ospf)#enable
switch(config-router-ospf)#area 0 enable
switch(config)#interface ip <ip interface number>
switch(config-ip-if)#ip ospf enable
```

Port mirroring

Port mirroring is used to monitor network traffic. The switch sends a copy of the network packets that are seen on one port or VLAN to a network monitoring connection on a different port.

Cisco IOS

Example A-22 shows how to configure port mirroring in Cisco IOS. Cisco generally refers to port mirroring as Switched Port Analyzer (SPAN).

Example: A-22 Configuring SPAN in Cisco IOS

```
configure terminal
monitor session 1 source interface gigabitEthernet 1/1 both
monitor session 1 destination interface gigabitEthernet 1/2
```

IBM isCLI

IBM isCLI supports a mirroring model that uses a total of three monitor ports. Each of these ports can receive traffic from any number of target ports. See Example A-23.

Example: A-23 Configuring port mirroring in IBM isCLI

```
configure terminal
port-mirroring monitor-port 2 mirroring-port 1 both
```

SNMP

SNMP community strings, SNMP v3 users, SNMP views, SNMP traps, and SNMP target servers can be configured using the following examples.

Cisco IOS

Example A-24 shows how to configure various SNMP access details in Cisco IOS.

Example: A-24 Cisco SNMP configuration

```
configure terminal
snmp-server location <location>
snmp-server contact <contact>
snmp-server community <community_string> <RO/RW> <acl>
snmp-server host <ip-address> <community-string>
snmp-server view <view> <MIB> <include/exclude>
snmp-server group <group_name> v3 <auth/noauth/priv> read <view> write <view>
access <acl>
snmp-server user <user> <group> v3 auth <md5/sha> <authentication password> access
<acl>
snmp-server enable traps <snmp_traps>
```

IBM isCLI

The IBM isCLI SNMP server has a three default SNMPv3 users, two SNMPv3 Groups and five SNMPv3 views enabled by default. Create news users with a USM user table index from <4-16>. To disable v1v2 only users, choose to allow only SNMPv3 requests. Example A-25 shows the output of **show snmp-server** with default SNMPv3 user settings.

Example: A-25 Showing snmp-server output in IBM isCLI

```
switch#show snmp-server
Current SNMP params:
  sysName:      "switch"
  sysLocation:  "SysLocation"
  sysContact:   "sysadmin@mars.ibm.com"
  Read community string: "public"
  Write community string: "private"
  SNMP state machine timeout: 5 minutes
  Trap source address: 0.0.0.0
  SNMP Trap source loopback interface not set
  Authentication traps enabled.
  All link up/down traps enabled.

Current SNMP trap hosts:

Current v1/v2 access enabled

Current SNMPv3 USM user settings:
  1: name adminmd5, auth md5, privacy des
  2: name adminsha, auth sha, privacy des
  3: name v1v2only, auth none, privacy none

Current SNMPv3 vacmAccess settings:
  1: group name admingrp, model usm
    level authPriv,
    read view iso, write view iso, notify view iso
  2: group name v1v2grp, model snmpv1
    level noAuthNoPriv,
    read view iso, write view iso, notify view v1v2only

Current SNMPv3 vacmSecurityToGroup settings:
  1: model usm, user name adminmd5, group name admingrp
  2: model usm, user name adminsha, group name admingrp
  3: model snmpv1, user name v1v2only, group name v1v2grp

Current SNMPv3 vacmViewTreeFamily settings:
  1: name v1v2only, subtree 1
    type included
  2: name v1v2only, subtree 1.3.6.1.6.3.15
    type excluded
  3: name v1v2only, subtree 1.3.6.1.6.3.16
    type excluded
  4: name v1v2only, subtree 1.3.6.1.6.3.18
    type excluded
  5: name iso, subtree 1
    type included
```

To configure parameters for the SNMP server in IBM isCLI, see Example A-26.

Example: A-26 The commands used to configure SNMP server in IBM isCLI

```
configure terminal
snmp-server location <location>
snmp-server contact <contact>
snmp-server read-community <community_string>
snmp-server write-community <community_string>
snmp-server host <ip-address> <community_string>
snmp-server user 4 name <name>
snmp-server user 4 authentication-protocol <md5/none/sha> authentication-password
<auth_password>
snmp-server group 3 group-name <name>
snmp-server access 4 level <authPriv/authNoPriv/noAuthNopPriv>
snmp-server access 4 read-view <view>
snmp-server access 4 write-view <view>
snmp-server version v3only
```

Spanning Tree Protocol (STP)

Both Cisco and IBM isCLI can run different versions of STP.

Cisco IOS

Example A-27 shows how to configure different STP modes in Cisco IOS.

Example: A-27 Configuring different STP modes in Cisco IOS

```
configure terminal
spanning-tree mode <mst/pvst/rapid-pvst>
```

IBM isCLI

Spanning-tree mode by default in IBM isCLI is rapid Per-VLAN Spanning Tree (PVRST). Using PVRST, each VLAN runs a separate instance of spanning tree. To configure other STP modes, see Example A-28.

Example: A-28 Configuring different STP modes in IBM isCLI

```
configure terminal
spanning-tree mode <disable/mst/pvrst/rstp>
```

SSH and Telnet

Use the following commands to configure SSH and Telnet.

Cisco IOS

SSHv1 or SSHv2 can be configured in Cisco IOS as shown in Example A-29.

Example: A-29 Cisco IOS ssh and telnet configuration

```
switch(config)#ip ssh authentication-retries 2
switch(config)#ip ssh version 2
switch(config)#line vty 0 4
switch(config-line)#transport input ssh
switch(config)#line vty 5 15
switch(config-line)#transport input telnet ssh
```

IBM isCLI

SSH is disabled by default. Enabling SSH generates all applicable keys automatically for the user. Telnet is enabled by default, but can be deactivated. Example A-30 for how to enable both SSH and Telnet.

Example: A-30 IBM isCLI ssh and telnet configuration

```
switch(config)#ssh enable
switch(config)#no access telnet enable
```

Syslog

Syslog can be configured to send log messages to a configured syslog server. Severity levels are configured from emergency-only =0 to full debug =7.

Cisco IOS

Example A-31 shows how to configure the Syslog level and how to configure a syslog server in Cisco IOS.

Example: A-31 Syslog configuration in Cisco IOS

```
configure terminal
logging monitor <0-7>
logging <server-ip>
```

IBM isCLI

Example A-32 shows how to configure the Syslog level and server in IBM isCLI.

Example: A-32 Syslog configuration in IBM isCLI

```
configure terminal
logging host 1 address <server ip>
logging host 1 severity <severity>
no logging log link
```

Port aggregation (static)

To create static port aggregation (or aggregation over Etherchannel) between a Cisco IOS switch and IBM System Networking switch, use the following commands.

Cisco IOS

Example A-33 shows how to configure a static port aggregation in Cisco IOS. The configuration is almost identical to creating an LACP aggregation except that the mode is set to on and not active.

Example: A-33 Configuring a static port aggregation in Cisco IOS

```
switch(config)#interface range gigabit 0/1 - 2
switch(config-if-range)#channel-group <number> mode on
switch(config)#interface port-channel <number>
switch(config-if)#no shutdown
```

IBM isCLI

Example A-34 shows how to configure a static port aggregation in IBM isCLI that connects to the aggregated port created on a Cisco IOS switch. Note that IBM isCLI does not support the Cisco proprietary aggregation protocol PAgP.

Example: A-34 Configuring a static port aggregation in IBM isCLI

```
switch(config)#portchannel <number> port <port>
switch(config)#portchannel <number> port <port>
switch(config)#portchannel <number> enable
```

Port aggregation (LACP)

To create a port aggregation using the LACP protocol (IEEE 802.3ad), use the following commands.

Cisco IOS

Example A-35 shows how to configure port aggregation using LACP in Cisco IOS. The configuration is almost identical to creating a static port aggregation except the mode is set to active instead of merely on.

Example: A-35 Configuring LACP port aggregation in Cisco IOS

```
switch(config)#interface range gigabit <0/X - Y>
switch(config-if-range)#channel-group <number> mode active
switch(config)#interface port-channel <number>
switch(config-if)#no shutdown
```

IBM isCLI

Choose the ports to be bundled in an LACP grouping, give them an arbitrary “key” value and enable the bundle with the **lacp mode active** command. Use different “key” values in different port aggregations. See Example A-36.

Example: A-36 Configuring LACP port aggregation in IBM isCLI

```
switch(config)# interface port <portX-portY>
switch(config-if)# lacp key <key>
switch(config-if)# lacp mode active
```

VLAN tagging (802.1q)

VLAN trunking is supported with the 802.1q protocol in both Cisco IOS and IBM isCLI.

Cisco IOS

Example A-37 show how to configure 802.1q VLAN trunking.

Example: A-37 Configuring multiple ports for VLAN tagging in Cisco IOS

```
interface range gig <0/X - Y>
switchport trunk encapsulation dot1q
switchport mode trunk
no switchport trunk native vlan
switchport trunk allowed van <vlan>
```

IBM isCLI

In IBM isCLI, enable tagging on the port itself. Use a pvid=1 if no systems on these ports need to have a “native” VLAN defined. Untagged frames on interfaces with tagging enabled need to be given a PVID (Port VLAN identifier) if the endstation device cannot tag or recognize 802.1q frames. Similar functions are enabled on Cisco equipment with “switchport trunk native van”. Multiple entries can be specified by using the syntax in Example A-38.

Example: A-38 Configuring multiple ports for VLAN tagging in IBM isCLI

```
8264(config)# interface port <portX-portY>
8264(config-if)# tagging
8264(config-if)# pvid 1
8264(config-if)# exit
```

Example A-39 shows how to create VLANs and assign ports in IBM isCLI.

Example: A-39 Creating VLANs and assigning ports in IBM isCLI

```
switch(config)# vlan <vlan>
switch(config-vlan)# enable
8264(config-vlan)# member <portX-portY>
```



Easy Connect

IBM Easy Connect is a simple configuration mode implemented on IBM System Networking Ethernet and Converged switches. It enables easy integration of IBM Flex/PureSystems with existing Cisco and other vendor data center networks. Easy Connect makes connecting to existing upstream networks simple while enabling advanced in-system connectivity at the network edge. It also allows administrators to allocate bandwidth and optimize performance. In short, it supports both your existing and future network.

This appendix includes the following sections:

- ▶ Introduction to IBM Easy Connect
- ▶ Single Mode
- ▶ Storage Mode
- ▶ Easy Connect Multi-Chassis Mode
- ▶ Customer examples with diagrams
- ▶ Easy Connect limitations

Introduction to IBM Easy Connect

Easy Connect configuration mode enables IBM PureSystems to meet the primary selection criteria for adding new integrated systems to existing data center networks. Instead of requiring complex network configuration for each individual server, Easy Connect mode allows connection to a complete, integrated multiprocessor chassis or rack. This complete system includes PureSystems compute, storage, system management, and networking resources. Easy Connect allows you to manage this scalable resource with the simplicity of a single network node.

The following IBM System Networking Ethernet switches support the Easy Connect feature:

1. IBM Flex System Fabric EN4093/EN0493R and Virtual Fabric 10 Gb Scalable Switches
2. IBM Flex System Fabric CN4093 10 Gb Converged Scalable Switch
3. IBM System Networking RackSwitch G8264CS
4. IBM RackSwitch G8264 or G8124E
5. IBM RackSwitch G8264 (not in FCoE mode)

Easy Connect mode provides transparent PureSystems connectivity to your existing Cisco or other vendor network. With Easy Connect enabled on the EN4093/R, CN4093, or G8264 switches, the core network sees a “big pipe” for compute traffic to and from the PureSystems chassis. The switch becomes a simple I/O module that connects servers and storage with the core network. It aggregates compute node ports. The switch behaves similarly to Cisco Fabric Extension (FEX) by appearing as a “dumb” device to the upstream network, with the main difference being that intra-chassis switching is supported. Unlike Cisco FEX, traffic does not have to be sent upstream if the network destination is housed in the same physical chassis.

The Spanning Tree Protocol is disabled on the supported IBM System Networking switch in all Easy Connect modes, eliminating the data center administrator’s spanning tree concerns. This loop-free topology requires no additional configuration after it is set up. It helps to provide economical bandwidth use with prioritized pipes and network virtualization for both Intel and Power Compute nodes.

Single Mode

Easy Connect Single Mode allows the IBM Flex System EN4093/R switch to act as a Fabric Extension module in a Cisco network. If you use Active/Passive NIC teaming with no NIC bonding (LACP or static PortChannel) on the Compute Nodes, your system is well suited for Single Mode.

Figure B-1 shows the Single Mode configuration.

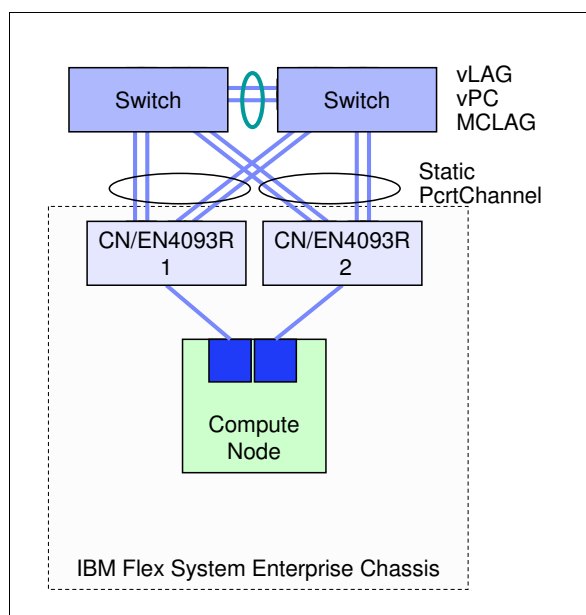


Figure B-1 IBM Easy Connect Single Mode diagram

Single Mode has the following important distinctions:

1. All local Layer-2 traffic pointing to the same I/O Bay in the Enterprise Chassis remains within the same chassis.
2. Because the CN4093 or EN4093/EN4093R I/O modules are not connected together with a Virtual Link Aggregation Group (vLAG), traffic that is destined for Compute Nodes using different I/O Bays within the same Enterprise Chassis must travel to the upstream switch, and then back down.
3. Each Enterprise Chassis appears as two separate devices to the upstream network when you are using two I/O modules.

To configure the CN4093 or EN4093/EN4093R I/O modules for Easy Connect Single Mode, complete the following steps:

1. Connect to the I/O module's CLI interface by using Telnet or SSH.
2. Change the configuration mode to the Industry Standard CLI (isCLI) if it is not already configured to do so as shown in Example B-1. Enable the CLI prompt in the last step if the Flex System Manager (FSM) is being used in the environment.

Example B-1 Changing the I/O module to use the isCLI

```
/boot/mode iscli
/boot/reset
/boot/prompt enable
```

3. If the I/O module is not already in a factory default configuration, reset it as shown in Example B-2 after you connect to it through Telnet/SSH.

Example B-2 Resetting the I/O module to a factory default configuration

```
EN4093> enable
EN4093# configure terminal
```

```
EN4093#(config) boot configuration-block factory
EN4093#(config) reload
```

4. After the I/O module returns to a factory default configuration, complete the steps shown in Example B-3 to enable Easy Connect Single Mode.

Example B-3 Implementing Easy Connect Single Mode

```
spanning-tree mode disable
portchannel 1 port ext1-ext10 enable
vnic enable
    vnic vnicgroup 1
    vlan 4091
    port INTA1-INTA14
    portchannel 1
    enable
    failover
    exit
write memory
```

5. Easy Connect Single Mode is now implemented.

Note: The IBM Virtual Fabric Switch Module (VFSM) for the IBM BladeCenter H or HT chassis is supported by Easy Connect Single and Storage Modes. The configuration steps are identical. This can also be done in a System x environment with rack servers by using the G8124, G8264, or G8264CS.

Easy Connect Single Mode has these important considerations and potential next steps:

- ▶ Configure Spanning-Tree BPDU Guard and Edge on the upstream switch for extra protection. These are enabled by default on Cisco Nexus 2000 Fabric Extender ports, and cannot be disabled.
- ▶ Setting a spanning-tree type network on an upstream Cisco Nexus port is not supported.

Storage Mode

Easy Connect Storage Mode allows the IBM Flex System EN4093/R switch to act as a Fabric Extension module in a Cisco network running on Fibre Channel over Ethernet (FCoE) connections. Storage Mode is nearly identical to Single Mode from a configuration standpoint. The only difference is that Converged Enhanced Ethernet (CEE) must be enabled in order for FCoE to function.

Storage Mode is illustrated in Figure B-2.

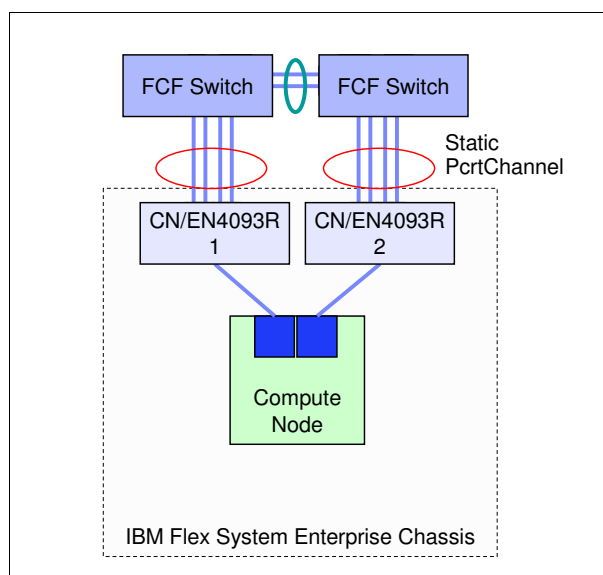


Figure B-2 BM Easy Connect Storage Mode diagram

The distinctions that are listed for Single Mode are the same for Storage Mode.

To configure the CN4093 or EN4093/EN4093R I/O modules for Easy Connect Storage Mode, complete the following steps:

1. Connect to the I/O module's CLI interface by using Telnet or SSH.
2. Change the configuration mode to the Industry Standard CLI (isCLI) if it is not already configured to do so as shown in Example B-4. Enable the CLI prompt in the last step if the Flex System Manager (FSM) is being used in the environment.

Example B-4 Changing the I/O module to use the isCLI

```
/boot/mode iscli
/boot/reset
/boot/prompt enable
```

3. If the I/O module is not already in a factory default configuration, reset it as shown in Example B-5 after you connect to it through Telnet/SSH.

Example B-5 Resetting the I/O module to a factory default configuration

```
EN4093> enable
EN4093# configure terminal
EN4093#(config) boot configuration-block factory
EN4093#(config) reload
```

4. Implement Storage Mode using the command shown in Example B-6. The only difference from Single Mode is highlighted in bold text.

Example B-6 Implementing Easy Connect Storage Mode

```
spanning-tree mode disable
portchannel 1 port ext1-ext10 enable
vnic enable
    vnic vnicgroup 1
    vlan 4091
```

```
port inta1-inta14
portchannel 1
enable
failover
exit
cee enable
write memory
```

5. Easy Connect Storage Mode is now implemented.

The same considerations that are listed for Single Mode and next steps apply to Storage Mode, except for the following caveat:

- IBM Networking OS 7.6 and earlier does not support FCoE traffic over multiple aggregated links, either using LACP or static PortChannels.

Easy Connect Multi-Chassis Mode

Easy Connect Multi-Chassis Mode allows IBM RackSwitch G8264 (acting as an aggregator for multiple chassis) and Flex System EN4093/R switches to act as Fabric Extension modules in a Cisco network.

If you use Active/Active NIC teaming with either Link Aggregation Control Protocol (LACP, or IEEE 802.3ad), or Static IP Hash on the Compute Node, your system is suited to Multi-Chassis Mode as illustrated in Figure B-3.

Multiple chassis: Alternatively, multiple chassis can connect to a pair of G8264s at the top-of-rack going out to your existing network.

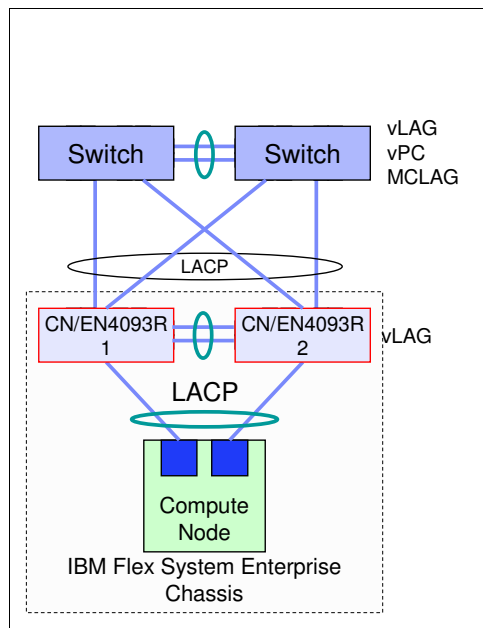


Figure B-3 IBM Easy Connect Multi-Chassis Mode diagram

Multi-Chassis Mode has the following important distinctions:

1. Because the CN4093 or EN4093/EN4093R I/O modules are connected together with a vLAG inter-switch link (ISL), all layer-2 traffic destined for Compute Nodes using either the same, or different I/O bays within the same Enterprise Chassis never leaves the chassis.
2. Each enterprise chassis appears as a single device to the upstream network when you use two I/O modules.
3. All operating systems (IBM AIX®, Linux, Windows, VMWare, VIO) within the IBM Flex System Enterprise Chassis must TAG VLANs.

Exception: If Flex System Manager is used, you must enable the Top-of-Rack Port “Native VLAN ID” with the VLAN that the FSM is configured on because the FSM cannot TAG.

4. Multi-Chassis Mode allows for pNIC or Switch Independent vNIC modes to be used on the Compute Node network adapters. If multiple vNIC Groups are used for either traffic separation or you are using IBM Virtual Fabric Mode, each vNIC Group requires its own uplink/PortChannel.

Restriction: At the time of writing, IBM Flex System POWER® Nodes support pNIC mode only.

5. Multi-Chassis Mode allows for the eventual implementation of IBM Virtual Fabric Mode.

Implementation with CN/EN4093/R

To configure the CN4093 or EN4093/R I/O modules for Easy Connect Multi-Chassis Mode, complete the following steps:

1. Restore the factory default configuration to the I/O module. Detailed steps for this are described in Example B-2 on page 317.
2. Disable the Spanning-Tree protocol globally.
3. Configure all the internal (INT) and external (EXT) CN4093 or EN4093/R ports by using the “tagvid-ingress” keyword. Use VLAN 4091 as the PVID.
4. Enable 802.1Q VLAN tagging on the external ports that are used as the vLAG Peer Link between the I/O modules. Use VLAN 4090 (vLAG ISL VLAN) as the PVID. Add VLAN 4091 as a tagged member.
5. Configure all required LACP aggregations (vLAG Peer Link, EXT, and INT ports).
6. Configure a superfluous IP address to be used by the management EXT port vLAG Health Check parameter. Consider using address 1.1.1.1 for the first I/O module, and 1.1.1.2 for the second I/O module.
7. Configure the vLAG ISL, Health Check peer-ip, and all associated vLAG pairs.
8. Easy Connect Multi-Chassis Mode is now implemented on the CN/EN4093/R.

Note: The IBM VFSM for the IBM BladeCenter H or HT chassis does not work in Multi-Chassis Mode because it does not support vLAG.

A sample script to enable Easy Connect Multi-Chassis Mode on the CN/EN4093/R I/O module is shown in Example B-7.

Example B-7 Sample script for Easy Connect Multi-Chassis Mode on CN/EN4093/R

```
spanning-tree mode disable
interface port ext9,ext10      --> ISL vLAG Peer-Link Ports
    pvid 4090
    tagging
    lacp key 1001
    lacp mode active
vlan 4090
    enable
    name Peer-Link
vlan 4091
    enable
    name Intel-Nodes
    member int1-int14,ext1-ext4,ext9,ext10
interface port int1-int14,ext1-ext4
    tagpvid-ingress
interface port ext1-ext4      --> uplink ports to AGG/Core
    lacp key 4091             --> use SAME key on both VFSM INTEL Uplinks (4091)
    lacp mode active
interface port inat1          --> INTa1 on both Switches will be in same
PortChannel using vLAG (lacp key MUST match)
    lacp key 101
    lacp mode active
interface port inat2
    lacp key 102
    lacp mode active
interface ip 127              --> IP 127 is dedicated to the MGT Port used for
vLAG health check
    ip address 1.1.1.1
    enable
vlag ena
vlag isl peer-ip 1.1.1.2      --> other switch will use 1.1.1.1
vlag isl vlan 4090
vlag isl adminkey 1001
vlag tier-id 10               --> each pair of switches connecting to each
other should be a different Tier-ID
vlag adminkey 4091 enable
vlag adminkey 101 enable
vlag adminkey 102 enable     --> repeat for each Server using 802.3ad / LACP
write memory
```

Implementation with G8264

If you are using a pair of IBM RackSwitch G8264 switches in the overall topology as shown in Figure B-4, the following section describes how Easy Connect can be implemented. Possible implementations include a pre-racked, pre-cabled IBM PureFlex System Express, Standard, or Enterprise rack configuration.

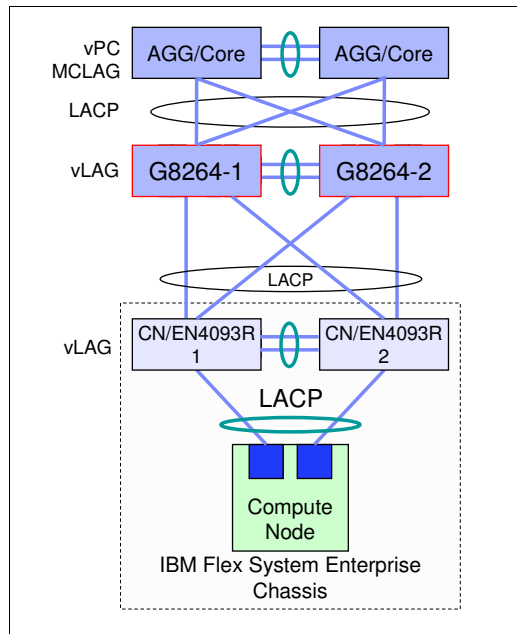


Figure B-4 IBM Easy Connect Multi-Chassis Mode with RackSwitch G8264

To configure the RackSwitch G8264 for Easy Connect Multi-Chassis Mode, complete the following steps:

1. Restore the factory default configuration to the G8264. Generalized steps for the EN4093/R can be used and are described in Example B-2 on page 317.
2. Disable the Spanning-Tree protocol globally.
3. Configure all the upstream and downstream G8264 ports by using the **tagpvid-ingress** keyword. Use VLAN 4091 as the PVID.
4. Enable 802.1Q VLAN tagging on the ports that you are using as the vLAG Peer Link between the G8264s. Use VLAN 4090 (vLAG ISL VLAN) as the PVID. Add VLAN 4091 as a tagged member.
5. Configure all required LACP aggregations (vLAG Peer Link, CN4093/EN4093/R facing ports).
6. Configure a superfluous IP address to be used by the management EXT port vLAG Health Check parameter. Consider using address 1.1.1.1 for the first I/O module, and 1.1.1.2 for the second I/O module.
7. Configure the vLAG ISL, Health Check peer-ip, and all associated vLAG pairs.
8. Easy Connect Multi-Chassis Mode is now implemented on the RackSwitch G8264.

A sample script to enable Easy Connect Multi-Chassis Mode on the RackSwitch G8264 is shown in Example B-8.

Example B-8 Sample script for Easy Connect Multi-Chassis Mode on RackSwitch G8264

```

spanning-tree mode disable          --> Optional
interface port 1,5                  --> 2x 40Gb ISL (e.g. between G8264's)
    tagging
    pvid 4090
    lacp key 4090
    lacp mode active
vlan 4090
    enable
    name Peer-Link
vlan 4091
    enable
    name "Transparent-Ports"
interface port 17-64                --> Uplinks and CN/EN4093/R facing Ports ONLY
    tagpvid-ingress
interface port 17,18                --> Uplink ports to AGG/Core
    lacp key 1001
    lacp mode active
interface port 19,20                --> Ports facing first PureFlex enclosure
    lacp key 1920
    lacp mode active
interface port 21,22                --> Ports facing second PureFlex enclosure
    lacp key 2122
    lacp mode active
vlag enable
vlag isl adminkey 4090
vlag tier-id 1
vlag adminkey 1001 ena              --> Uplink PortChannel to AGG/Core
vlag adminkey 1920 ena
vlag adminkey 2122 ena              --> Repeat for each Port-Channel to each
CN/EN4093/R
write memory

```

Easy Connect Multi-Chassis Mode has the following consideration:

- ▶ Configure Spanning-Tree BPDU Guard and Edge on the upstream switch for extra protection.

Customer examples with diagrams

The following section lists common implementation scenarios with Easy Connect for various industries that have purchased IBM PureFlex System hardware. Requirements are listed as dictated by the customer, and a network diagram to fit those requirements is displayed.

Telecommunications customer

This customer had the following requirements:

- ▶ No Spanning Tree or any other protocols that are seen by the network.
- ▶ Upstream connection must be into a Cisco Nexus 2000 Fabric Extender that is not running vPC.
- ▶ The EN4093/R I/O modules in the IBM Flex System Enterprise Chassis must be transparent devices that require no management by any group after initial setup.

Figure B-5 shows how Easy Connect satisfies all of the telecommunications customer's requirements.

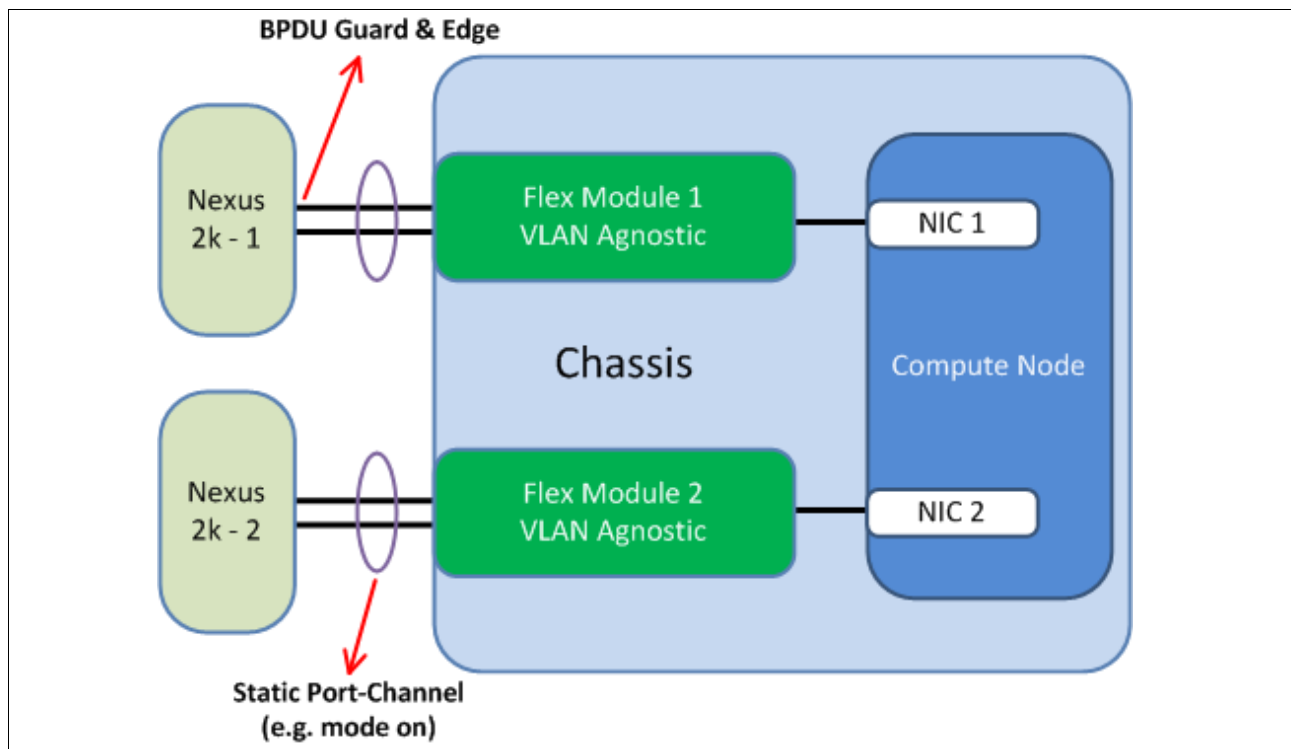


Figure B-5 Telecommunications customer network diagram

State government customer

This customer had the following requirements:

- ▶ Use LAN on Motherboard (or LoM) in Virtual Fabric Mode so bandwidth can be adjusted dynamically for each vNIC as required.
- ▶ Dedicated uplink vPC PortChannel from each EN4093/R for each vNIC Group for separation of traffic.
- ▶ The EN4093/R I/O modules in the IBM Flex System Enterprise Chassis must be transparent devices that require no management by any group after initial setup.

Figure B-6 shows how Easy Connect satisfies all of the state government customer's requirements.

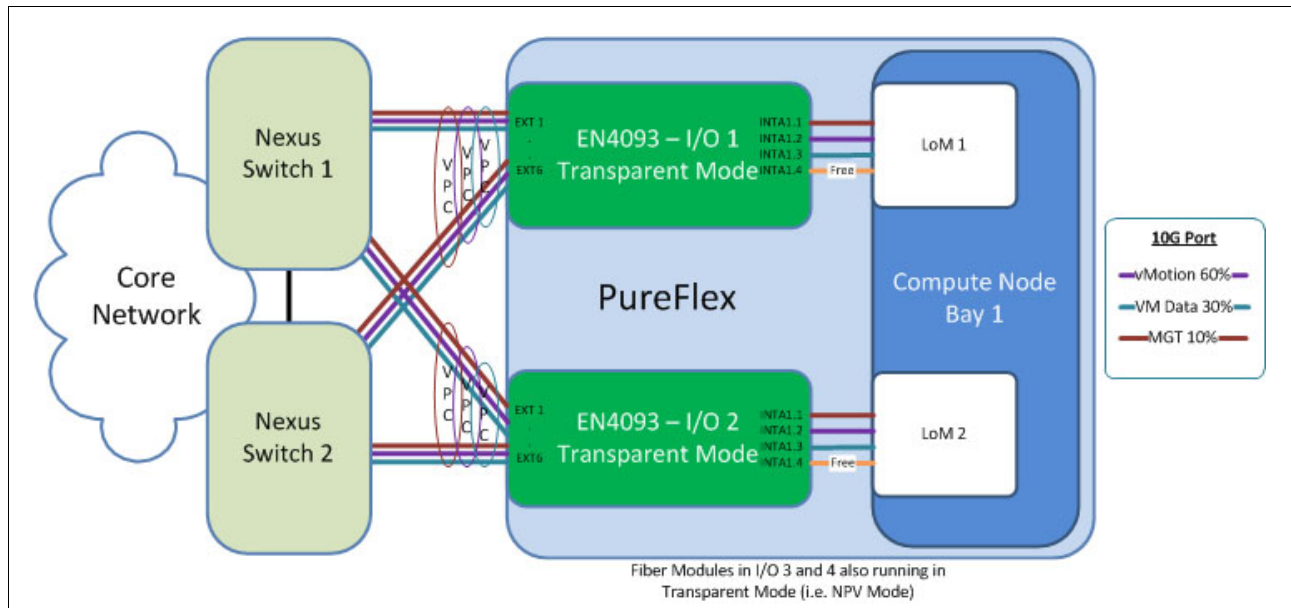


Figure B-6 State government customer network diagram

Medical center customer

This customer had the following requirements:

- ▶ Separation of and Dedicated Fibre Channel and Ethernet from each Compute Node and IBM Flex System Enterprise Chassis.
- ▶ Total hardware redundancy that includes both NIC and ASIC on each Compute Node using the CN4054 mezzanine adapter.
- ▶ Transparency on both Ethernet (Easy Connect) and Fibre Channel (NPV).

Figure B-7 shows how Easy Connect satisfies all of the medical center customer's requirements.

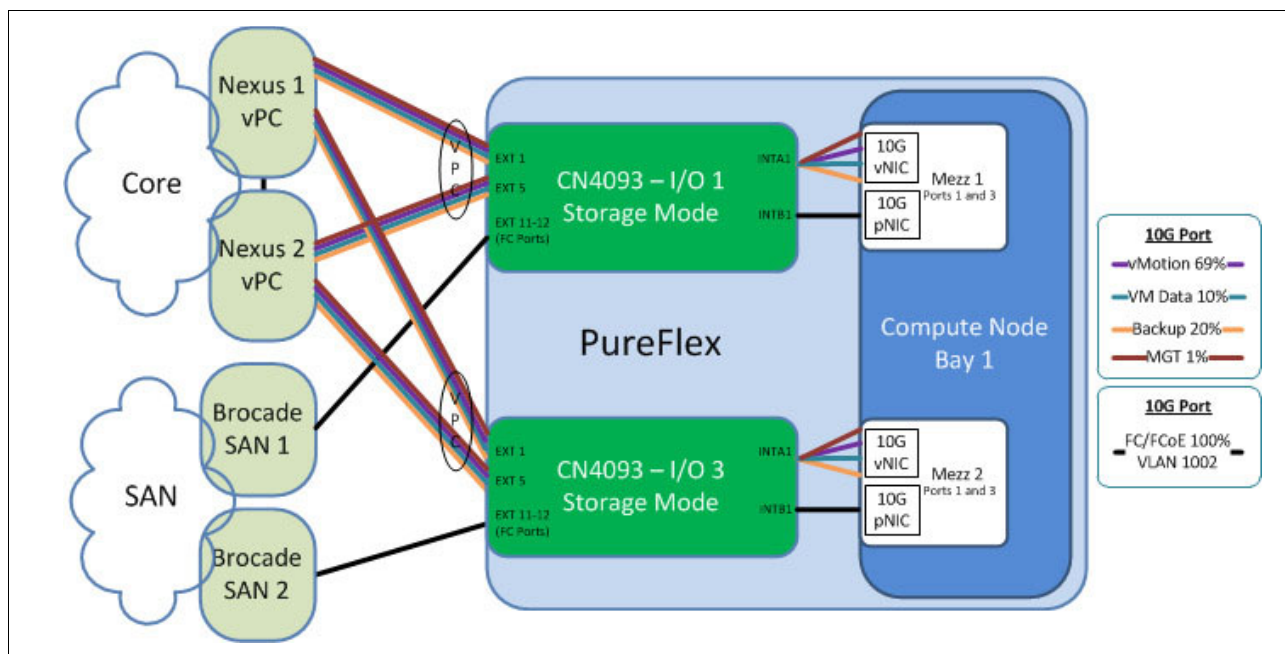


Figure B-7 Medical center customer network diagram

Easy Connect limitations

When configured for any Easy Connect mode, the following stand-alone features are not supported:

- ▶ Basic Routing
- ▶ Border Gateway Protocol (BGP)
- ▶ Edge Virtual Bridging / 802.1QBG
- ▶ IGMP Relay, IGMP Querier, IGMP Multicast Snooping and IGMPv3
- ▶ Stacking
- ▶ OSPF and OSPFv3
- ▶ Policy-Based Routing
- ▶ RIP
- ▶ Routed Ports
- ▶ Virtual Router Redundancy Protocol (VRRP)
- ▶ VMReady across the data center

Additionally, if multi-tenant security is a concern within the same IBM Flex System Enterprise Chassis, Easy Connect might not be recommended because each vNIC group is a single broadcast domain.

Related publications

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

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- ▶ *Connecting an IBM PureFlex System to the Network*, TIPS0941
- ▶ *IBM and Cisco: Together for a World Class Data Center*, SG24-8105
- ▶ *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984
- ▶ *IBM Flex System and PureFlex System Network Implementation*, SG24-8089
- ▶ *IBM PureFlex System Solutions for Managed Service Providers*, REDP-4994
- ▶ *IBM System Networking RackSwitch G8264*, TIPS0815
- ▶ *Implementing Systems Management of IBM PureFlex System*, SG24-8060
- ▶ *Moving to IBM PureFlex System: x86-to-x86 Migration*, REDP-4887

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

In addition to the power of your favorite search engine, these websites are also relevant as further information sources:

- ▶ IBM PureFlex Systems
<http://www.ibm.com/systems/pureflex/index.html>
- ▶ IBM System Networking
<http://www.ibm.com/systems/networking/>
- ▶ Cisco Systems
<http://www.cisco.com/>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



IBM Flex System and PureFlex System Network Implementation with Cisco Systems

(0.5" spine)
0.475" <-> 0.873"
250 <-> 459 pages



IBM Flex System and PureFlex System Network Implementation with Cisco Systems



Redbooks®

Connect IBM PureFlex Systems to a Cisco Network

Troubleshoot and maintain the IBM EN4093

Understand Cisco IOS and IBM N/OS differences

To meet today's complex and ever-changing business demands, you need a solid foundation of server, storage, networking, and software resources. It must be simple to deploy and able to quickly and automatically adapt to changing conditions. You also need access to, and the ability to take advantage of, broad expertise and proven best practices in systems management, applications, hardware maintenance, and more.

IBM® PureFlex System is part of the IBM PureSystems family of expert integrated systems. It 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. With the PureFlex System, 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.

In this IBM Redbooks publication, the examples use a Cisco Nexus 5000 Series Switch, although any configurations should also apply to the Cisco Nexus 7000 Series Switch too. However, it is wise to check as there might be minor differences.

This book also covers the different variations for the implementation of these use cases when you use Cisco Catalyst Series Switches.

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