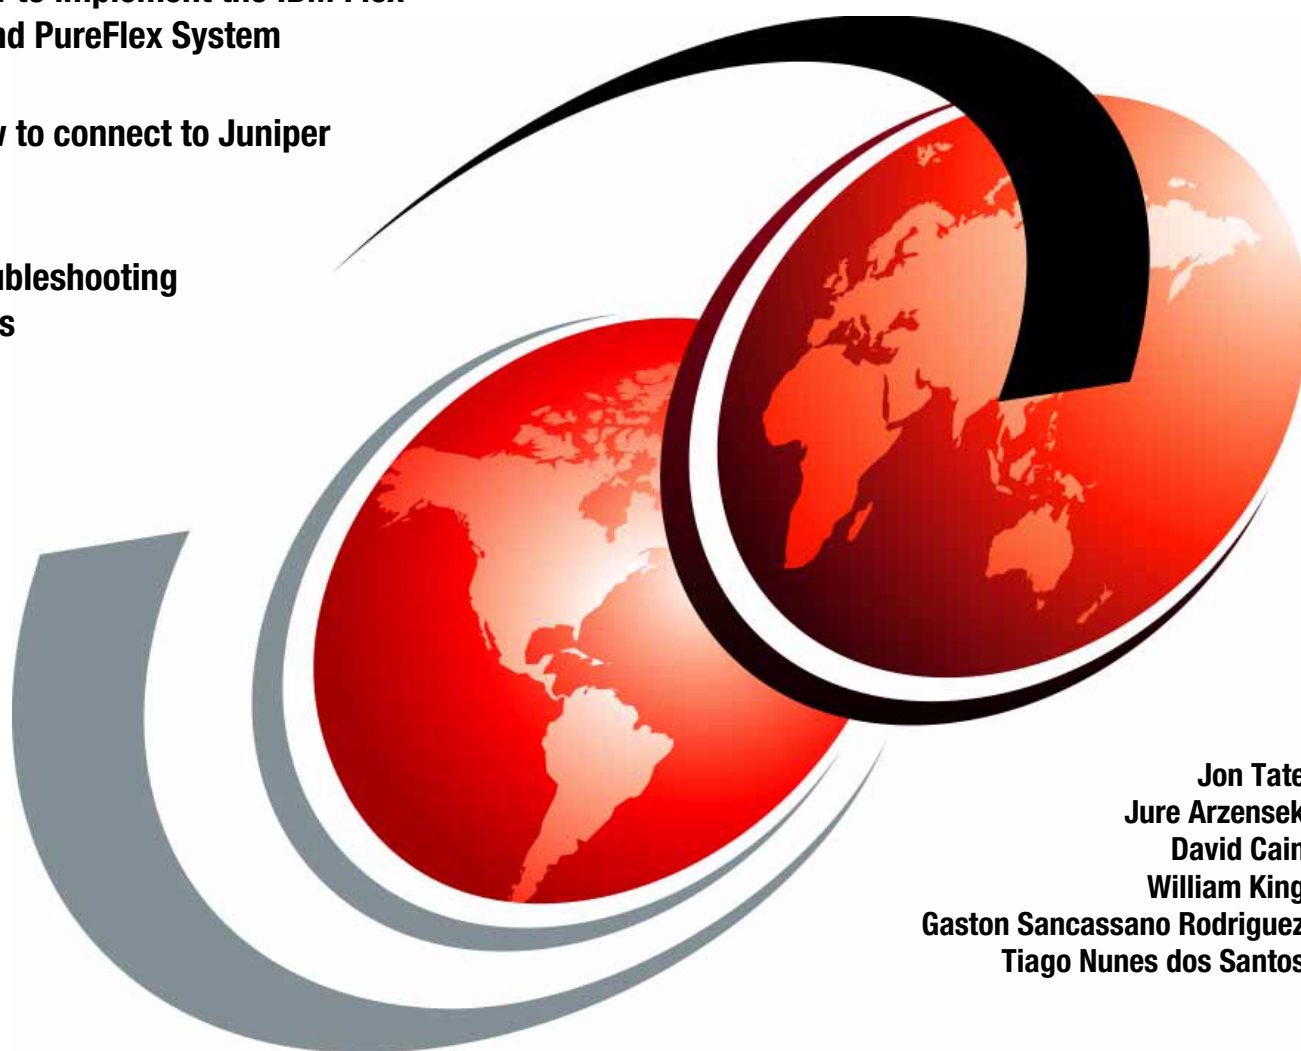


IBM Flex System and PureFlex System Network Implementation with Juniper Networks

Learn how to implement the IBM Flex System and PureFlex System

Learn how to connect to Juniper Networks

Learn troubleshooting techniques



Jon Tate
Jure Arzenek
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 Juniper Networks**

July 2013

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

First Edition (July 2013)

This edition applies to the IBM PureFlex System and Juniper EX4500 and QFX3500 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	xi
 Chapter 1. Introduction	 1
1.1 The IBM Flex System and IBM PureFlex System families	3
1.2 The goal of this book	3
1.3 Networking equipment	4
1.3.1 IBM System Networking	4
1.3.2 Juniper switches	5
 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	16
2.2 Physical layer on IBM Flex System Enterprise Chassis	18
2.3 IBM Flex System Ethernet I/O modules	21
2.3.1 IBM Flex System EN2092 1Gb Ethernet Scalable Switch	21
2.3.2 IBM Flex System Fabric EN4093 10Gb Scalable Switch	25
2.3.3 IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch	28
2.3.4 IBM Flex System EN4091 10Gb Ethernet Pass-thru module	34
2.3.5 Cables and transceivers for I/O modules	35
2.4 IBM Flex System Ethernet adapters	36
2.4.1 IBM Flex System CN4054 10Gb Virtual Fabric adapter	36
2.4.2 IBM Flex System EN4054 4-port 10 Gb Ethernet Adapter	38
2.4.3 IBM Flex System CN4058 8-port 10Gb Converged Adapter	39
2.4.4 IBM Flex System EN2024 4-port 1Gb Ethernet Adapter	42
2.4.5 IBM Flex System EN4132 2-port 10Gb Ethernet Adapter	44
2.4.6 IBM Flex System EN4132 2-port 10Gb RoCE Adapter	45
 Chapter 3. Layer 2 Overview	 49
3.1 Layer 2 Network protocols and technologies	50
3.1.1 Basic Frame Forwarding Concept	50
3.1.2 Virtual local area network and tagging	51
3.1.3 Spanning tree	52
3.1.4 Dynamic Link Aggregation Control Protocol	54
3.1.5 Virtual Link Aggregation Groups	55
3.1.6 Juniper Virtual Chassis	55
3.1.7 Link Layer Discovery Protocol	57
3.1.8 Layer 2 fail over	57
 Chapter 4. Layer 3 Overview	 59

4.1 Layer 3	60
4.1.1 Static routes	60
4.1.2 Default gateways	60
4.1.3 ECMP static routes	60
4.1.4 Routing Information Protocol v2	60
4.1.5 Open Shortest Path First for IPv4	61
4.1.6 Border Gateway Protocol	63
4.1.7 IPv6	64
4.1.8 Open Shortest Path First for IPv6	65
4.1.9 Virtual Router Redundancy Protocol	66
Chapter 5. Connecting IBM PureFlex System to a Juniper Network	69
5.1 Introduction	70
5.2 High availability overview	70
5.2.1 Looped and blocking design	71
5.2.2 Non-looped, single upstream device design	72
5.2.3 Non-looped, multiple upstream devices design	72
5.3 Fully redundant with Virtualized Chassis technology	74
5.3.1 Components used	74
5.3.2 Network topology and physical setup	74
5.3.3 EN4093 flex_1 configuration	76
5.3.4 G8264tor_1 configuration	81
5.3.5 Juniper EX4500-VC switch configuration	84
5.3.6 Verification and show command output	85
5.3.7 Full configuration files	106
5.4 Fully redundant with traditional spanning tree protocol	130
5.4.1 Topology and requirements	130
5.4.2 Components used	130
5.4.3 Network diagram and physical setup	130
5.4.4 EN4093 flex_1 configuration	132
5.4.5 G8264tor_1 configuration	137
5.4.6 EX4500_1 STP primary switch configuration	140
5.4.7 EX4500_2 STP secondary switch configuration	141
5.4.8 Verification and show command output	143
5.4.9 Full configuration files	166
5.5 Fully redundant with OSPF	190
5.5.1 Topology and requirements	190
5.5.2 Network diagram and physical setup	191
5.5.3 EN4093 flex_1 configuration	192
5.5.4 G8264tor_1 configuration	196
5.5.5 G8264tor_2 configuration	201
5.5.6 EX4500-VC switch configuration	206
5.5.7 Verification and show command output	207
5.5.8 Full configuration files	232
Chapter 6. IBM PureFlex Systems Troubleshooting and Maintenance	257
6.1 Troubleshooting	258
6.1.1 Basic troubleshooting procedures	258
6.1.2 Connectivity troubleshooting	262
6.1.3 Port mirroring	263
6.1.4 Serial cable troubleshooting procedures	265
6.2 Configuration management	266
6.2.1 Configuration files	266

6.2.2	Configuration blocks	266
6.2.3	Managing configuration files	266
6.2.4	Resetting to factory defaults	269
6.2.5	Password recovery	274
6.3	Firmware management	274
6.3.1	Firmware images	275
6.3.2	Upgrading the firmware with ISCLI	276
6.3.3	Recovering from a failed firmware upgrade	280
6.4	Logging and reporting	283
6.4.1	System logs	283
6.4.2	SNMP	285
6.4.3	Remote Monitoring	289
6.4.4	Using sFlow to monitor traffic	291

Appendix A. Fully Redundant Implementation with IBM vLAG and Juniper MC-LAG Technology

	Technology	293
A.1	Components used	294
A.2	Network topology and physical setup	294
A.3	Switches configuration	296
A.3.1	IBM EN4093flex and G8264tor switches configuration	296
A.3.2	Juniper QFX3500-1 switch configuration	296
A.3.3	Juniper QFX3500-2 switch configuration	298
A.4	Verification and show command output	299
A.4.1	QFX3500 output	299
A.4.2	Full configuration files	304

Appendix B. Junos Operating System to IBM Networking Operating System Command Comparison

	Comparison	307
B.1	Host name and banner	308
B.2	Management	308
B.2.1	Juniper Junos OS	308
B.2.2	IBM Networking OS	309
B.3	SSH and Telnet	309
B.3.1	Juniper Junos OS	309
B.3.2	IBM Networking OS	309
B.4	Local authentication	310
B.4.1	Juniper Junos OS	310
B.4.2	IBM Networking OS	310
B.5	Syslog	311
B.5.1	Juniper Junos OS	311
B.5.2	IBM Networking OS	311
B.6	Network Time Protocol	311
B.6.1	Juniper Junos OS	311
B.6.2	IBM Networking OS	312
B.7	Simple Network Management Protocol	312
B.7.1	Juniper Junos OS	312
B.7.2	IBM Networking OS	312
B.8	Link Layer Discovery Protocol	313
B.8.1	Juniper Junos OS	313
B.8.2	IBM Networking OS	313
B.9	Spanning Tree Protocol	313
B.9.1	Juniper Junos OS	313
B.9.2	IBM Networking OS	314
B.10	Interface speed and duplex	314

B.10.1	Juniper Junos OS	314
B.10.2	IBM Networking OS	314
B.11	VLAN	314
B.11.1	Juniper Junos OS	314
B.11.2	IBM Networking OS existing syntax	314
B.11.3	IBM Networking OS isCLI syntax	315
B.12	VLAN Tagging (802.1q)	315
B.12.1	Juniper Junos OS	315
B.12.2	IBM Networking OS existing syntax	315
B.12.3	IBM Networking OS isCLI syntax	316
B.13	Trunking and link aggregation: Static	316
B.13.1	Juniper Junos OS	316
B.13.2	IBM Networking OS existing syntax	316
B.13.3	IBM Networking OS isCLI syntax	317
B.14	Trunking and link aggregation: Link Aggregation Control Protocol	317
B.14.1	Juniper Junos OS	317
B.14.2	IBM Networking OS existing syntax	317
B.14.3	IBM Networking OS isCLI syntax	318
B.15	External authentication	318
B.15.1	Juniper Junos OS	318
B.15.2	IBM Networking OS	318
B.16	Bridge Protocol Data Unit Guard	319
B.16.1	Juniper Junos OS	319
B.16.2	IBM Networking OS	319
B.17	Dynamic Host Configuration Protocol snooping	319
B.17.1	Juniper Junos OS	319
B.17.2	IBM Networking OS	320
B.18	Port mirroring	320
B.18.1	Juniper Junos OS	320
B.18.2	IBM Networking OS	320
B.19	Open Shortest Path First configuration	321
B.19.1	Juniper Junos OS	321
B.19.2	IBM isCLI	321
Appendix C. Easy Connect		323
C.1	Introduction to IBM Easy Connect	324
C.2	Single Mode	324
C.2.1	Implementation	325
C.3	Storage Mode	326
C.3.1	Implementation	327
C.4	Easy Connect Multi-Chassis Mode	328
C.4.1	Implementation with CN/EN4093/R	329
C.4.2	Implementation with G8264	330
C.5	Customer examples with diagrams	332
C.5.1	Telecommunications customer	332
C.5.2	State government customer	332
C.5.3	Medical center customer	333
C.6	Easy Connect limitations	335
Related publications		337
IBM Redbooks intellectual property		337
Online resources		337
Help from IBM		337

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®
DB2®
developerWorks®
DS4000®
Extreme Blue®
IBM®
IBM Flex System™

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

Redbooks (logo) ®
System Storage®
System x®
Tivoli®
VMready®
xSeries®
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 that is simple to deploy and can 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, which is a part of the IBM PureSystems™ family of expert integrated systems, combines advanced IBM hardware and software along with patterns of expertise and integrates them into three optimized configurations that are simple to acquire and deploy so you can achieve faster time to value.

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

In this IBM Redbooks® publication, we use EX4500 core switches to demonstrate interoperability with the System Networking switches (RackSwitch™ G8264 top of rack switch and the Flex system fabric EN4093 10Gb scalable switch).

We also describe a redundant environment using QFX3500 switches running IBM Virtual-Link Aggregation Group (MC-LAG/vLAG) and Juniper Multi- Chassis-Link Aggregation Group.

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 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 9 years of experience in the Datacenter, with expertise in ethernet switching, storage, SAN, security, virtualization, xSeries®, 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 the year 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 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 IBM's 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

Pushkar Patil
Tim Shaughnessy
IBM San Jose

Scott Lorditch
IBM Denver

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

Greg Bassett
Vaishali Ghiya
Jeremy Wallace
Juniper Networks

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

In this book, we highlight the products and solutions that help to address some of the following challenges and trends that are facing data center network architects and engineers today:

- ▶ An ever increasing network traffic load that is driven by the use of multimedia data (including audio, video, and streaming), cloud oriented storage and big data analytics, and so on, are drivers for increased bandwidth and lower latency.
- ▶ The current economic climate is making clients minimize their infrastructure costs while also trying to keep and improve the services that are offered. Data center administrators are encouraged to use virtualization techniques to maximize resource usage.
- ▶ Data center consolidation with server virtualization, can mean that hardware from multiple vendors might be required to function together to produce a consistently smooth and transparent network service to run on.
- ▶ Ten Gigabit Ethernet is rapidly becoming the standard for data center networking. This transformation provides opportunities, but if it is not implemented diligently, it can cause potential complications.
- ▶ Stored information is becoming more pervasive, traversing from data center to hand-held device. Clients want to, and can, access their data wherever they are. Because data no longer stays geographically local, a heavier burden is placed on server-side network security mechanisms.
- ▶ Today, stored data and the information it contains are the currency of any company. The servers that are needed for processing this data can be purchased (or virtualized) by the dozen, moved from LAN to LAN, and easily replaced. By contrast, data storage restoration is expensive, time-consuming and might not even be possible in some cases. The data usually must be available constantly. Therefore, storage networks are designed from the outset to incorporate a high degree of redundancy and availability. The merger of traditional storage networks with Ethernet that use protocols such as FCoE or iSCSI means that these same high availability and redundant designs also must be incorporated into the Ethernet fabric.

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.

Administering a mixed network infrastructure is complex, often time-consuming and sometimes unsuccessful. One solution might be to adopt a single vendor, which, in theory, automatically eliminates any interoperability barriers. However, should a vendor's competitor present new exclusive solutions, the client might be locked in and miss out on a cutting-edge technology.

For this reason, IBM is strongly in favor of an open standards-based approach working closely with its business partners, the IEEE and other leading networking organizations and institutes. IBM System Networking's Ethernet sales amount to more than 15 million Ethernet switch ports deployed worldwide. Clients recognize the benefits of system networking solutions that enable a combination of best-of-breed components that are combined with an open approach, which allows these switches to easily connect into existing core network infrastructures.

This chapter includes the following topics:

- ▶ The IBM Flex System and IBM PureFlex System families
- ▶ The goal of this book
- ▶ Networking equipment

1.1 The IBM Flex System and IBM PureFlex System families

The IBM Flex System™ and IBM PureFlex System products start a new era of computing and represent the next generation of Smarter Computing, which offers 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 usage 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 zEnterprise® with zBX is focused on mainframe applications with a simplified workload-oriented management approach, PureSystems offers various implementation possibilities that are focused on a cloud-oriented customer strategy. This new platform is built over the following blocks concept to deliver value:

- ▶ Management
- ▶ Compute Nodes
- ▶ Storage Networking

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.

1.2 The goal of this book

The goal of this book is to demonstrate the interoperability of the IBM PureFlex System with an upstream core network that consists of Juniper switches. The practical scenarios that are demonstrated here reflect environments that are encountered within the industry.

1.3 Networking equipment

In this section, we present the IBM System Networking and Juniper Switch product families.

1.3.1 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 systems networking products includes Top-of-Rack (TOR) switches and the embedded switches in the IBM PureFlex System and IBM Flex System families. The IBM System Networking business is focused on driving data center networking by using the latest in Ethernet technologies engineered with IBM innovation.

The physical layout of most corporate networks evolved over time. Classic hub and router topologies gave way to faster switched topologies, particularly now that switches are increasingly intelligent. IBM System Networking switches are intelligent and fast enough to perform routing functions on a par with wirespeed 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 4k VLANs per switch.

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 currently supports the following standards and technologies:

- ▶ Virtual local area network (VLAN):
 - PVID: Port VLAN IDs
 - VLAN tagging
 - Protocol-based VLANs (PVLANS)
- ▶ Spanning Tree Protocol (STP):
 - Rapid Spanning Tree Protocol (RSTP)
 - Per-VLAN Rapid Spanning Tree Protocol (PVRST)
 - Multiple Spanning Tree Protocol (MSTP)
- ▶ Internet Protocol (IP) routing:
 - Static routes
 - Equal-Cost Multi-Path static routes (ECMP)
 - Routing Information Protocol (RIP): RIPv1 and RIPv2
- ▶ Open Shortest Path First (OSPF)
- ▶ Border Gateway Protocol (BGP):
 - eBGP
 - iBGP
- ▶ IP multicast
- ▶ Internet Group Management Protocol (IGMP)
- ▶ Protocol Independent Multicast (PIM):
 - PIM Sparse Mode
 - PIM Dense Mode
- ▶ Internet Protocol version 6 (IPv6)

- ▶ Neighbor Discovery (ND) protocol
- ▶ Port mirroring
- ▶ ACL-based mirroring
- ▶ sFlow monitoring
- ▶ Remote Monitoring (RMON)
- ▶ Trunking:
 - Static trunk groups (portchannel)
 - Dynamic Link Aggregation Control Protocol (LACP) trunk groups
- ▶ LACP IEEE 802.3ad
- ▶ Virtual Link Aggregation Groups (VLAG)
- ▶ Fast Uplink Convergence
- ▶ Network interface controller (NIC) teaming and Layer 2 failover
- ▶ Virtual Router Redundancy Protocol (VRRP)
- ▶ Active Multipath Protocol (AMP)

The stacking possibility features having a group of up to eight IBM System Networking switches that work together as a unified system. Stacking is supported on the EN4093/EN4093R, the RackSwitch G8264 switches, and on G800.

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

In this book, we use the IBM Flex System Fabric EN4093 10Gb Scalable Switch, which is described at this website:

http://www.ibm.com/systems/flex/networking/ethernet/en4093_10gb_vf/

1.3.2 Juniper switches

This section describes the Juniper switches that are found in the IBM portfolio.

Juniper EX-series switches

Juniper EX-series Ethernet switches were originally launched in 2008 and provide a complete range of Ethernet switches from branch to core devices. Juniper EX-series run the JunOS operating system.

EX2200-series

The EX2200 series offers 24- and 48-port fixed configuration 1Gb ports with and without Power over Ethernet (PoE) support. More details are available here:

<http://www.redbooks.ibm.com/redbooks.nsf/RedbookAbstracts/tips0816.html>

EX4200-series offered as the IBM J48E Ethernet Switch

The EX4200 series offers 24- and 48-port fixed configuration 1 Gb ports with and without PoE support. EX4200 series switches can have 1 Gb or 10 Gb uplinks. It comes in RJ45 and SFP versions. More details are available here:

<http://www.redbooks.ibm.com/redbooks.nsf/RedbookAbstracts/tips0810.html>

EX4500-series

The EX4500 series offers a compact, scalable high-performance platform that contains 10 Gbps ports.

Juniper QFabric switches

Juniper QFabric consists of three parts: Qfabric nodes, Qfabric interconnect, and Qfabric director. Qfabric creates a single-tier architecture in the data center improving speed, scalability, performance, and efficiency. Qfabric switches are designed to connect into the Qfabric as Qfabric nodes.

QFX-3500 Switch

The QFX3500 switch can act as a stand-alone 48-port 10 GbE top of rack switch with 4x 40 GbE uplinks. It has FCoE and FC gateway functionality. The QFX3500 also can be converted into a Qfabric node to connect into the Qfabric.

For more information about the Juniper product family, see this website:

<http://www.juniper.net>

Use cases approach

In this book, we use two EX4500 core switches to demonstrate interoperability with the System Networking switches (RackSwitch G8264 top of rack switch and the Flex system fabric EN4093 10Gb scalable switch). The three scenarios that are presented include a high availability fully redundant scenario that uses Juniper's virtual chassis technology with vLAG ISL virtual group technology on the IBM System Networking switches. A second scenario demonstrates interoperability in a more traditional STP environment. The third scenario demonstrates Layer 3 OSPF interoperability.

In Appendix A, "Fully Redundant Implementation with IBM vLAG and Juniper MC-LAG Technology" on page 293 we describe a redundant environment that uses two QFX3500 switches that are running Multi-chassis LAG instead of the Juniper Virtual chassis technology.



Layer 1 Overview

In this chapter, we provide details about IBM PureFlex System networking from the physical layer perspective. We begin by explaining Layer 1 networking concepts and terminology, and continue with description of IBM PureFlex System networking components (mid-plane connections, switches, adapters). Finally, we list transceivers and cables that are used with IBM PureFlex System networking options.

This chapter includes the following topics:

- ▶ 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. We describe 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 that is used for Ethernet twisted-pair cabling is likely the connector that most people recognize and associate 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 shown in Figure 2-2. This figure shows T568B pin layout. T568A pin layout is similar. The only difference between T568A and T568B is that orange and green pairs are swapped.




Figure 2-2 Straight through Ethernet cable with T568B pin layout

An Ethernet operating in 10/100 Mb mode only uses 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 two like devices are directly connected, such as 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 shown 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), where the X refers to 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 explained in this chapter: single-mode and multi-mode. The difference between the modes is the wavelength of the light used for the transmission, as shown in Figure 2-4 on page 11.




Figure 2-4 Multimode versus single-mode optic signaling

Single-mode fiber

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




Figure 2-5 Single-mode fiber cable

Single-mode fiber cabling allows for much longer cable lengths than multi-mode. For example, when 10GBASE-ER transceivers and suitable single mode fibre cable are used, it is possible to reach distances up to 40 km (24.85 mi).

Multimode fiber

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




Figure 2-6 Multimode fiber cable

The color of the outer coating is sometimes used to identify if a cable is a multi-mode 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 multi-mode fiber for civilian applications. This guideline is not always implemented (as shown in Figure 2-7), which shows a blue cable. The reliable method is to look at the specifications of the cable that are printed on the outer coating of the cabling (see a Figure 2-8 and Figure 2-9).



Figure 2-7 Blue 62.5 micron MMF cable




Figure 2-8 Yellow SMF cable



Figure 2-9 Orange 50 micron MMF cable

With multi-mode cabling, 10 Gbps Ethernet supports cable lengths of up to 550 m (1804.46 ft.), and 40 Gbps Ethernet supports cable lengths of up to 125 m (410.10 ft.).

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. LC connector often is used to connect fiber cable to an SFP transceiver.




Figure 2-10 LC fiber connector

EN4093/EN4093R scalable switch and CN4093 converged scalable switch contain external ports for 40 Gb QSFP+ transceivers. Cables that connect to QSFP+ transceivers have MTP connectors. Figure 2-11 shows an example of MTP connector.




Figure 2-11 MTP connector

Older connector types that were commonly used are the SC connector (as shown in Figure 2-12 on page 14), and the ST connector (as shown in Figure 2-13 on page 14). SC connectors were used to connect fiber cables to GBICs, which were widely used on previous generations of networking gear. Today, GBICs are uncommon because they were replaced with various forms of SFP transceivers.




Figure 2-12 SC fiber connector

ST connectors also are uncommon today, but they might still be found in older network infrastructure.



Figure 2-13 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, which provides flexibility of use for single or multi-mode implementations.

In today's Ethernet networks, we often use the following types of transceivers: SFP, SFP+, XFP, and QSFP. Figure 2-14 shows SFP, SFP+, and XFP transceivers (MMF and SMF varieties).




Figure 2-14 From left to right: SFP-MMF, SFP-SMF, SFP+-MMF, XFP-MMF, and XFP-SMF

A QSFP transceiver is shown in Figure 2-15.



Figure 2-15 QSFP transceiver

Table 2-2 shows a 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, 8Gb FC, OTU2
XFP	10 Gbps	10 Gb Ethernet, 10 Gb FC, SONET, OTU2
QSFP	40 Gbps	40 Gb Ethernet, 20Gb/40Gb InfiniBand

IBM EN4093/EN4093R and CN4093 switches support SFP, SFP+, and QSFP+ transceivers.

Direct Attached Cable

Direct Attached Cable (DAC) is a twinaxial (twinax) cable that can be used in 10 GbE and 40 GbE environments. The 10 GbE DAC has SFP+ housing on each end, which means you can plug it directly into SFP+ slot on a switch. The 40 GbE DAC has QSFP+ housing on each end, and can attach directly into QSFP+ slot on a switch.

DAC can be passive or active. Passive DAC contains no active components. EN4093/EN4093R and CN4093 switches support passive DAC cable lengths up to five meters. Active DAC contains active electronic components in SFP+ housing for enhanced signal quality.

QSFP+ DAC cable lengths up to 3 m (9.84 ft.) are supported on EN4093/EN4093R and CN4093 switches.

DAC cables are a cost-effective alternative to fiber cables for distances of 5 m (16.40) or less.

Figure 2-16 shows a DAC example: 3 m IBM Passive DAC SFP+ cable, P/N 90Y9430.



Figure 2-16 3 m IBM Passive DAC SFP+ cable, P/N 90Y9430

2.1.4 Physical configuration parameters

When we describe the physical layer (Layer 1) properties, we consider elements such as line speed and duplex.

Speed

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

Duplex

Duplex modes are full or half duplex. *Half duplex* is when a device can only send or receive at a time (see Figure 2-17), while *full duplex* devices can send and receive at the same time (see Figure 2-18).

Half duplex is supported only in older 10 and 100 Mbps devices.




Figure 2-17 Half-duplex mode




Figure 2-18 Full-duplex mode

Flow control

EN4093/EN4093R and CN4093 switches support Ethernet flow control, which is defined by the IEEE 802.3x standard. They also support Priority-based Flow Control (PFC), which is defined by the IEEE 802.1Qbb standard.

Under standard flow control, when a port becomes busy, the switch manages congestion by pausing all the traffic on the port, regardless of the traffic type. PFC provides more granular flow control, which allows the switch to pause specified types of traffic on the port, while other traffic on the port continues. PFC pauses traffic based on 802.1p priority values in the VLAN tag.

For any port, only one flow control method can be implemented at any time: PFC or standard IEEE 802.3x flow control.

When Converged Enhanced Ethernet (CEE) is off (the default), 802.3x standard flow control is enabled on all switch ports. When CEE is turned on, standard flow control is disabled on all ports. In its place, PFC is enabled on all ports for 802.1p priority value 3.

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 negotiating with each other.

One step to reduce negotiation problems is to ensure that both devices on a switch segment are configured the same. Configure both devices for autonegotiation or hardcode (manually configure) 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. We describe the physical connectivity between network adapters that are installed in compute nodes and switches that are installed in I/O bays. We also list and describe the switches and network adapters that are available for IBM Flex System.

The Ethernet networking I/O architecture for the IBM Flex System Enterprise Chassis includes various connectivity options for compute nodes that are 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 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 2-19.




Figure 2-19 Rear view of the Enterprise Chassis showing I/O module bays

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

- ▶ Network cards installed in a node
- ▶ I/O module installed in the I/O bay
- ▶ Port licenses enabled on the I/O module

Figure 2-20 shows how the mid-plane lanes connect between the node bays upfront and the I/O bays in the rear. The concept of an I/O module partition also is shown in Figure 2-20. 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 node bay in the front. By adding an upgrade license to the I/O module, it is possible to 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 other partitions, that partition is enabled through an upgrade on the I/O module. The node needs a card 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.

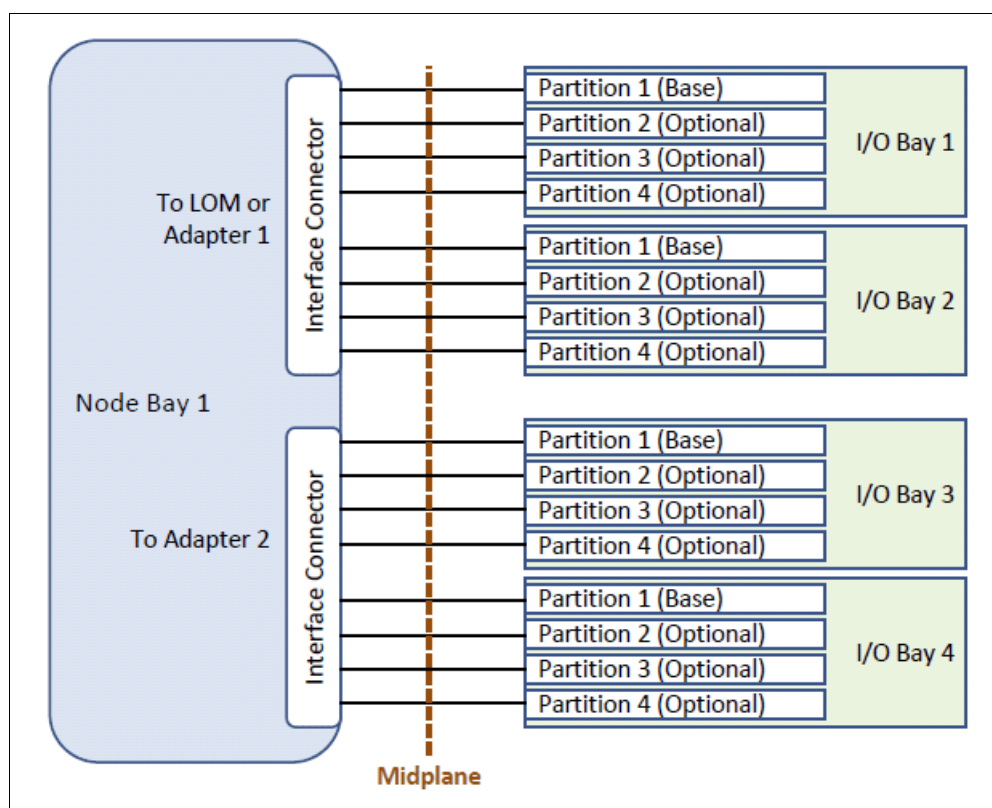


Figure 2-20 Sixteen lanes total of a single half-wide node bay toward the I/O bays

For example, if a dual port LAN on system board (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 2-21 on page 20.

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.

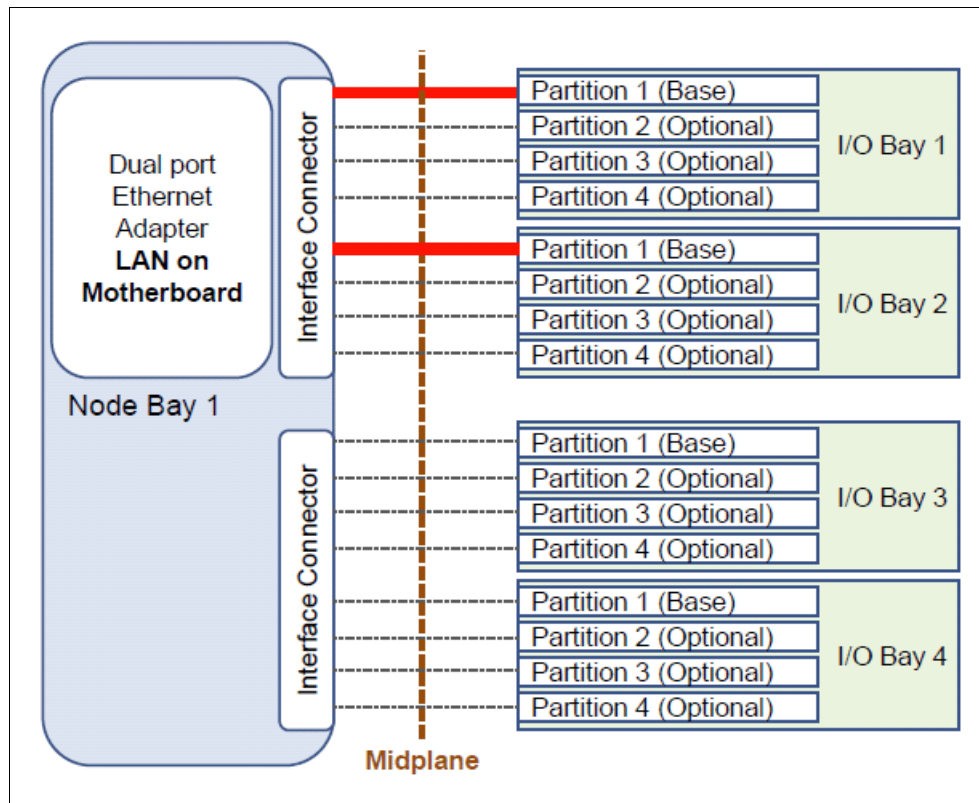


Figure 2-21 Dual port LOM connecting to partition on I/O bays 1 and 2 (all other lanes unused)

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

Currently, the nodes are limited to a maximum of two quad port adapters on a single half-wide node, thus setting the connection limit 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 upfront. By adding port licenses, a single EN2092 1 Gb Ethernet Switch can offer two 1 Gb ports to each half-wide node bay. The EN4093/EN4093R 10 Gb Scalable Switch and CN4093 10 Gb Converged 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 10 Gb Ethernet Pass-thru I/O module only can 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. The EN4093/EN4093R 10Gb Scalable Switch, CN4093 10 Gb Converged Scalable Switch and the EN2092 1 Gb Ethernet Switch support more than the base partition. Table 2-3 on page 22 and Table 2-4 on page 25 show the available I/O module partition upgrades.

At the time of this 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.

The following Ethernet I/O modules are available for deployment with the Enterprise Chassis:

- ▶ IBM Flex System EN2092 1Gb Ethernet Scalable Switch
- ▶ IBM Flex System Fabric EN4093 10Gb Scalable Switch
- ▶ IBM Flex System CN4093 10Gb Converged 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 primarily is a 1 Gb switch, which offers up to 28 x 1 Gb downlinks to the internal nodes, with a total combination of up to 20 x 1 Gb RJ45 uplinks and four 10 Gb uplinks with “pay-as-you-grow” scalability.

Figure 2-22 shows a view of the EN2092 1Gb Ethernet Switch.



Figure 2-22 The EN2092 1Gb Ethernet Switch

Ports that are enabled and available depend on the features that are 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:
 - A total of 28 internal full-duplex Gigabit ports with 14 ports enabled by default; an optional Feature on Demand (FoD) capability license is required to activate the other 14 ports.
 - Two internal full-duplex 1 GbE ports that are 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 direct-attach cables (DAC). 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 another 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 that includes the following limits:
 - 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
- 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 based on source/destination IP or MAC addresses or both
- Fast port forwarding and fast uplink convergence for rapid STP convergence
- ▶ Availability and redundancy:
 - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy.
 - IEEE 802.1D 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), which 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 and 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 4095 VLANs are supported per switch at any time, 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 re-marking 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
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 (separate sFLOW analyzer required elsewhere)
 - POST diagnostic tests

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

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

2.3.2 IBM Flex System Fabric EN4093 10Gb Scalable Switch

The EN4093/EN4093R 10Gb Scalable Switch is primarily a 10 Gb switch that 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 2-23.



Figure 2-23 The IBM Flex System Fabric EN4093 10Gb Scalable Switch

More information about the available upgrade options for this module is provided in Table 2-4.

Table 2-4 IBM Flex System Fabric EN4093 10Gb 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
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 enabled with Upgrade 2 require a 6-port adapter card, which is unavailable as of this writing.

The IBM Flex System Fabric EN4093 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 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 direct-attach cables (DAC). A total of 10 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. You also can use QSFP+ break-out cables (40 Gb QSFP+ to 4 x 10 Gb SFP+) to use the 40 GbE port as four 10 GbE ports.
- One RS-232 serial port (mini-USB connector) that provides another 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
 - Media access control (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 or MAC addresses or both
 - Fast port forwarding and fast uplink convergence for rapid STP convergence
- Availability and redundancy:
 - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy
 - IEEE 802.1D 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 4095 VLANs supported per switch at any given time, with VLAN numbers that range from 1 to 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 re-marking 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)
 - 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 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 *IBM Flex System Fabric EN4093 and EN4093R 10Gb Scalable Switches*, TIPS0864 at this website:

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

2.3.3 IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch

The IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch provides unmatched scalability, performance, convergence, and network virtualization, while delivering innovations to help address a number of networking concerns and providing capabilities that help you prepare for the future.

The switch offers full Layer 2/3 switching and FCoE Full Fabric and Fibre Channel NPV Gateway operations to deliver a converged and integrated solution. It is installed within the I/O module bays of the IBM Flex System Enterprise Chassis. The switch can help you migrate to a 10 Gb or 40 Gb converged Ethernet infrastructure and offers virtualization features, such as Virtual Fabric and IBM VMready, and the ability to work with IBM Distributed Virtual Switch 5000V.

Figure 2-24 shows the IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch.




Figure 2-24 IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch

The CN4093 switch is initially licensed for 14 10 GbE internal ports, two external 10 GbE SFP+ ports, and six external Omni Ports enabled.

The following ports can be enabled:

- ▶ Another 14 internal ports and two external 40 GbE QSFP+ uplink ports with Upgrade 1.
- ▶ Another 14 internal ports and six more external Omni Ports with the Upgrade 2 license options.
- ▶ Upgrade 1 and Upgrade 2 can be applied on the switch independently from each other or in combination for full feature capability.

Table 2-5 shows the part numbers for ordering the switches and the upgrades.

Table 2-5 CN4093 part numbers and feature codes

Description	Part number	Feature code (x-config/e-config)
IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch	00D5823	A3HH/ESW2
IBM Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 1)	00D5845	A3HL/ESU1
IBM Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 2)	00D5847	A3HM/ESU2
IBM Flex System Management Serial Access Cable	90Y9338	N/A

QSFP+ or SFP+ transceivers nor cables are included with the switch. They must be ordered separately.

The switch does not include a serial management cable. However, IBM Flex System Management Serial Access Cable, 90Y9338, is supported and contains two cables: a mini-USB-to-RJ45 serial cable and a mini-USB-to-DB9 serial cable. Either of these cables can be used to connect to the switch locally for configuration tasks and firmware updates.

The following base switch and upgrades are available:

- ▶ 00D5823 is the part number for the physical device, which comes with 14 internal 10 GbE ports enabled (one to each node bay), two external 10 GbE SFP+ ports that are enabled to connect to a top-of-rack switch or other devices. A total of six Omni Ports are enabled to connect to Ethernet or Fibre Channel networking infrastructure, depending on the SFP+ cable or transceiver that is used.

- ▶ 00D5845 (Upgrade 1) can be applied on the base switch when you need more uplink bandwidth with two 40 GbE QSFP+ ports that can be converted into 4x 10 GbE SFP+ DAC links with the optional break-out cables. This upgrade also enables 14 more internal ports, for a total of 28 ports, to provide more bandwidth to the compute nodes by using four-port expansion cards.
- ▶ 00D5847 (Upgrade 2) can be applied on the base switch when you need more external Omni Ports on the switch, or if you want more internal bandwidth to the node bays. The upgrade enables the remaining six external Omni Ports and 14 more internal 10 Gb ports (for a total of 28 internal ports) to provide more bandwidth to the compute nodes by using four-port expansion cards.
- ▶ Both 00D5845 (Upgrade 1) and 00D5847 (Upgrade 2) can be applied on the switch at the same time so that you can use six ports on an eight-port expansion card, and use all the external ports on the switch.

The IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch has the following features and specifications:

- ▶ Internal ports:
 - A total of 42 internal full-duplex 10 Gigabit ports. (By default, 14 ports are enabled. Optional FoD licenses are required to activate the remaining 28 ports.)
 - Two internal full-duplex 1 GbE ports are connected to the Chassis Management Module.
- ▶ External ports:
 - Two ports for 1 Gb or 10 Gb Ethernet SFP+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, 10GBASE-LR, or SFP+ copper direct-attach cables (DACs)). These two ports are enabled by default. SFP+ modules and DACs are not included and must be purchased separately.
 - Twelve IBM Omni Ports. Each of them can operate as 10 Gb Ethernet (support for 10GBASE-SR, 10GBASE-LR, or 10 GbE SFP+ DACs), or auto-negotiating as 4/8 Gb Fibre Channel, depending on the SFP+ transceiver that is installed in the port. The first six ports are enabled by default. An optional FoD license is required to activate the remaining six ports. SFP+ modules and DACs are not included and must be purchased separately.

Important: Omni Ports do not support 1 Gb Ethernet operations.

- Two ports for 40 Gb Ethernet QSFP+ transceivers or QSFP+ DACs. (Ports are disabled by default; an optional FoD license is required to activate them.) Also, you can use break-out cables to break out each 40 GbE port into four 10 GbE SFP+ connections. QSFP+ modules and DACs are not included and must be purchased separately.
- One RS-232 serial port (mini-USB connector) that provides another 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 the 10 Gb core infrastructure.
 - Non-blocking architecture with wire-speed forwarding of traffic and aggregated throughput of 1.28 Tbps on Ethernet ports.
 - Media access control (MAC) address learning: Automatic update, and support for up to 128,000 MAC addresses.

- Up to 128 IP interfaces per switch.
- Static and LACP (IEEE 802.3ad) link aggregation, up to 220 Gb of total uplink bandwidth per switch, up to 64 trunk groups, and up to 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 that are based on source/destination IP or MAC addresses, or both.
- Fast port forwarding and fast uplink convergence for rapid STP convergence.
- ▶ Availability and redundancy:
 - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy.
 - IEEE 802.1D STP for providing L2 redundancy.
 - IEEE 802.1s MSTP for topology optimization. Up to 32 STP instances are supported by a single switch.
 - IEEE 802.1w RSTP provides rapid STP convergence for critical delay-sensitive traffic, such as voice or video.
 - 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 4095 VLANs supported per switch, with VLAN numbers from 1 - 4095, where 4095 is used for management module's connection only.
 - 802.1Q VLAN tagging support on all ports.
 - Private VLANs.
- ▶ Security
 - VLAN-based, MAC-based, and IP-based 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 re-marking based on defined policies.
 - Eight WRR priority queues per port for processing qualified traffic.
- ▶ IP v4 Layer 3 functions:
 - Host management.
 - IP forwarding.

- IP filtering with ACLs, with 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), for up to 2048 entries in a routing table.
- Support for DHCP Relay.
- Support for IGMP snooping and IGMP relay.
- Support for PIM in PIM-SM and PIM-DM.
- ▶ IP v6 Layer 3 functions:
 - IPv6 host management (except for a default switch management IP address).
 - IPv6 forwarding.
 - Up to 128 static routes.
 - Support for OSPF v3 routing protocol.
 - IPv6 filtering with ACLs.
- ▶ Virtualization:
 - vNICs: Ethernet, iSCSI, or FCoE traffic is supported on vNICs.
 - 802.1Qbg Edge Virtual Bridging (EVB) is an emerging IEEE standard for allowing networks to become virtual machine (VM)-aware:
 - Virtual Ethernet Bridging (VEB) and Virtual Ethernet Port Aggregator (VEPA) are mechanisms for switching between VMs on the same hypervisor.
 - Edge Control Protocol (ECP) is a transport protocol that operates between two peers over an IEEE 802 LAN, which provides reliable and in-order delivery of upper layer protocol data units.
 - Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP) allows centralized configuration of network policies that persists with the VM, independent of its location.
 - EVB Type-Length-Value (TLV) is used to discover and configure VEPA, ECP, and VDP.
 - 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 that is based on the 802.1p priority value in each packet's VLAN tag.
 - Enhanced Transmission Selection (ETS) (IEEE 802.1Qaz) provides a method for allocating link bandwidth that is based on the 802.1p priority value in each packet's VLAN tag.
 - Data center Bridging Capability Exchange Protocol (DCBX) (IEEE 802.1AB) allows neighboring network devices to exchange information about their capabilities.
- ▶ Fibre Channel over Ethernet (FCoE):
 - FC-BB5 FCoE specification compliant.
 - Native FC Forwarder switch operations.
 - End-to-end FCoE support (initiator to target).
 - FCoE Initialization Protocol (FIP) support.

- ▶ Fibre Channel:
 - Omni Ports support 4/8 Gb FC when FC SFPs+ are installed in these ports.
 - Full Fabric mode for end-to-end FCoE or NPV Gateway mode for external FC SAN attachments (support for IBM B-type, Brocade, and Cisco MDS external SANs).
 - Fabric services in Full Fabric mode:
 - Name Server
 - Registered State Change Notification (RSCN)
 - Login services
 - Zoning
- ▶ Manageability:
 - Simple Network Management Protocol (SNMP V1, V2, and V3)
 - HTTP browser GUI
 - Telnet interface for CLI
 - SSH
 - Secure FTP (sFTP)
 - Service Location Protocol (SLP)
 - 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.
 - RMON agent to collect statistics and proactively monitor switch performance.
 - Port mirroring for analyzing network traffic that passes through a switch.
 - Change tracking and remote logging with syslog feature.
 - Support for sFLOW agent for monitoring traffic in data networks (separate sFLOW analyzer is required elsewhere).
 - POST diagnostic tests.

The following features are not supported by IPv6:

- ▶ Default switch management IP address
- ▶ SNMP trap host destination IP address
- ▶ Bootstrap Protocol (BOOTP) and DHCP
- ▶ RADIUS, TACACS+, and LDAP
- ▶ QoS metering and re-marking ACLs for out-profile traffic
- ▶ VMware Virtual Center (vCenter) for VMready
- ▶ Routing Information Protocol (RIP)
- ▶ Internet Group Management Protocol (IGMP)
- ▶ Border Gateway Protocol (BGP)
- ▶ Virtual Router Redundancy Protocol (VRRP)
- ▶ sFLOW

For more information, see *IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch*, TIPS0910, which can be found at this website:

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

2.3.4 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 does not have a management interface. It 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-through 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 2-25.



Figure 2-25 The IBM Flex System EN4091 10Gb Ethernet Pass-thru

The part number for the EN4091 10Gb Ethernet Pass-thru module is listed in Table 2-6. There are no upgrades available for this I/O module at the time of this writing.

Table 2-6 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-through 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 direct-attach cables (DAC). 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 *IBM Flex System EN4091 10Gb Ethernet Pass-thru Module*, TIPS0865:

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

2.3.5 Cables and transceivers for I/O modules

Table 2-7 lists supported cables and transceivers for IBM PureFlex System Ethernet I/O modules.

Table 2-7 Modules and cables supported in Ethernet I/O modules

Part number	Description	EN2092 1 GbE Switch	EN4093 10 GbE Switch	CN4093 10 GbE Switch	EN4091 10 GbE Pass- through
44W4408	10 GBase-SR SFP+ (MMFiber)	Yes	Yes	Yes	Yes
46C3447	10 GBase-SR SFP+ (MMFiber)	Yes	Yes	Yes	Yes
90Y9412	IBM SFP+ LR (SMFiber)	Yes	Yes	Yes	Yes
81Y1622	1000Base-SX SFP (MMFiber)	Yes	Yes	Yes	Yes
81Y1618	1000Base-T SFP	Yes	Yes	Yes	Yes
90Y9424	1000Base-LX SFP	Yes	Yes	Yes	Yes
49Y7884	IBM QSFP+ 40 Gbase-SR	No	Yes	Yes	No
90Y9427	1m IBM Passive DAC SFP+	Yes	Yes	Yes	No
90Y9430	3m IBM Passive DAC SFP+	Yes	Yes	Yes	No
90Y9433	5m IBM Passive DAC SFP+	Yes	Yes	Yes	No
49Y7886	1m 40 Gb QSFP+ to 4 x 10 Gb SFP+ Cable	No	Yes	Yes	No
49Y7887	3m 40 Gb QSFP+ to 4 x 10 Gb SFP+ Cable	No	Yes	Yes	No
49Y7888	5m 40 Gb QSFP+ to 4 x 10 Gb SFP+ Cable	No	Yes	Yes	No
90Y3519	10m IBM MTP Fiber Optical Cable	No	Yes	Yes	No
90Y3521	30m IBM MTP Fiber Optical Cable	No	Yes	Yes	No
49Y7890	1m QSFP+ to QSFP+ DAC	No	Yes	Yes	No
49Y7891	3m QSFP+ to QSFP+ DAC	No	Yes	Yes	No
95Y0323	IBM 1m 10 GBase Copper SFP+ TwinAx (Active)	No	No	No	Yes
95Y0326	IBM 3m 10 GBase Copper SFP+ TwinAx (Active)	No	No	No	Yes
95Y0329	IBM 5m 10 GBase Copper SFP+ TwinAx (Active)	No	No	No	Yes
81Y8295	1m 10 GE Twinax Act Copper SFP+ DAC (active)	No	No	No	Yes
81Y8296	3m 10 GE Twinax Act Copper SFP+ DAC (active)	No	No	No	Yes
81Y8297	5m 10 GE Twinax Act Copper SFP+ DAC (active)	No	No	No	Yes
44X1964	IBM 8Gb SFP+ Software Optical Transceiver	No	No	Yes	No

All Ethernet I/O modules are restricted to using the SFP, SFP+ and QSFP modules that are listed in Table 2-7; 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 adapters differ in physical port speeds (1 Gbps versus 10 Gbps) and in functions they support (base Ethernet connectivity versus converged networks and virtual NIC support).

The following Ethernet I/O adapter adapters are described in this section:

- ▶ IBM Flex System CN4054 10Gb Virtual Fabric adapter
- ▶ IBM Flex System EN4054 4-port 10 Gb Ethernet adapter
- ▶ IBM Flex System CN4058 8-port 10Gb Converged adapter
- ▶ IBM Flex System EN2024 4-port 1Gb Ethernet adapter
- ▶ IBM Flex System EN4132 2-port 10Gb Ethernet adapter
- ▶ IBM Flex System EN4132 2-port 10Gb RoCE adapter

2.4.1 IBM Flex System CN4054 10Gb Virtual Fabric adapter

The IBM Flex System CN4054 10 Gb 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 2-26.




Figure 2-26 IBM Flex System CN4054 10 Gb Virtual Fabric adapter

The ordering information for the CN4054 adapter is listed in Table 2-8.

Table 2-8 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 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
- ▶ Support 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.
- ▶ The following TCP 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
 - Data transmit and receive function is handled by adapter
 - Supported with iSCSI

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

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

2.4.2 IBM Flex System EN4054 4-port 10 Gb Ethernet Adapter

The IBM Flex System EN4054 4-port 10Gb Ethernet Adapter from Emulex enables the installation of four 10 Gb ports of high-speed Ethernet into an IBM Power Systems™ compute node. These ports interface with chassis switches or pass-through modules, which enables connections within and external to the IBM Flex System Enterprise Chassis.

The firmware for this four-port adapter is provided by Emulex, while the AIX® driver and AIX tool support are provided by IBM.

Figure 2-27 shows the IBM Flex System EN4054 four-port 10Gb Ethernet Adapter.




Figure 2-27 IBM Flex System EN4054 4-port 10 Gb Ethernet Adapter

The IBM Flex System EN4054 4-port 10Gb Ethernet Adapter has the following features and specifications:

- ▶ Four-port 10 Gb Ethernet adapter
- ▶ Dual-ASIC Emulex BladeEngine 3 controller
- ▶ Connection to 1 Gb or 10 Gb data center infrastructure (1 Gb and 10 Gb auto-negotiation)
- ▶ PCI Express 3.0 x8 host interface (The p260 and p460 support PCI Express 2.0 x8.)
- ▶ Full-duplex capability
- ▶ Bus-mastering support
- ▶ Direct memory access (DMA) support
- ▶ PXE support
- ▶ IPv4/IPv6 TCP and UDP checksum offload:
 - Large send offload
 - Large receive offload
 - Receive-Side Scaling (RSS)
 - IPv4 TCP Chimney offload
 - TCP Segmentation offload
- ▶ VLAN insertion and extraction
- ▶ Jumbo frames up to 9000 bytes
- ▶ Load balancing and failover support, including adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), teaming support, and IEEE 802.3ad
- ▶ Enhanced Ethernet (draft):
 - Enhanced Transmission Selection (ETS) (P802.1Qaz)
 - Priority-based Flow Control (PFC) (P802.1Qbb)
 - Data Center Bridging Capabilities eXchange Protocol, CIN-DCBX, and CEE-DCBX (P802.1Qaz)
- ▶ Supports Serial over LAN (SoL)
- ▶ Total Max Power: 23.1 W

For more information, see IBM Redbooks Product Guide *IBM Flex System CN4054 10Gb Virtual Fabric Adapter and EN4054 4-port 10Gb Ethernet Adapter*, TIPS0868, which is available at this website:

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

2.4.3 IBM Flex System CN4058 8-port 10Gb Converged Adapter

The IBM Flex System CN4058 8-port 10Gb Converged Adapter is an eight-port 10Gb converged network adapter (CNA) for Power Systems compute nodes that supports 10 Gb Ethernet and FCoE.

With hardware protocol offloads for TCP/IP and FCoE standard, the CN4058 eight-port 10Gb Converged Adapter provides maximum bandwidth with minimal usage of processor resources. This situation is key in IBM Virtual I/O Server (VIOS) environments because it enables more VMs per server, which provides greater cost savings to optimize return on investment (ROI). With eight ports, the adapter uses all of the capabilities of all Ethernet switches in the IBM Flex System portfolio.

Figure 2-28 shows the CN4058 8-port 10Gb Converged Adapter.




Figure 2-28 IBM Flex System CN4058 8-port 10 Gb Converged Network Adapter

The IBM Flex System CN4058 8-port 10Gb Converged Adapter includes the following features:

- ▶ Eight-port 10 Gb Ethernet adapter
- ▶ Dual-ASIC controller using the Emulex XE201 (Lancer) design
- ▶ PCIe Express 2.0 x8 host interface (5 GTps)
- ▶ MSI-X support
- ▶ IBM Fabric Manager support

The adapter includes the following Ethernet features:

- ▶ IPv4/IPv6 TCP and UDP checksum offload, Large Send Offload (LSO), Large Receive Offload, Receive Side Scaling (RSS), and TCP Segmentation Offload (TSO)
- ▶ VLAN insertion and extraction
- ▶ Jumbo frames up to 9000 bytes
- ▶ Priority Flow Control (PFC) for Ethernet traffic
- ▶ Network boot
- ▶ Interrupt coalescing
- ▶ Load balancing and failover support, including adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), link aggregation, and IEEE 802.1AX

The adapter has the following FCoE features:

- ▶ Common driver for CNAs and HBAs
- ▶ 3,500 N_Port ID Virtualization (NPIV) interfaces (total for adapter)
- ▶ Support for FIP and FCoE Ether Types
- ▶ Fabric Provided MAC Addressing (FPMA) support
- ▶ 2048 concurrent port logins (RPIs) per port
- ▶ 1024 active exchanges (XRIs) per port

Important: The CN4058 does not support iSCSI hardware offload.

The adapter supports the following IEEE standards:

- ▶ PCI Express base spec 2.0, PCI Bus Power Management Interface rev. 1.2, and Advanced Error Reporting (AER)
- ▶ IEEE 802.3ap (Ethernet over Backplane)
- ▶ IEEE 802.1q (VLAN)
- ▶ IEEE 802.1p (QoS/CoS)
- ▶ IEEE 802.1AX (Link Aggregation)
- ▶ IEEE 802.3x (Flow Control)
- ▶ Enhanced I/O Error Handling (EEH)
- ▶ 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)

To use the capabilities of the CN4058 adapter, I/O modules should be upgraded as follows to maximize the number of active internal ports:

- ▶ For CN4093, EN4093, and EN4093R switches
Upgrade 1 and 2 are required for the CN4093, EN4093, and EN4093R to use six ports on the adapter. If only Upgrade 1 is applied, only four ports per adapter are connected. If neither upgrade is applied, only two ports per adapter are connected.
- ▶ For the EN4091 Pass-thru
The EN4091 Pass-thru has only 14 internal ports and therefore supports only ports 1 and 2 of the adapter.
- ▶ For the EN2092
Upgrade 1 of the EN2092 is required to use four ports of the adapter. If Upgrade 1 is not applied, only two ports per adapter are connected.

Important: With the switches that are available for Flex System, at most six of the eight ports of the CN4058 adapter are connected.

The CN4058 supports FCoE to FC and FCoE targets.

The IBM Flex System CN4058 8-port 10Gb Converged Adapter supports the following operating systems:

- ▶ VIOS 2.2.2.0 or later is required to assign the adapter to a VIOS partition
- ▶ AIX Version 6.1 with the 6100-08 Technology Level Service Pack 3
- ▶ AIX Version 7.1 with the 7100-02 Technology Level Service Pack 3
- ▶ IBM i 6.1 is supported as a VIOS client
- ▶ IBM i 7.1 is supported as a VIOS client
- ▶ Red Hat Enterprise Linux 6.3 for POWER®, or later, with current maintenance updates available from Red Hat

- SUSE Linux Enterprise Server 11 Service Pack 2 with additional driver updates provided by SUSE

For more information, see IBM Redbooks Product Guide *IBM Flex System CN4058 8-port 10Gb Converged Adapter*, TIPS0909, which is available at this website:

<http://www.redbooks.ibm.com/abstracts/tips0909.html>

2.4.4 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 2-29.




Figure 2-29 IBM Flex System EN2024 4-port 1 Gb Ethernet Adapter

The ordering part number for the EN2024 adapter is listed in Table 2-9.

Table 2-9 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 *IBM Flex System EN2024 4-port 1Gb Ethernet Adapter*, TIPS0845, which is available at this website:

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

2.4.5 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 that are 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 2-30.




Figure 2-30 The EN4132 2-port 10Gb Ethernet Adapter for IBM Flex System

The ordering part number for the EN4132 adapter is listed in Table 2-10.

Table 2-10 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 *IBM Flex System EN4132 2-port 10Gb Ethernet Adapter*, TIPS0873, which is available at this website:

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

2.4.6 IBM Flex System EN4132 2-port 10Gb RoCE Adapter

The IBM Flex System EN4132 2-port 10Gb RoCE Adapter for Power Systems compute nodes delivers high bandwidth and provides RDMA over Converged Ethernet (RoCE) for low latency application requirements.

Clustered IBM DB2® databases, web infrastructure, and high frequency trading are just a few applications that achieve significant throughput and latency improvements, which results in faster access, real-time response, and more users per server. This adapter improves network performance by increasing available bandwidth while it decreases the associated transport load on the processor.

Figure 2-31 shows the EN4132 2-port 10Gb RoCE Adapter.




Figure 2-31 IBM Flex System EN4132 2-port 10Gb RoCE Adapter

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

- ▶ RDMA over Converged Ethernet (RoCE)
EN4132 2-port 10Gb RoCE Adapter, which is based on Mellanox ConnectX-2 technology, uses the InfiniBand Trade Association's RDMA over Converged Ethernet (RoCE) technology to deliver similar low latency and high performance over Ethernet networks. By using Data Center Bridging capabilities, RoCE provides efficient, low-latency RDMA services over Layer 2 Ethernet. The RoCE software stack maintains existing and future compatibility with bandwidth and latency-sensitive applications. With link-level interoperability in the existing Ethernet infrastructure, network administrators can use existing data center fabric management solutions.
- ▶ Sockets acceleration
Applications that use TCP/UDP/IP transport can achieve industry-leading throughput over InfiniBand or 10 GbE adapters. The hardware-based stateless offload engines in ConnectX-2 reduce the processor impact of IP packet transport, which allows more processor cycles to work on the application.
- ▶ I/O virtualization
ConnectX-2 with Virtual Intelligent Queuing (Virtual-IQ) technology provides dedicated adapter resources and ensured isolation and protection for virtual machines within the server. I/O virtualization with ConnectX-2 gives data center managers better server usage while it reduces cost, power, and cable complexity.

The IBM Flex System EN4132 2-port 10Gb RoCE Adapter has the following specifications (based on Mellanox Connect-X2 technology):

- ▶ PCI Express 2.0 (1.1 compatible) through an x8 edge connector with up to 5 GTps
- ▶ 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)
- ▶ 128 MAC/VLAN addresses per port
- ▶ RoHS-6 compliant

The adapter meets the following IEEE specifications:

- ▶ IEEE 802.3ae 10 Gigabit Ethernet
- ▶ IEEE 802.3ad Link Aggregation and Failover
- ▶ IEEE 802.3az Energy Efficient Ethernet
- ▶ IEEE 802.1Q, .1p VLAN tags and priority
- ▶ IEEE 802.1Qau Congestion Notification
- ▶ IEEE P802.1Qbb D1.0 Priority-based Flow Control
- ▶ IEEE 1588 Precision Clock Synchronization
- ▶ Jumbo frame support (10 KB)

The EN4132 2-port 10Gb RoCE Adapter supports the following operating systems:

- ▶ AIX V7.1 with the 7100-02 Technology Level, or later
- ▶ AIX V6.1 with the 6100-08 Technology Level, or later
- ▶ SUSE Linux Enterprise Server 11 Service Pack 2 for POWER, with current maintenance updates available from SUSE to enable all planned functionality
- ▶ Red Hat Enterprise Linux 6.3, or later

For more information, see *IBM Flex System EN4132 2-port 10Gb RoCE Adapter*, TIPS0913, which is available at this website:

<http://www.redbooks.ibm.com/abstracts/tips0913.html>



Layer 2 Overview

This chapter describes the Layer 2 fundamental networking protocols and terminology that is related to topics that are covered in this book.

3.1 Layer 2 Network protocols and technologies

For the scope of this book, this section provides information about the following topics:

- ▶ Basic Frame Forwarding
- ▶ Virtual local area network (VLAN) and Tagging
- ▶ Spanning Tree: STP, RSTP, MST, VSTP
- ▶ Link Aggregation
- ▶ Virtual Link Aggregation (VLAG)
- ▶ Juniper Virtual Chassis (VC)
- ▶ Link Layer Discovery Protocol (LLDP)
- ▶ IBM Layer 2 fail over

3.1.1 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:

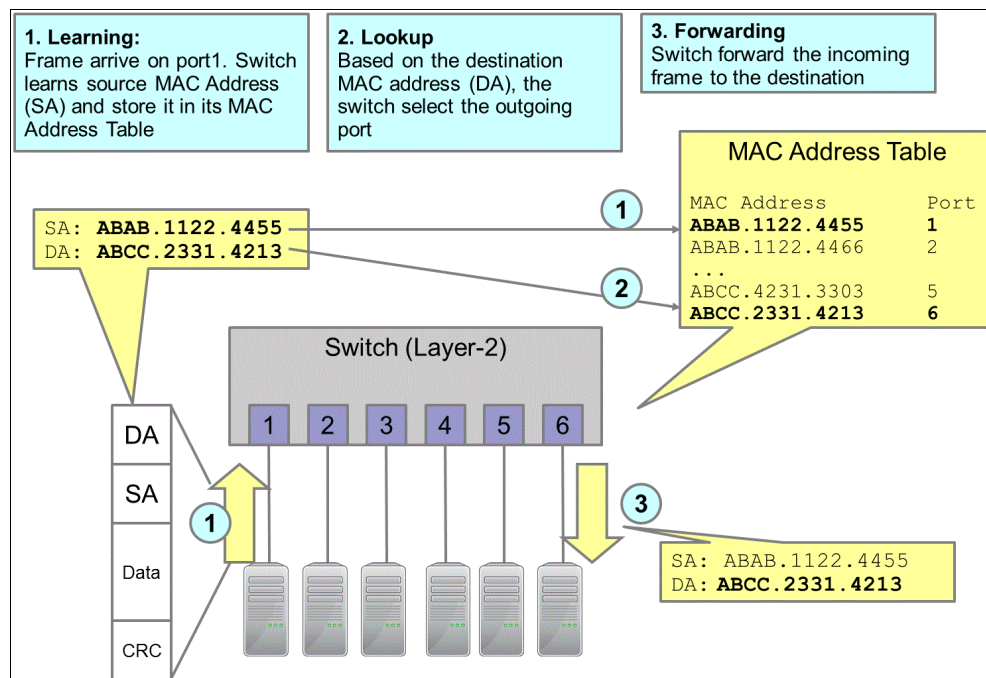


Figure 3-1 Frame forwarding principle

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

- ▶ Learning

Ethernet Frame arrives on port1. Switch learns source MAC Address (SA) and stores it in its MAC Address Table that this address belongs to port 1.

- ▶ Lookup

Based on the destination MAC address (DA), the switch lookup in its MAC address table and select the outgoing port (port 6)

- Forwarding

The switch forwards the Ethernet frame to the destination MAC address via port 6.

Should the switch not know the destination address; it forwards the packet on all ports except the port from which it was received.

3.1.2 Virtual local area network and tagging

A virtual local area network (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 shown in Figure 3-2.




Figure 3-2 Virtual local area network

As shown in Figure 3-3 on page 52, the following methods are available for sharing VLANs across devices:

- The use of dedicated cabling for each VLAN to keep isolated
This 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.
- Marking packets through tagging, so that a single interconnect can be used to transport data for various VLANs
This method is highly scalable because only a single link is required to provide connectivity to many VLANs, which provides for better usage of the link capacity when VLAN traffic is not uniform.




Figure 3-3 VLAN Tagging

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 another 4 bytes of information to be added to each Ethernet frame. A frame including this extra information is known as a *tagged frame*.

The 4-byte tag includes the following 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 that is 3-bits long to allow a priority value of eight different values to be included in the tag and has 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 of which VLAN the frame is a member, with 4096 different VLANs possible.

3.1.3 Spanning tree

Because of the history of LANs and Ethernet, there are some shortcomings in the protocol. In particular, Ethernet was not designed to use frame forwarding. Therefore, the frame format does not include a hop count field (or time-to-live) which would allow 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 simplest approach to prevent looping packets is to create a network topology where frames with a given target can take only one path on each individual switch element. For Ethernet, the tree topology was chosen, which is the simplest topology that ensures this requirement. Bridges and Switches were enhanced to support a topology configuration protocol called Spanning Tree Protocol (STP).

STP provides Layer 2 loop prevention by deactivating redundant routes between network elements. Over the years, it was further enhanced and is now used in different forms, such as:

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

STP (IEEE802.1D)

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

The following types of BPDUs are available:

- Configuration

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)

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)

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 elect a root bridge, identify root ports for each switch, identify designated ports for each physical LAN segment, and 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.

Root bridge election: The root bridge election is an 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.

Rapid Spanning Tree Protocol (IEEE802.1w)

Rapid Spanning Tree Protocol (RSTP) provides better re-convergence time than the original STP. RSTP identifies certain links as point-to-point. When a point-to-point link fails, the alternative link can make the transition to the forwarding state.

RSTP adds new bridge port roles to speed convergence that follow a link failure. RSTP bridge port includes the following 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 is connecting to this port.
- Alternative port: Provides an alternative root port.
- Backup port: Provides an alternative designated port.

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

Multi-instance Spanning Tree Protocol (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 the same spanning topology while many links in the network might 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 a Multi-instance Spanning Tree Protocol (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 functionality facilitates better load sharing across redundant links.

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

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

VLAN Spanning Tree Protocol

In Juniper Networks environment, if compatibility to Cisco proprietary PVRST protocol is required, VLAN Spanning Tree Protocol (VSTP) can be configured. VSTP allows EX series switches to run one or more STP or RSTP instances for each VLAN on which VSTP is enabled. For networks with multiple VLANs, this enables more intelligent tree spanning because each VLAN can have interfaces enabled or disabled depending on the paths that are available to that specific VLAN.

By default, VSTP runs RSTP, but you cannot have both stand-alone RSTP and VSTP running simultaneously on a switch. Also, usage of STP can be forced if the network includes old switches. VSTP can be enabled for up to 253 VLANs.

3.1.4 Dynamic Link Aggregation Control Protocol

Link Aggregation Control Protocol (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 do not 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 interconnected by the aggregated link uses a hashing algorithm to determine on which of the member links frames will be transmitted. 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.

3.1.5 Virtual Link Aggregation Groups

Virtual Link Aggregation Groups (VLAGs) are 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 that is 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, which block redundant uplink paths. This has the unwanted consequence of reducing the available bandwidth between the layers. In addition, STP can be slow to resolve topology changes that occur during a link failure and result in considerable MAC address flooding.

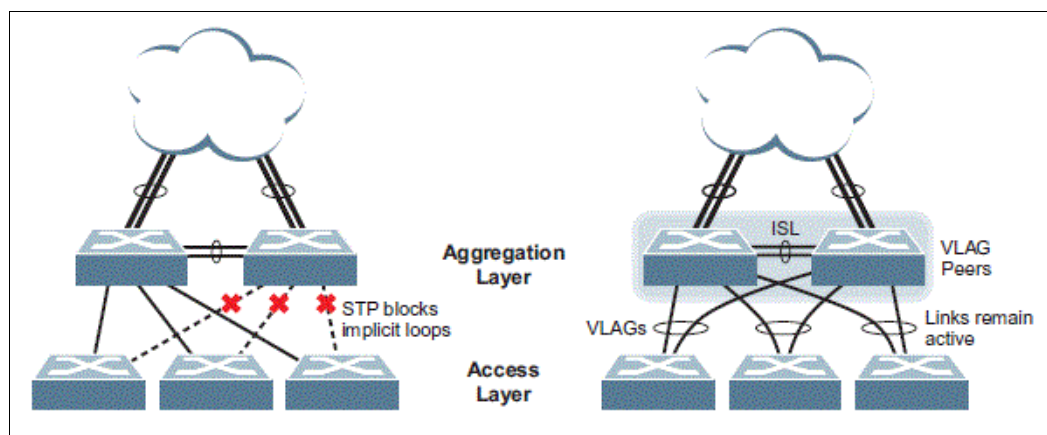


Figure 3-5 Spanning-tree versus VLAG

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.

3.1.6 Juniper Virtual Chassis

Juniper Networks Virtual Chassis (VC) technology is a feature of the Juniper Networks EX series switches that allow the interconnection and operation of switches as a unified, single, high-bandwidth device. Up to 10 switches can be interconnected via dedicated VC ports on each device or through optional uplink module ports that are configured as VC ports, with a combined backplane bandwidth of up to 128 Gbps.

Solutions that use the EX4200 line with VC technology combine the scalability and compact form factor of stand-alone switches with the high availability, high backplane bandwidth characteristics, and high port densities of traditional chassis-based switches. VC configurations enable economical deployments of switches that deliver network availability in locations where installation might otherwise be cost prohibitive or physically impossible.

In a VC configuration, all member switches are managed and monitored as a single logical device. This approach simplifies network operations, allows the separation of placement and logical groupings of physical devices, and provides efficient use of resources. The VC solution offers the same Routing Engine redundancy features as other Juniper Networks chassis-based switches and routers, including graceful Routing Engine switchover (GRES) for hitless failover.

For resiliency and redundancy, the VC configuration includes a master and a backup switch. The master and backup can be elected dynamically, or statically configured so as to position them in pre-determined locations within the VC (for example, master at top and backup at bottom).

Each remaining switch serves as a line card switch but is ready to be selected as a backup switch if the master or backup switch fails. Switches also can be selectively prioritized in a VC configuration to assign master and backup roles, and to determine the order in which the remaining switches are elected if the master and backup switches fail.

Management of the VC configuration is performed through the master switch. A Virtual Management Ethernet (VME) interface allows remote management by connecting to the out-of-band management port of any member switch through a single IP address.

In addition, the VC configuration uses a single Juniper Networks Junos operating system image file and a single configuration file. The Junos OS of all member switches in a VC configuration can be upgraded simultaneously from the master switch with a single command.

VC technology includes the following key benefits:

- ▶ Simplified overall system maintenance and management through a single management interface.
- ▶ Pay-as-you-grow scalability, from 24 - 480 10/100/1000 Mbps ports.
- ▶ Extension of the VC configuration by up to 50 km with redundant fiber links between physical devices.
- ▶ Redundant, internal hot-swappable power supplies.
- ▶ Hot-swappable fan tray with redundant blowers.
- ▶ Consistent modular Junos OS control plane feature implementation.
- ▶ Dual routing engines with GRES.

Figure 3-6 shows a typical VC configuration that uses five EX4200 switches.



Figure 3-6 Juniper Virtual Chassis

3.1.7 Link Layer Discovery Protocol

Link Layer Discovery Protocol (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).

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

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

3.1.8 Layer 2 fail over

The primary application for Layer 2 fail over 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 also can be configured as Active-Active pairing, in which both NICs would be 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 shows 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 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.

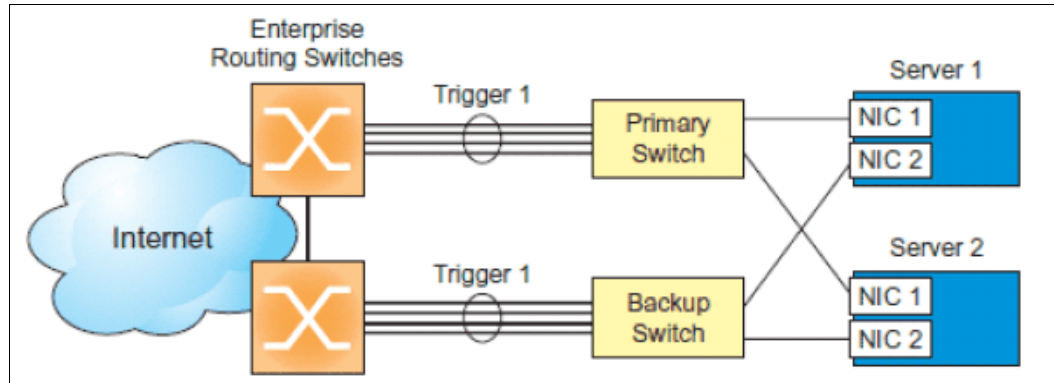


Figure 3-7 Basic Layer 2 Failover

This feature also is 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 provides an overview of Layer 3.

4.1 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, which frees up the upstream router to handle only in-bound and out-bound traffic. Because IBM System networking switches use application-specific integrated circuits (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, saves the network administrator from having to reconfigure each endpoint with new IP addresses and is achieved without any loss of performance.

4.1.1 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 next-hop in the network.

4.1.2 Default gateways

IBM System Networking switches can be configured with up to four of the following 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 virtual local area networks (VLANs) on the switch. The use of multiple gateways for the same IP address route also is used to configure equal-cost multi-paths (ECMPs).

4.1.3 ECMP static routes

ECMP is a forwarding mechanism that can be used to equally distribute loads 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 performs periodic health checks (ping) on each ECMP gateway. If a gateway fails, it is removed from the routing table.

4.1.4 Routing Information Protocol v2

The goal of any routing protocol is to populate a device's routing table with valid, loop-free routes. Routing protocols are essential in today's large and complex networks. Routing Information Protocol v2 (RIPv2) enabled routers share and track available routes. RIPv2 is a distance vector protocol. Routers that use distance vector protocols do not have knowledge of the entire path to a destination, or the topology of a network. Instead, they only 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 is easy to configure; however, because routers must share their entire routing tables regularly, RIPv2 does not scale well. Both System Networking switches and Juniper switches support RIPv2. RIPv2 is an IETF standard.

For more information, see this website:

<http://www.ietf.org/rfc/rfc2453>

4.1.5 Open Shortest Path First for IPv4

Open Shortest Path First (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. For more information, see this website:

<http://www.ietf.org/rfc/rfc2328>

OSPF networks scale well because they can be logically divided into routing areas. Each area is identified by a 32-bit number that is 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 other area must be connected to area 0.0.0.0.

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 (as shown in Figure 4-1):

- ▶ **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 with more capabilities. Routes that are originating from within the NSSA can be propagated to adjacent transit and backbone areas. External routes from outside the autonomous systems (AS) 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.

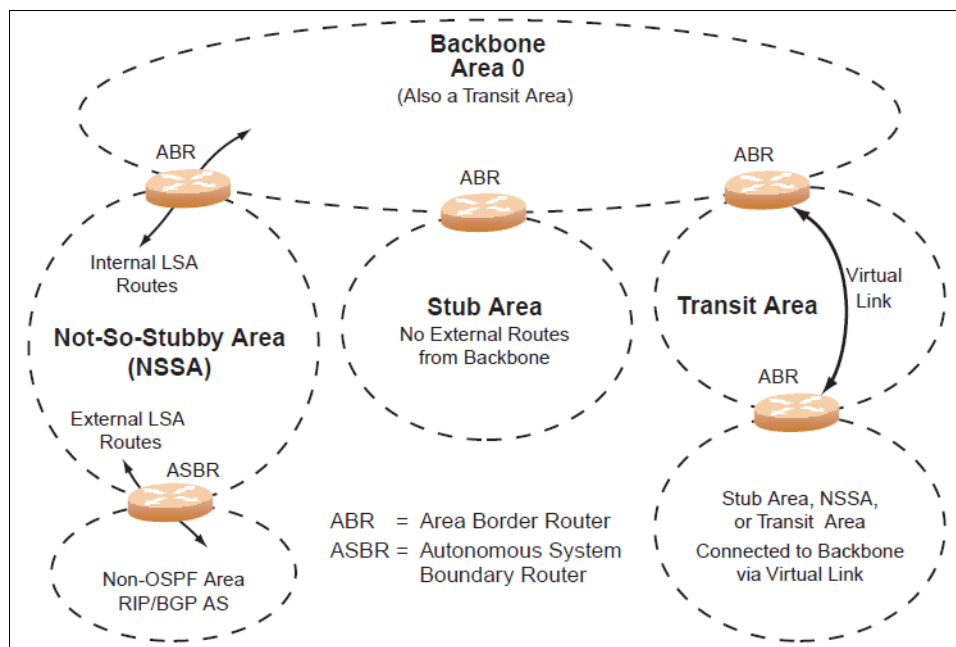


Figure 4-1 OSPF area types

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 other's 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 and area number. Area stub-flags 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.

Link-State Database

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 Link-State Database (LSDB) that describes 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 also can 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, which allow OSPF routing devices to be aware of devices that otherwise participate in OSPF (because they do not support it or 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 that is removed from the route table, an update message that contains the route to withdraw is sent if the route already was sent to an adjacent neighbor.

OSPF router types

As shown in Figure 4-2 on page 63, 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 that are 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.

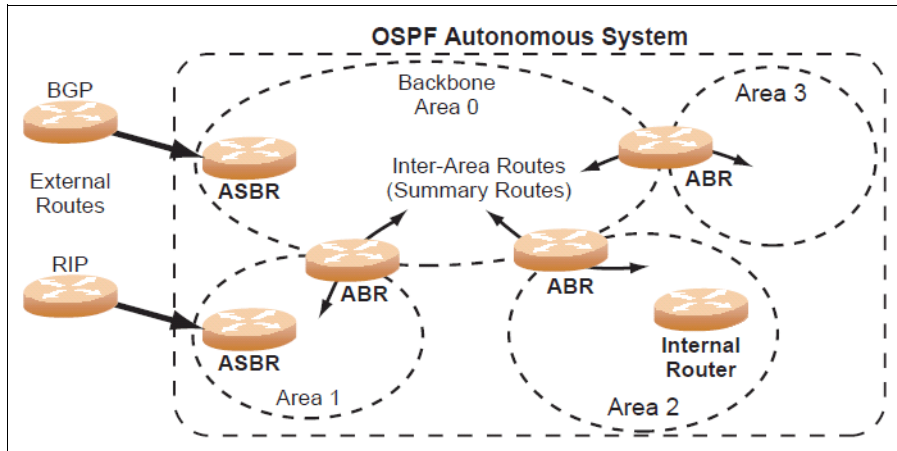


Figure 4-2 OSPF router types

Shortest Path First

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

4.1.6 Border Gateway Protocol

Border Gateway Protocol (BGP) version 4 is an open standard, after several reiterations it is now published as RFC 4271. For more information, see this website:

<http://www.ietf.org/rfc/rfc4271>

BGP is the lifeblood of the internet. It exchanges routing information between all the major Internet Service Providers (ISPs) and it is an Exterior Gateway Protocol (EGP), which means it exchanges routing information between AS. This is different from Interior Gateway Protocols (IGP) such as RIPv2, EIGRP, and OSPF which support routing within an AS.

Autonomous systems

An AS is defined as a connected group of one or more IP prefixes run by one or more network operators which have a single and clearly defined routing policy. In real terms, this tends to be an ISP's network 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 allocates. The ASN is a globally unique identifier. BGP keeps a list of every AS that a path passes through, which 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 goes to great extremes to ensure that only one is actually used. There is a list of weighted steps that are used to determine which routes are preferred and which routes removed.

An iBGP is a type of internal routing protocol you can use to perform active routing inside your network. It also carries AS path information, which is important when you are an ISP or performing 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 that is shown between the two routers in AS 20 is not present (see Figure 4-3), the top router does not learn the route to AS 50. Also, the bottom router does not learn the route to AS 11, even though the two AS 20 routers are connected through the IBM System Networking switch.




Figure 4-3 Diagram showing iBGP and eBGP

Typically, an AS has one or more *border routers*, which are peer routers that exchange routes with other autonomous systems, 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 that is destined for any address in 192.204.4.0/24, you know how to carry that data to its destination.

4.1.7 IPv6

The IPv6 protocol is an IETF standard. For more information, see this website:

<http://www.ietf.org/rfc/rfc2460>

The IPv4 protocol was originally developed in the early 1980s. IPv6 has several improvements over IPv4 and resolves some unforeseen issues.

Address size

Clearly the obvious one to start with 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

Address usage

IPv6 design aim was not to create only a massive pool of IP addresses, but to 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. Here, 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 three bits of an IPv6 address are reserved to define the type of IPv6 address that is used.

Address Hierarchy

IPv6 has a hierarchy of address block allocation. The 64 bits that are used for the Subnet identifier are broken down into the following levels:

- ▶ Top Level Aggregation Identifier (TLA ID, 13 bits) is assigned to major service providers.
- ▶ Next Level Aggregation Identifier (NLA ID, 24 bits) is assigned to minor service providers.
- ▶ Site Level Aggregation Identifier (SLA ID, 16 bits) that are assigned to an organization's and company's 16 bits provides 645,535 subnets.

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

Address auto-configuration and 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 in "Address usage" on page 65, the link-local address is: fe80::202:55ff:fe54:af3a/64. All devices on a subnet can communicate 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 includes the following advantages over IPv4 systems that are used:

- ▶ A host always has the same unique global and link-local IPv6 address.
- ▶ Address allocation is built into the system, there are no separate DHCP server issues.
- ▶ Host and server devices can have their addresses auto-configured.
- ▶ Router advertisements that are received at the same time also give 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, printing.

4.1.8 Open Shortest Path First for IPv6

Open Shortest Path First (OSPFv3) works in a similar way to OSPFv2. It is defined as an IETF standard. For more information, see this website:

<http://www.ietf.org/rfc/rfc5340>

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, it is necessary to define an arbitrary 32-bit router ID that is expressed in dotted decimal (IPv4) format.

4.1.9 Virtual Router Redundancy Protocol

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. To achieve this goal usually requires redundancy for all vital network components.

First Hop redundancy protocols (FHRP) protect the default gateway address that is used on a subnetwork by allowing two or more routers to provide backup for that address. In the event of failure of the active router, a backup router takes over the function of that address. FHRP protocols also are applied to other services that require redundancy for a single IP address. Virtual Router Redundancy Protocol (VRRP) is a widely implemented FHRP protocol. It is an open standard defined in RFC 2338. For more information, see this website:

<http://www.ietf.org/rfc/rfc2338>

Some modifications are defined in RFC 3768. For more information, see this website:

<http://www.ietf.org/rfc/rfc3768>

Table 4-2 shows the important aspects of the VRRP protocol.

Table 4-2 Table highlighting the important aspects of the VRRP protocol

Protocol	VRRP
Availability	Open standard supported by Juniper and IBM System Networking devices
RFC	2338 and 3768
Virtual IP address	Can use same IP address as physical port
Master/backups	One master, one standby, all others listening
Hello packets	Multicast ip 224.0.0.18
MAC address	00-00-5E-00-01-XX where XX is the virtual group ID
Failback configuration	Master claims Virtual IP address when available

The virtual router consists of a user-configured virtual router ID (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 00-00-5E-00-01. 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) and 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 router or 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 that is addressed to it. As the backup router is using the same MAC address, hosts do not need to send ARP requests and packets are processed with the minimum of disruption.

Configuration tip: It is a recommended best practice to configure all VRRP options before the virtual IP address is configured. This minimizes disruption and state changes in the network.

A backup router can stop receiving advertisements for one of two reasons: the master can be down or all communications links between the master and the backup can be down. If the master fails, it is desirable for the backup (or one of the backups, if there is more than one) to become the master.

Two masters: If the master is healthy but communication between the master and the backup fails, there are 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.

Active-active redundancy

In an active-active configuration (as shown in 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.

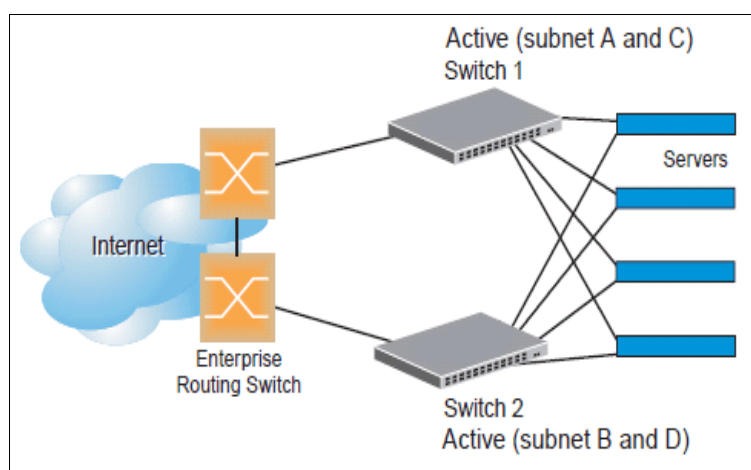


Figure 4-4 VRRP Configured as active-active

Although this example shows only two switches, there is no limit to the number of switches that are 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.

VRRP HA with vLAGs

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

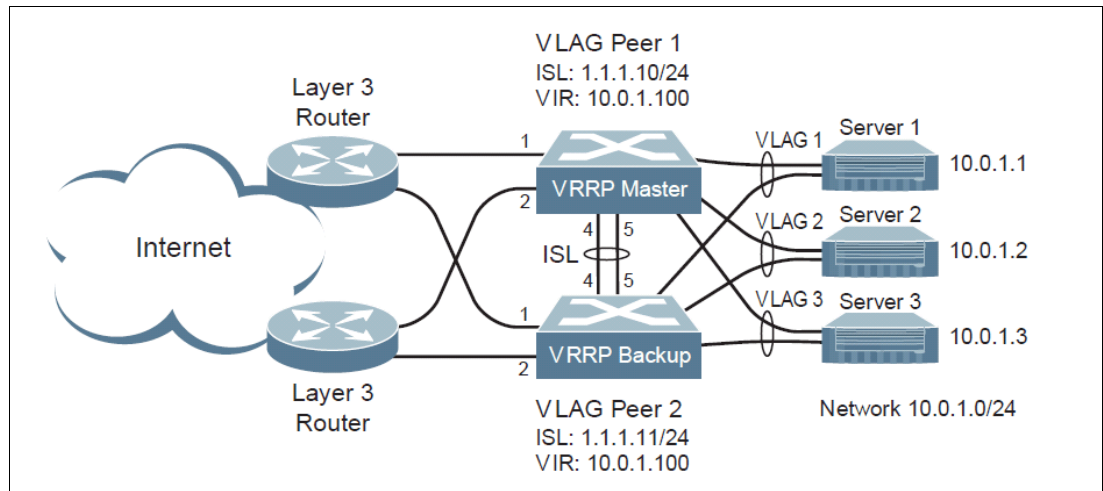


Figure 4-5 Active-active VRRP with vLAG



Connecting IBM PureFlex System to a Juniper Network

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, best practice, industry-standard interoperability use cases with an upstream Juniper network are described here.

Important: This information applies to PureFlex Express, Standard, Enterprise, and Open Choice configurations.

This chapter includes the following topics:

- ▶ Introduction
- ▶ High availability overview
- ▶ Fully redundant with Virtualized Chassis technology
- ▶ Fully redundant with traditional spanning tree protocol
- ▶ Fully redundant with OSPF

5.1 Introduction

The use cases that are described in this chapter were selected primarily based on input from IBM System Networking Consulting Engineers as to what was observed most often in the field during customer engagements.

Important: While these implementation scenarios were tested and verified to be compatible with an upstream Juniper network in a lab environment, these are not the only design options that are available to the network architect, and should be used as general guidance only. Customers should consult with their IBM Account Representative to work with our Worldwide System Networking Consulting Engineers for a more in-depth design discussion should a unique topology be required.

To facilitate the explanation of the recommended scenarios in the following sections, we must first review traditional, highly available network implementations and describe their own unique characteristics for a background framework as to why these scenarios were 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. Our focus is to provide HA 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, but some are specific to the IBM Flex System Enterprise Chassis. We review 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 network interface controllers (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 (on switches and the node adapters), the midplane, and the I/O modules.

The first step in achieving HA 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 (ToR) 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, it is necessary to consider 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 that are 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.

We describe several of these features next.

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.

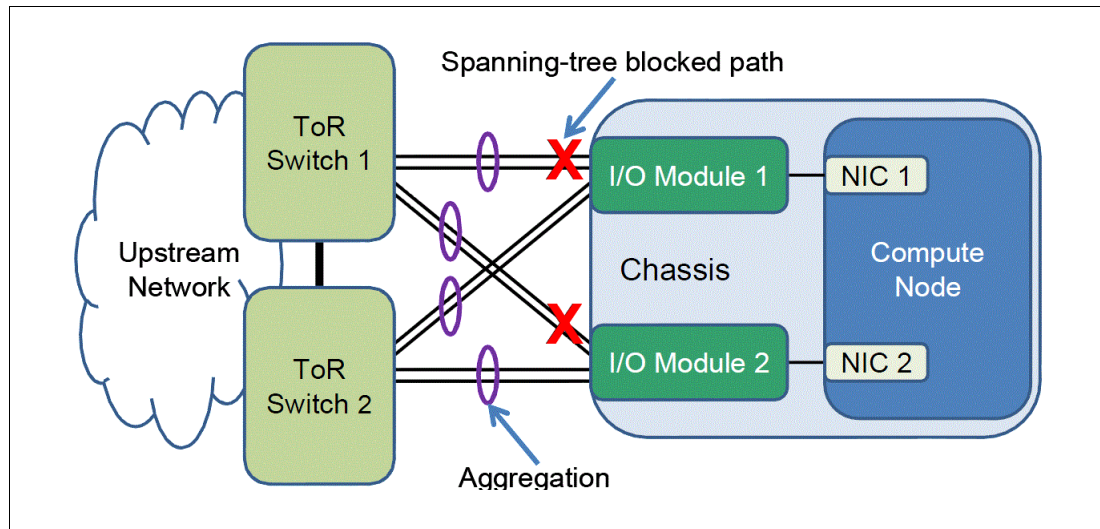


Figure 5-1 Looped and blocking design, no host NIC teaming

The looped and blocking design shows each I/O module in the Enterprise Chassis with two direct aggregations to a pair of upstream ToR switches. The specific number and speed of the external ports that are used for link aggregation in this and other designs that are shown in this section depend on the redundancy and bandwidth requirements of the customer. This topology is a bit complicated and is suggested for environments in which hosts need network redundancy, but they are not themselves performing 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, a recommended good L2 design is to pursue 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 take on the Looped and blocking design that is shown in Figure 5-2 is the Non-looped, single upstream device HA design.

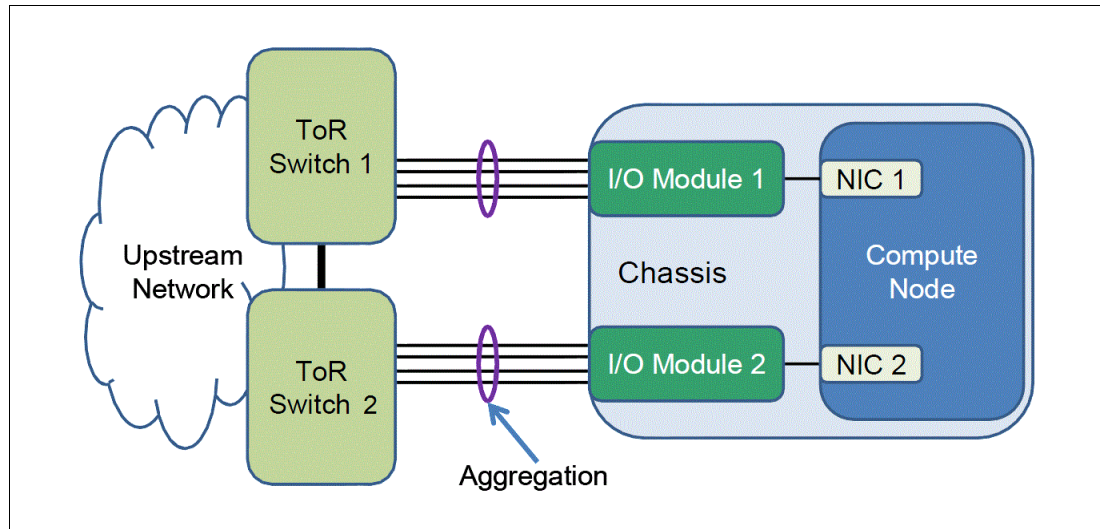


Figure 5-2 Non-looped, single upstream device design, with host NIC teaming

Figure 5-2 shows each I/O module in the Enterprise Chassis directly connected to a single ToR switch through aggregated links. The use of this topology is suggested 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 and bonding also are used to fail the traffic over to the other NIC in the team, which ensures 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 for 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, as shown in Figure 5-3 on page 73.

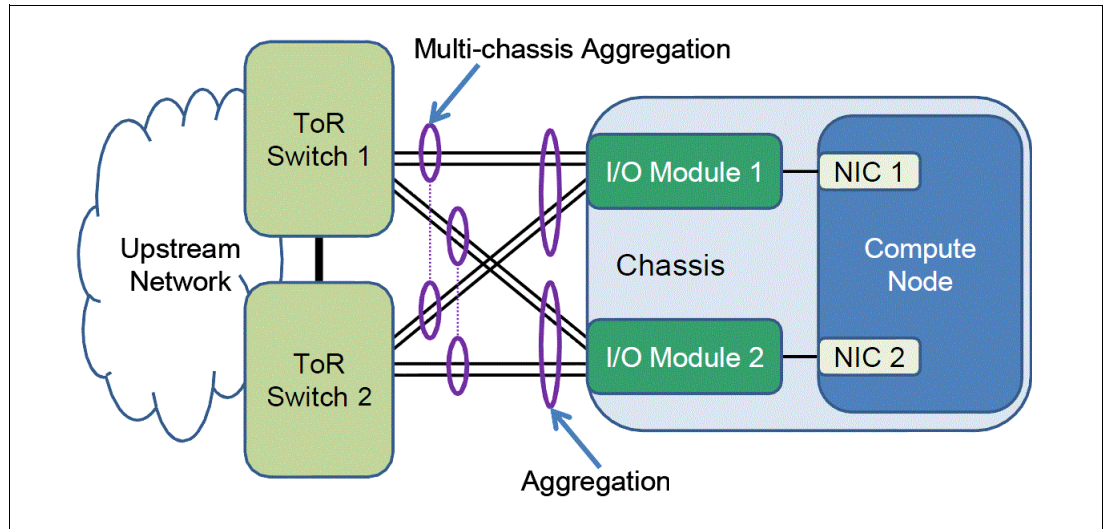


Figure 5-3 Non-looped, multiple upstream devices design

The non-looped, multiple upstream devices design brings the best of both the looped and blocking design and the non-looped, single upstream device design in a robust, stable implementation, which is suitable for use with hosts that have 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 functionality, including IBM System Networking products. The I/O modules in our implementation scenarios deploy the IBM Virtual Link Aggregation Group (vLAG) technology to the upstream ToR switch infrastructure to appear 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.

Next, we examine some recommended implementation scenarios.

5.3 Fully redundant with Virtualized Chassis technology

This implementation scenario incorporates switch virtualization features that allow a downstream switch the ability to be connected to two upstream, virtualized switches through the means of aggregated links. Inter-Switch Links (ISLs) between the same or similar products on the aggregation or access-layer provide a loop-free design that is 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 and internally prevent implicit loops. This design is recommended for customers that want to use a best-practice implementation on a Juniper network using next generation networking features, such as Juniper Virtual Chassis (VC) or Multichassis Link Aggregation Groups (MC-LAG) technology.

This approach includes 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

5.3.1 Components used

Two of each of the following components are used:

- ▶ Juniper EX4500-40F
- ▶ IBM G8264 RackSwitch
- ▶ IBM Flex System Fabric EN4093 10Gb Scalable Switch

5.3.2 Network topology and physical setup

Figure 5-4 on page 75 shows the network topology for the fully redundant scenario with virtualization technology (VC/vLAG).

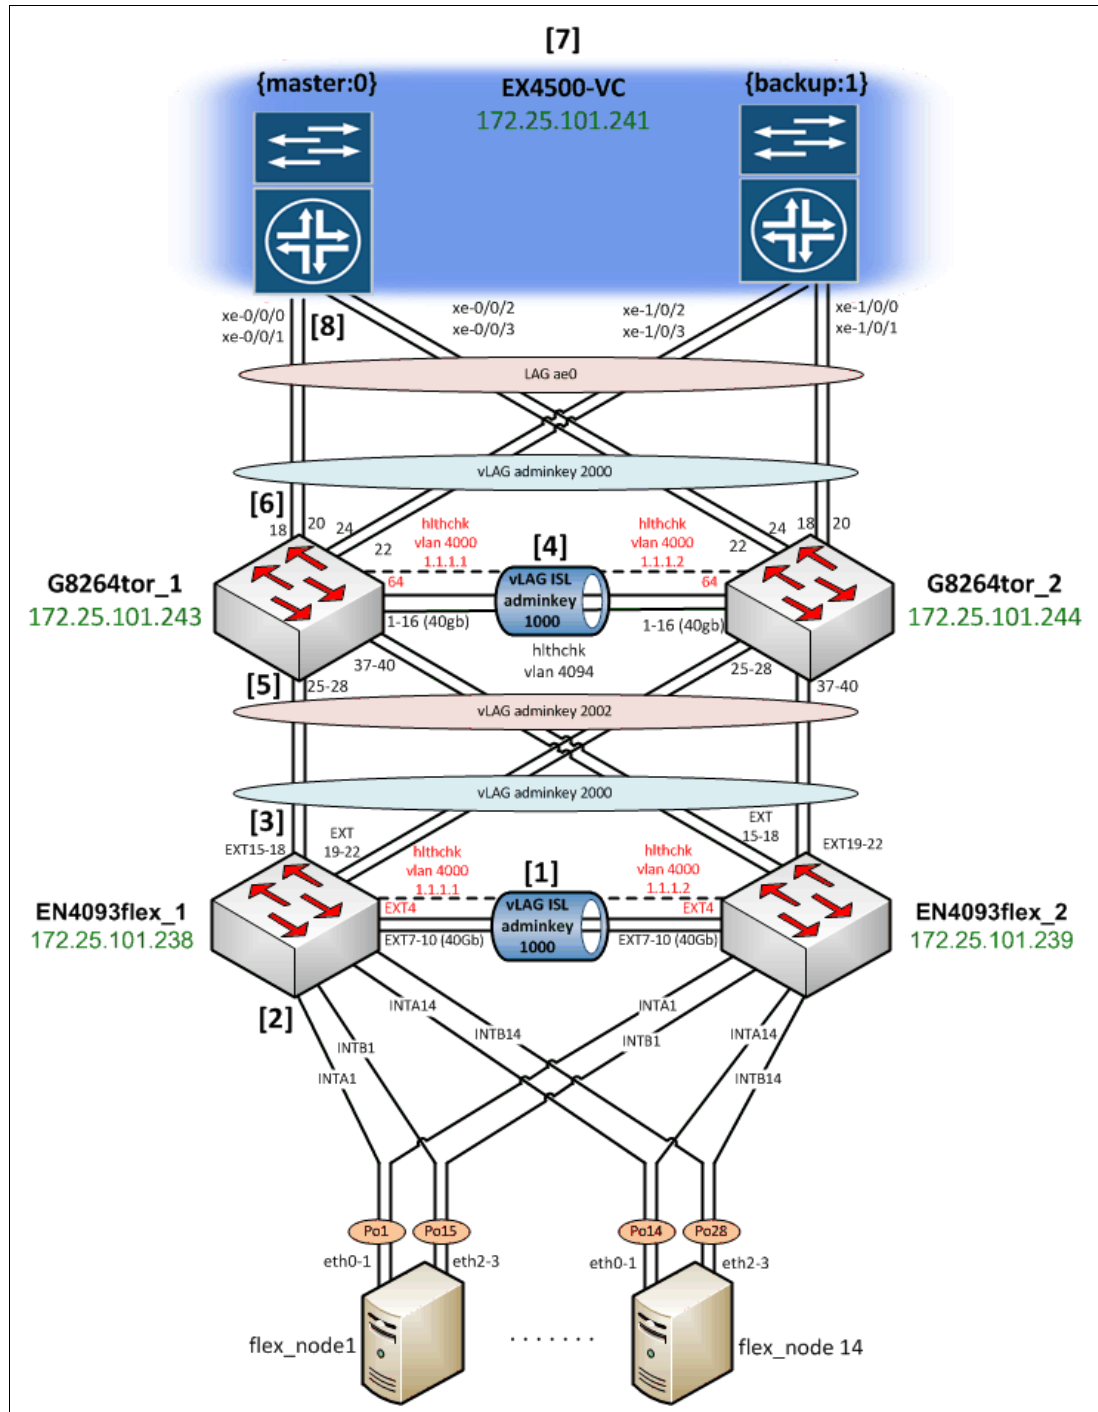


Figure 5-4 Network Topology diagram for fully redundant scenario with VC/vLAG

Start by verifying the physical cabling between the EN4093 switches and G8264s. In our lab environment, we used four IBM QSFP+ DAC Break Out Cables from the EN4093 switches to the upstream G8264s. This requires that the EN4093 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 160Gb ISL between the G8264 switches and 10 Gb SFP+ DAC cables were used for all other connections in the diagram.

5.3.3 EN4093 flex_1 configuration

We begin the implementation of this scenario on the IBM Flex System Fabric EN4093 switches, working our way northward in Figure 5-4 on page 75. Each step provides the commands necessary and are reflective of the numbering schema in the diagram to aid the user in what is configured.

General configuration

Complete the following steps to set up the general configuration:

1. Create the ISL Healthcheck, ISL data, and Data VLANs, as shown in Example 5-1, and give them descriptive names, assign them to spanning-tree groups, and enable them. You can elect to have the switch create STP instances for you; we chose to manually create them instead.

Example 5-1 Create ISL hlthchk, DATA and ISL VLANs on EN4093 flex_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 & Data VLANs in Example 5-2 so that we can verify connectivity between the various pieces of equipment when the configuration is verified. 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, with the last octet borrowed from the network diagram's Management address to quickly aid in the identification of which piece of equipment we are verifying connectivity to.

Example 5-2 Create IP interfaces and assign VLANs and IP addresses on EN4093 flex_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 the ISL between EN4093 flex switches

Complete the following steps to configure the ISL between EN4093 flex switches:

1. Configure the eventual ISL in Example 5-3 between the EN4093 switches by configuring them to have a default (untagged) VLAN of 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 5-3 Initial ISL configuration on EN4093 flex_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
```

2. Create the dedicated health check VLAN and physical interface in Example 5-4 to be used for heartbeats between the EN4093 switches. We chose EXT4 as a dedicated interface and VLAN 4000 to serve as the health check for the ISL.

Example 5-4 Create vLAG health check on EN4093 flex_1

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

3. Disable STP between the EN4093 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) and reference the LACP key that configured in the previous step.

Example 5-5 Disable STP and activate ISL vLAG on EN4093 flex_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

Complete the following steps to configure downstream internal node ports.

1. Configure 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 EN4093 flex

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

2. For redundancy, we created two port channels on each of the 14 nodes. Each port-channel aggregates two ports, one from each EN4093 flex switch. Port channels 1 - 14 to match the A internally labeled ports, and port channel 15 - 28 to match the B ports, as shown in Example 5-7.

Example 5-7 Node-facing port channel creation and vLAG activation, on EN4093 flex_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

Complete the following steps to configure upstream, G8264t or facing ports and Layer 2 failover:

1. 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 EN4093 flex_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
```

2. Activate the vLAG feature for the upstream EN4093 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 bundling ports EXT15-22 together as one.

Example 5-9 Activating the upstream G8264tor-facing vLAG on EN4093 flex_1

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

3. Enable Layer 2 failover as shown in Example 5-10, which effectively shuts down the links to the compute nodes should the uplinks for the EN4093 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, which in our case is connected to the other EN4093 switch in the Enterprise Chassis and ensures that redundancy is maintained.

Example 5-10 Enable Layer 2 failover for the compute nodes on EN4093 flex_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 EN4093_flex2 on the other I/O module. The only difference between the EN4093 flex_1 switch and EN4093 flex_2 switch is the vLAG health check peer address and the Data and ISL hlthchk vlan ip addresses. To verify EN4093 flex switch configuration, run the show commands that are described 5.3.6, “Verification and show command output” on page 85.

5.3.4 G8264tor_1 configuration

Next is the configuration of the RackSwitch G8264.

General configuration

Complete the following steps to set up the general configuration:

1. Begin by creating 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 Create 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 that is shown in Figure 5-4 on page 75. IP gateway 4 is the upstream router interface for our 172 management network.

Example 5-12 Create IP interfaces

```
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 the ISL between G8264tor switches

Complete the following steps to configure the ISL between G8264tor switches:

1. Configure the ISL between the G8264 switches, as shown in Example 5-13. Make the default (untagged) VLAN 4094 (ISL). Assign 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
```

2. 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, which references the LACP key that was configured in the previous step.

Example 5-14 Disable STP and activate 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 EN4093 flex-facing ports

Complete the following steps to configure downstream EN4093 flex-facing ports:

1. Configure the downstream EN4093 flex-facing ports, as shown in Example 5-15 on page 83. Make the default (untagged) VLAN 4092 (data vlan), with 802.1q tagging enabled and tag the PVID. Add the ability for all member ports to be on VLAN 4092 and use an LACP key of 2002 to bundle the ports together.

Example 5-15 Configure downstream EN4093 flex-facing ports

```
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
```

2. Activate the vLAG feature for the downstream EN4093 flex-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 Activate downstream EN4093 flex-facing vLAG on G8264tor_1

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

Configuring upstream EX4500-VC facing ports

1. Configure the upstream EX4500-VC facing ports that are 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 Configure upstream EX4500-VC facing ports

```
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
```

2. Activate the vLAG feature for the ports bundled with LACP key 2000, which the EX4500-VC sees as a single, virtualized entity, as shown in Example 5-18.

Example 5-18 Activate the upstream EX4500-VC facing vLAG, adminkey 2000 on G8264tor_1, step [6] continued

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

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 the G8264tor switch configuration, run the show commands that are described in 5.3.6, “Verification and show command output” on page 85.

5.3.5 Juniper EX4500-VC switch configuration

We now configure the Juniper EX4500-VC core switch.

Configuring VC on Juniper switches

Complete the following steps to configure VC on Juniper switches:

1. Interconnect the member switches by using the dedicated Virtual Chassis Ports (VCPs) that are on the rear side of the switch. Power on only the switch that you plan to use as the master switch (for management purposes, we recommend arranging the switches in member ID sequence from top to bottom).
2. Configure mastership priority for the master, disable the split, and merge features (recommended for a two-member VC), as shown in Example 5-19.

Example 5-19 Virtual Chassis configuration

```
set virtual-chassis member 0 mastership-priority 255
set no-split-detection
```

3. Power on the second EX4500 switch. This switch has the backup role, and now both switches work as a single switch.

Important: If you do not edit the VC configuration file, a nonprovisioned configuration is generated by default. The mastership priority value for each member switch is 128. The master role is selected by default.

General configuration

Complete the following steps to set up the general configuration:

1. Assign IP addresses for the management Port and configure an IP gateway for our 172 management network, as shown in Example 5-20.

Example 5-20 Management IP address and IP gateway configuration on EX4500-VC

```
set interfaces vme unit 0 family inet address 172.25.101.241/16
set routing-options static route 0.0.0.0/0 next-hop 172.25.1.1
```

2. Create the DATA (4092) VLAN and create the Routed VLAN Interfaces (RVI) for our Data VLAN, as shown in Example 5-21.

Example 5-21 Data (4092) vlan and RVI configuration on EX4500-VC

```
set vlans DATA vlan-id 4092
set vlans DATA 13-interface vlan.4092
set interfaces vlan unit 4092 family inet address 10.1.4.241/24
```

3. Disable RSTP protocol, enable VSTP protocol (to be compatible with ToR and Flex switches PVRST protocol) and set the spanning-tree priority for the DATA (4092) VLAN to be 4 k so it becomes the root bridge for Layer 2 spanning-tree functionality, as shown in Example 5-22.

Example 5-22 VSTP and STP priority configuration on EX4500-VC

```
delete protocols rstp
set protocols vstp vlan DATA bridge-priority 4k
```

Configuring downstream G8264tor facing ports

Configure the downstream logical and physical interfaces to be all on the same aggregated interface by using LACP aggregation, as shown in Example 5-23.

Example 5-23 Downstream G8264tor facing interfaces on EX4500-VC, step [8]

```
set chassis aggregated-devices ethernet device-count 1
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic slow
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members DATA

set interfaces xe-0/0/0 description "Connection to 8264_1 Port 18"
set interfaces xe-0/0/0 ether-options 802.3ad ae0
set interfaces xe-0/0/1 description "Connection to 8264_1 Port 20"
set interfaces xe-0/0/1 ether-options 802.3ad ae0
set interfaces xe-0/0/2 description "Connection to 8264_2 Port 22"
set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/3 description "Connection to 8264_2 Port 24"
set interfaces xe-0/0/3 ether-options 802.3ad ae0
set interfaces xe-1/0/0 description "Connection to 8264_2 Port 18"
set interfaces xe-1/0/0 ether-options 802.3ad ae0
set interfaces xe-1/0/1 description "Connection to 8264_2 Port 20"
set interfaces xe-1/0/1 ether-options 802.3ad ae0
set interfaces xe-1/0/2 description "Connection to 8264_1 Port 22"
set interfaces xe-1/0/2 ether-options 802.3ad ae0
set interfaces xe-1/0/3 description "Connection to 8264_1 Port 24"
set interfaces xe-1/0/3 ether-options 802.3ad ae0
```

5.3.6 Verification and show command output

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

We begin by showing helpful commands from the EN4093 switches and work our way up the Network Topology diagram all the way to the Juniper EX4500-VC switch.

EN4093 output

Here we list output from the switch with host name EN4093f1ex_1. Similar or identical output exists for the switch with host name EN4093f1ex_2.

Show version

The command output in Example 5-24 shows information regarding the switch that we used and the associated code and firmware level at that time.

Example 5-24 EN4093 flex_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 5-25 shows output regarding VLAN assignment for all the various ports on the switch.

Example 5-25 EN4093 flex_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-INTA14 INTB1-INTB14 EXT7-EXT10 EXT15-EXT22
4094	ISL	ena	dis	EXT7-EXT10
4095	Mgmt VLAN	ena	ena	EXTM MGT1

Show interface status

Because we only have one compute node in our chassis (in slot 1), this explains why all the other internal ports are listed as “down” from a link perspective in the output that is shown in Example 5-26.

Example 5-26 EN4093 flex_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 5-27 shows our physical topology and verifies that cables are plugged into the ports we specified in our Network Topology diagram and the configuration that specified in Appendix A, “Fully Redundant Implementation with IBM vLAG and Juniper MC-LAG Technology” on page 293.

Example 5-27 EN4093 flex_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 5-28 shows the command output regarding the status of the ISL between the EN4093 switches and the ports that comprise the ISL.

Example 5-28 EN4093 flex_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 5-29 shows that the vLAG between the EN4093 switches and G8264 switches is up and operational as referenced by the LACP admin key of 2000. Our ISL between the EN4093 switches also is up.

Example 5-29 EN4093 flex_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
```

EN4093 flex_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, role election is performed. The switch with primary role controls the centralized operation. Role election is non-preemptive (if a primary 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 priority is the same, switch with smaller system MAC address is the vLAG primary switch. It is possible to configure vLAG priority to anything between 0 - 65535 and priority was left at the default value of 0 in all examples.

show vlag adminkey 2000

The output in Example 5-30 shows that the vLAG is formed and enabled by using LACP reference key 2000.

Example 5-30 EN4093 flex_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 5-31 shows which ports are participating in an LACP aggregation and which reference keys are used on those specific interfaces.

Example 5-31 EN4093 flex_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

Failover output that shows which ports are monitored and which ports are shutdown should an issue be encountered is shown in Example 5-32. In our case, our upstream to G8264 links are monitored with LACP reference key 2000. Our control ports are the downstream internal I/O module ports that are used by the Compute Nodes.

Example 5-32 EN4093 flex_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, we issued ping commands to devices in the lab infrastructure on VLAN 4092 (Data VLAN), as shown in Example 5-33. IP address 10.4.1.10 represents a Compute Node with an operating system installed, flex_node1 on the Network Topology diagram.

Example 5-33 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

Here we list output from the switch with host name G8264tor_1. Similar or identical output exists for the switch with host name G8264tor_2.

Show version

The output in Example 5-34 shows information regarding the switch that we used and the associated code and firmware level at that time.

Example 5-34 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 5-35 shows VLAN assignment for all the various ports on the switch.

Example 5-35 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

Example 5-36 shows the full interface table, listing port status, speed, description, and so on for the G8264tor_1 switch.

Example 5-36 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 EX4500
19	19	1G/10G	full	no	no	down	19
20	20	10000	full	no	no	up	VLAG to EX4500
21	21	1G/10G	full	no	no	down	21
22	22	10000	full	no	no	up	VLAG to EX4500
23	23	1G/10G	full	no	no	down	23
24	24	10000	full	no	no	up	VLAG to EX4500
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 5-37 shows our physical topology and verifies that cables are plugged into the ports we specified in our Network Topology diagram and the configuration specified in Appendix A, “Fully Redundant Implementation with IBM vLAG and Juniper MC-LAG Technology” on page 293.

Example 5-37 G8264tor_1 show lldp remote-device output

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
-----	-----	-----	-----	-----
11	1	08 17 f4 33 75 00	11	G8264TOR-2
12	2	08 17 f4 33 75 00	12	G8264TOR-2
MGT	3	fc cf 62 40 a6 00	20	BNT-AS-PM
13	4	08 17 f4 33 75 00	13	G8264TOR-2
14	5	08 17 f4 33 75 00	14	G8264TOR-2
15	6	08 17 f4 33 75 00	15	G8264TOR-2
16	7	08 17 f4 33 75 00	16	G8264TOR-2
1	8	08 17 f4 33 75 00	1	G8264TOR-2
2	9	08 17 f4 33 75 00	2	G8264TOR-2
3	10	08 17 f4 33 75 00	3	G8264TOR-2
4	11	08 17 f4 33 75 00	4	G8264TOR-2
5	12	08 17 f4 33 75 00	5	G8264TOR-2
6	13	08 17 f4 33 75 00	6	G8264TOR-2
7	14	08 17 f4 33 75 00	7	G8264TOR-2
8	15	08 17 f4 33 75 00	8	G8264TOR-2
9	16	08 17 f4 33 75 00	9	G8264TOR-2
18	17	28 c0 da 3c 64 40	505	EX4500-VC
10	18	08 17 f4 33 75 00	10	G8264TOR-2
20	19	28 c0 da 3c 64 40	507	EX4500-VC
22	20	28 c0 da 3c 64 40	618	EX4500-VC
24	21	28 c0 da 3c 64 40	619	EX4500-VC
26	22	6c ae 8b bf 6d 00	57	en4093flex_1
25	23	6c ae 8b bf 6d 00	58	en4093flex_1
28	24	6c ae 8b bf 6d 00	59	en4093flex_1
27	25	6c ae 8b bf 6d 00	60	en4093flex_1
37	26	6c ae 8b bf fe 00	57	en4093flex_2

39	27	6c ae 8b bf fe 00	58	en4093flex_2
38	28	6c ae 8b bf fe 00	59	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

Example 5-38 command output shows the status of the ISL between the G8264 switches and the ports that comprise the ISL.

Example 5-38 G8264tor_1 show vlag isl output

ISL_ID	ISL_Vlan	ISL_Trunk	ISL_Members	Link_State	Trunk_State
65	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 5-39 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 EX4500-VC switch, which is referenced by the LACP admin key of 2000. Our ISL between the G8264 switches also is up.

Example 5-39 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 65
ISL state Up
Startup Delay Interval: 120s (Finished)

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

vLAG 66: config with admin key 2002, associated trunk 67, state formed
```

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, role election is performed. The switch with primary role controls the centralized operation. Role election is non-preemptive (if a primary exists) and 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 priority is the same, switch with smaller system MAC address is the vLAG primary switch. It is possible to configure vLAG priority to anything between 0 - 65535, priority was left at the default value of 0 in all examples.

show vlag adminkey 2002

The output in Example 5-40 shows that the downstream vLAG towards the EN4093 switches is formed and enabled by using LACP reference key 2002.

Example 5-40 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 5-41 shows that the upstream vLAG towards the EC4500-VC switch is formed and enabled by using LACP reference key 2000.

Example 5-41 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 5-42 shows which ports are participating in an LACP aggregation and which reference keys are used on those specific interfaces.

Example 5-42 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, we issued several ping commands to devices in the lab infrastructure on VLAN 4092 (Data VLAN), as shown in Example 5-43. IP address 10.4.1.10 represents a Compute Node with an operating system installed, flex_node1 on the Network Topology diagram.

Example 5-43 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 10.1.4.238 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 1 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.

G8264TOR-1#ping 10.1.4.241 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 1 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 2 msec.
10.1.4.241: #5 ok, RTT 2 msec.
Ping finished.
```

EX4500-VC output

Here we list output from the switch with host name EX4500-VC.

show system software

Example 5-44 shows information regarding the switch that we used and the associated code and firmware level at that time.

Example 5-44 EX4500-VC show system software output

fpc0:

Information for fips-mode-powerpc:

Comment:

Junos OS FIPS mode utilities [12.1R3.5]

Information for jbase:

Comment:

Junos OS Base OS Software Suite [12.1R3.5]

Information for jcrypto-ex:

Comment:

Junos OS Crypto Software Suite [12.1R3.5]

Information for jdocs-ex:

Comment:

Junos OS Online Documentation [12.1R3.5]

Information for jkernel-ex:

Comment:
Junos OS Kernel Software Suite [12.1R3.5]

Information for jpfe-ex45x:
Comment:
Junos OS Packet Forwarding Engine Enterprise Software Release
Copyright (c) 1996-2012, Juniper Networks, Inc.
All rights reserved.

Junos OS Packet Forwarding Engine Enterprise Software Suite for EX45xx series
[12.1R3.5]

Information for jroute-ex:
Comment:
Junos OS Routing Software Suite [12.1R3.5]

Information for jswitch-ex:
Comment:
Junos OS Enterprise Software Suite [12.1R3.5]

Information for junos:
Comment:
Junos OS Base OS boot [12.1R3.5]

Information for jweb-ex:
Comment:
Junos OS Web Management [12.1R3.5]

fpc1:

Information for fips-mode-powerpc:
Comment:
Junos OS FIPS mode utilities [12.1R3.5]

Information for jbase:
Comment:
Junos OS Base OS Software Suite [12.1R3.5]

Information for jcrypto-ex:
Comment:
Junos OS Crypto Software Suite [12.1R3.5]

Information for jdocs-ex:
Comment:
Junos OS Online Documentation [12.1R3.5]

Information for jkernel-ex:
Comment:
Junos OS Kernel Software Suite [12.1R3.5]

Information for jpfe-ex45x:
Comment:
Junos OS Packet Forwarding Engine Enterprise Software Release
Copyright (c) 1996-2012, Juniper Networks, Inc.
All rights reserved.

Junos OS Packet Forwarding Engine Enterprise Software Suite for EX45xx series
[12.1R3.5]

Information for jroute-ex:
Comment:
Junos OS Routing Software Suite [12.1R3.5]

Information for jswitch-ex:
Comment:
Junos OS Enterprise Software Suite [12.1R3.5]

Information for junos:
Comment:
Junos OS Base OS boot [12.1R3.5]

Information for jweb-ex:
Comment:
Junos OS Web Management [12.1R3.5]

show virtual chassis status

The output shown in Example 5-45 show the status, members roles, members priorities, and so on of the EX4500 VC.

Example 5-45 EX4500-VC show VC status output

Virtual Chassis ID: cb81.1dce.51ee
Virtual Chassis Mode: Enabled

Member ID	Status	Serial No	Model	Mstr prio	Role	Mixed Neighbor List Mode ID	Interface
0 (FPC 0)	Prsnt	DE0210512973	ex4500-40f	255	Master*	N 1	vcp-1
						1	vcp-0
1 (FPC 1)	Prsnt	GG0210271860	ex4500-40f	128	Backup	N 0	vcp-1
						0	vcp-0

Member ID for next new member: 2 (FPC 2)

show virtual-chassis vc-port

Example 5-46 shows the status, speed, and physical connections of the dedicated VC ports.

Example 5-46 EX4500-VC show virtual-chassis vc-port output

fpc0:

Interface or PIC / Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0	Dedicated	2	Up	32000	1	vcp-1
vcp-1	Dedicated	2	Up	32000	1	vcp-0

fpc1:

Interface or PIC / Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0	Dedicated	2	Up	32000	0	vcp-1
vcp-1	Dedicated	2	Up	32000	0	vcp-0

show vlan

Example 5-47 shows the VLAN assignments for the various ports on the switch.

Example 5-47 EX4500-VC show vlan output

Name	Tag	Interfaces
DATA	4092	ae0.0*
default		None

show interface terse

Example 5-48 shows the full interface table (logical and physical), which lists administrative status, link status, IP address, and so on.

Example 5-48 EX4500-VC show interface terse output

Interface	Admin	Link	Proto	Local	Remote
xe-0/0/0	up	up			
xe-0/0/0.0	up	up	aenet	--> ae0.0	
xe-0/0/1	up	up			
xe-0/0/1.0	up	up	aenet	--> ae0.0	
xe-0/0/2	up	up			
xe-0/0/2.0	up	up	aenet	--> ae0.0	
xe-0/0/3	up	up			
xe-0/0/3.0	up	up	aenet	--> ae0.0	
xe-0/0/8	up	down			
xe-0/0/9	up	down			
xe-0/0/10	up	down			
xe-0/0/11	up	down			
xe-0/0/12	up	down			
xe-0/0/13	up	down			
xe-0/0/14	up	down			
xe-0/0/15	up	down			

xe-0/0/16	up	down		
xe-0/0/17	up	down		
xe-0/0/38	up	up		
xe-1/0/0	up	up		
xe-1/0/0.0	up	up	aenet	--> ae0.0
xe-1/0/1	up	up		
xe-1/0/1.0	up	up	aenet	--> ae0.0
xe-1/0/2	up	up		
xe-1/0/2.0	up	up	aenet	--> ae0.0
xe-1/0/3	up	up		
xe-1/0/3.0	up	up	aenet	--> ae0.0
xe-1/0/8	up	down		
xe-1/0/9	up	down		
xe-1/0/10	up	down		
xe-1/0/11	up	down		
xe-1/0/12	up	down		
xe-1/0/13	up	down		
xe-1/0/14	up	down		
xe-1/0/15	up	down		
xe-1/0/16	up	down		
xe-1/0/17	up	down		
xe-1/0/38	up	up		
vcp-0	up	up		
vcp-0.32768	up	up		
vcp-1	up	up		
vcp-1.32768	up	up		
ae0	up	up		
ae0.0	up	up	eth-switch	
ae1	up	down		
bme0	up	up		
bme0.32768	up	up	inet	128.0.0.1/2 128.0.0.16/2 128.0.0.32/2
			tnp	0x10
bme0.32770	down	up	eth-switch	
dsc	up	up		
gre	up	up		
ipip	up	up		
lo0	up	up		
lsi	up	up		
me0	up	up		
me0.0	up	up	eth-switch	
mtun	up	up		
pimd	up	up		
pime	up	up		
tap	up	up		
vlan	up	up		
vlan.4092	up	up	inet	10.1.4.12/24
vme	up	up		
vme.0	up	up	inet	172.25.101.241/16

show lldp neighbors

Example 5-49 lists the LLDP information and serves as a means to verify our physical connectivity.

Example 5-49 EX4500-VC show lldp neighbors output

Local Interface	Parent Interface	Chassis Id	Port info	SystemName
xe-1/0/0.0	ae0.0	08:17:f4:33:75:00	Ethernet18	G8264TOR-2
xe-1/0/1.0	ae0.0	08:17:f4:33:75:00	Ethernet20	G8264TOR-2
xe-0/0/2.0	ae0.0	08:17:f4:33:75:00	Ethernet22	G8264TOR-2
xe-0/0/3.0	ae0.0	08:17:f4:33:75:00	Ethernet24	G8264TOR-2
xe-0/0/0.0	ae0.0	08:17:f4:33:9d:00	Ethernet18	G8264TOR-1
xe-0/0/1.0	ae0.0	08:17:f4:33:9d:00	Ethernet20	G8264TOR-1
xe-1/0/2.0	ae0.0	08:17:f4:33:9d:00	Ethernet22	G8264TOR-1
xe-1/0/3.0	ae0.0	08:17:f4:33:9d:00	Ethernet24	G8264TOR-1
vme.0	-	fc:cf:62:40:a6:00	24	BNT-AS-PM

Ping output for equipment on VLAN 4092

To verify connectivity, we issued several ping commands to devices in the lab infrastructure on VLAN 4092 (Data VLAN), as shown in Example 5-50.

Example 5-50 Ping verification for equipment on VLAN 4092

```
admin@EX4500-VC> 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=64 time=1.277 ms
64 bytes from 10.1.4.10: icmp_seq=1 ttl=64 time=2.320 ms
64 bytes from 10.1.4.10: icmp_seq=2 ttl=64 time=0.782 ms
64 bytes from 10.1.4.10: icmp_seq=3 ttl=64 time=0.786 ms
64 bytes from 10.1.4.10: icmp_seq=4 ttl=64 time=0.788 ms
64 bytes from 10.1.4.10: icmp_seq=5 ttl=64 time=0.756 ms
64 bytes from 10.1.4.10: icmp_seq=6 ttl=64 time=0.775 ms
64 bytes from 10.1.4.10: icmp_seq=7 ttl=64 time=0.751 ms
64 bytes from 10.1.4.10: icmp_seq=8 ttl=64 time=0.799 ms
64 bytes from 10.1.4.10: icmp_seq=9 ttl=64 time=0.754 ms
^C
--- 10.1.4.10 ping statistics ---
10 packets transmitted, 10 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.751/0.979/2.320/0.472 ms

admin@EX4500-VC> 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=255 time=2.385 ms
64 bytes from 10.1.4.243: icmp_seq=1 ttl=255 time=1.213 ms
64 bytes from 10.1.4.243: icmp_seq=2 ttl=255 time=1.416 ms
64 bytes from 10.1.4.243: icmp_seq=3 ttl=255 time=1.377 ms
64 bytes from 10.1.4.243: icmp_seq=4 ttl=255 time=1.139 ms
64 bytes from 10.1.4.243: icmp_seq=5 ttl=255 time=1.263 ms
64 bytes from 10.1.4.243: icmp_seq=6 ttl=255 time=2.459 ms
^C
--- 10.1.4.243 ping statistics ---
7 packets transmitted, 7 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.139/1.607/2.459/0.523 ms

admin@EX4500-VC> 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=1.401 ms
64 bytes from 10.1.4.238: icmp_seq=1 ttl=254 time=2.832 ms
64 bytes from 10.1.4.238: icmp_seq=2 ttl=254 time=1.839 ms
64 bytes from 10.1.4.238: icmp_seq=3 ttl=254 time=1.840 ms
64 bytes from 10.1.4.238: icmp_seq=4 ttl=254 time=1.959 ms
64 bytes from 10.1.4.238: icmp_seq=5 ttl=254 time=2.149 ms
^C
--- 10.1.4.238 ping statistics ---
6 packets transmitted, 6 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.401/2.003/2.832/0.433 ms

```

5.3.7 Full configuration files

In this section, we describe the configuration that was used on all of the devices in the Network Topology diagram.

EN4093 flex-1

Example 5-51 lists the configuration for the EN4093 flex-1 switch.

Example 5-51 EN4093 flex-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 h1thchk"
    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
!
ip gateway 1 address 10.1.4.241
ip gateway 1 enable
!
!
!
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
4002479906b4f177661b42d268b55d2cfea55ca43558622
!
ntp trusted-key 8811
!
end

```

EN4093 flex_2

Example 5-52 lists the configuration for the EN4093flex_2 switch.

Example 5-52 EN4093 flex_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 h1thchk"
    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 h1thchk"
    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
!
ip gateway 1 address 10.1.4.241
ip gateway 1 enable
!
!
!
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 dd204ddaa96d7fdd5a63571d5ce8f3ac7dfd1c4
!
ntp trusted-key 8811
!
end
```

G8264tor_1

Example 5-53 lists the configuration for the G8264tor_1 switch.

Example 5-53 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 EX4500"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 20
    name "VLAG to EX4500"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 22
    name "VLAG to EX4500"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 24
    name "VLAG to EX4500"
    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 1
!    addr <default>
!    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 1 address 10.1.4.241
ip gateway 1 enable
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
!
!
end

```

G8264tor_2

Example 5-54 lists the configuration for the G8264tor_2 switch.

Example 5-54 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 EX4500"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 20
    name "VLAG to EX4500"
    tagging
    tag-pvid
    pvid 4092
    exit
!

```

```

interface port 22
    name "VLAG to EX4500"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 24
    name "VLAG to EX4500"
    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 1

```

```

!      addr <default>
!      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 1 address 10.1.4.241
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
!
!
end

```

EX4500-VC switch

Example 5-55 lists the configuration of the Juniper EX4500-VC switch.

Example 5-55 EX4500-VC switch configuration

```

set version 12.1R3.5
set system host-name EX4500-VC
set system root-authentication encrypted-password
"$1$16b7xXmg$rEkslG/xYqFKDv4y2rHCt/"
set system login user admin uid 2000
set system login user admin class super-user
set system login user admin authentication encrypted-password
"$1$440v.vyA$Lzz2LAAf3QyMlqBuxUtXk0"
set system services ftp
set system services ssh
set system services netconf ssh
set system syslog user * any emergency
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/0 description "Connection to 8264_1 Port 18"
set interfaces xe-0/0/0 ether-options 802.3ad ae0
set interfaces xe-0/0/1 description "Connection to 8264_1 Port 20"
set interfaces xe-0/0/1 ether-options 802.3ad ae0
set interfaces xe-0/0/2 description "Connection to 8264_2 Port 22"
set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/3 description "Connection to 8264_2 Port 24"

```

```

set interfaces xe-0/0/3 ether-options 802.3ad ae0
set interfaces xe-1/0/0 description "Connection to 8264_2 Port 18"
set interfaces xe-1/0/0 ether-options 802.3ad ae0
set interfaces xe-1/0/1 description "Connection to 8264_2 Port 20"
set interfaces xe-1/0/1 ether-options 802.3ad ae0
set interfaces xe-1/0/2 description "Connection to 8264_1 Port 22"
set interfaces xe-1/0/2 ether-options 802.3ad ae0
set interfaces xe-1/0/3 description "Connection to 8264_1 Port 24"
set interfaces xe-1/0/3 ether-options 802.3ad ae0
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic slow
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members DATA
set interfaces ae0 unit 0 family ethernet-switching native-vlan-id 1
set interfaces vlan unit 4092 family inet address 10.1.4.241/24
set interfaces vme unit 0 family inet address 172.25.101.241/16
set routing-options static route 0.0.0.0/0 next-hop 172.25.1.1
set protocols igmp-snooping vlan all
set protocols vstp vlan DATA bridge-priority 4k
set protocols vstp vlan native-vlan bridge-priority 4k
set protocols lldp interface all
set protocols lldp-med interface all
set ethernet-switching-options storm-control interface all
set vlans DATA vlan-id 4092
set vlans DATA l3-interface vlan.4092
set virtual-chassis member 0 mastership-priority 255

```

5.4 Fully redundant with traditional spanning tree protocol

This section describes the various aspects of the implementation scenario.

5.4.1 Topology and requirements

This implementation scenario uses a more traditional, classic network design with the spanning tree protocol serving as a protection against L2 loops. Customers with upstream Juniper equipment that might not have the ability to implement a VC (for example, EX4500 switches without a VC module), or if they are more comfortable with STP can choose this implementation scenario.

This approach includes the following advantages:

- ▶ Almost plug and play if VLAN Spanning-Tree Protocol (VSTP) is used on Juniper equipment and Per-Vlan-Rapid Spanning-Tree protocol (PVRST+) on IBM equipment (default selection as of recent software versions of N/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 Juniper switch.

This approach includes the following disadvantages:

- ▶ Links are blocked by spanning-tree to prevent L2 loops, which wastes valuable bandwidth.
- ▶ Can take slightly longer convergence times in the event of a link failure.
- ▶ Troubleshooting problems with spanning-tree can be more difficult for less experienced network architects.

5.4.2 Components used

Two of each of the following components were used:

- ▶ Juniper EX4500-40F
- ▶ IBM G8264 RackSwitch
- ▶ IBM Flex System Fabric EN4093 10Gb Scalable Switch

5.4.3 Network diagram and physical setup

Figure 5-5 on page 131 shows the Network topology for the fully redundant scenario with spanning tree.

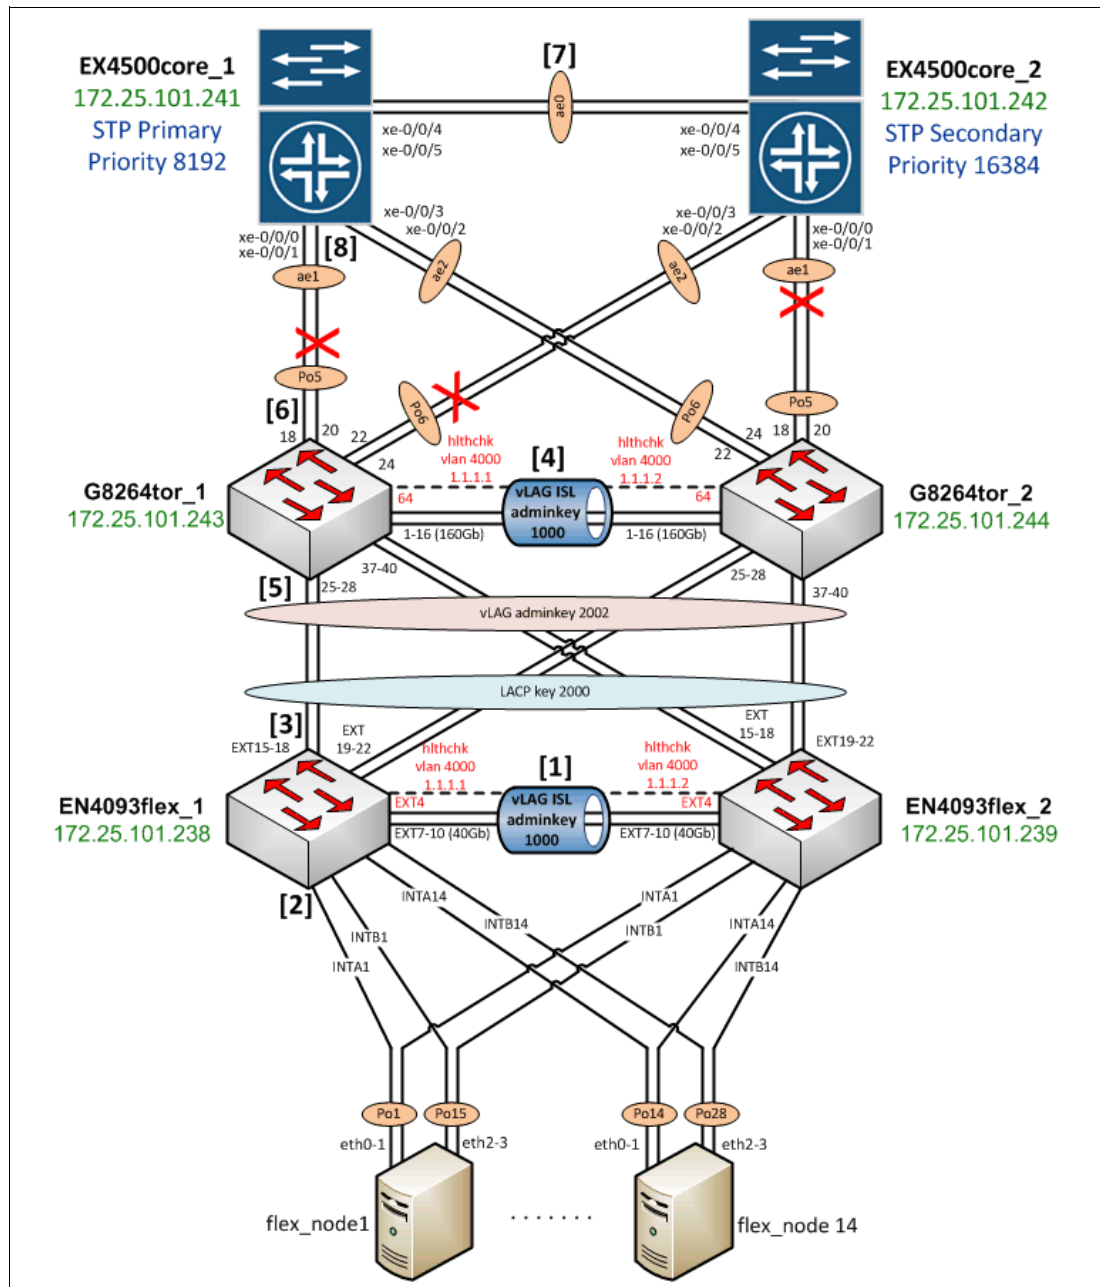


Figure 5-5 Network Topology diagram for fully redundant scenario by using spanning-tree

Start by verifying the physical cabling between the EN4093 flex switches and G8264tor switches. In our lab environment, we used four IBM QSFP+ DAC Break Out Cables from the EN4093 switches to the upstream G8264s. This requires that the EN4093 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 160Gb ISL between the G8264 switches and 10Gb SFP+ DAC cables were used for all other connections in Figure 5-5.

5.4.4 EN4093 flex_1 configuration

We begin the implementation of this scenario on the IBM Flex System Fabric EN4093 switches, working our way northward on the diagram in Figure 5-5 on page 131. Each step provides the commands necessary and are reflective of the numbering schema in the diagram to aid the user in what is configured.

General configuration

Complete the following steps to set up the general configuration:

1. Create the ISL Healthcheck, ISL data, and Data VLANs as shown in Example 5-56, giving them descriptive names, assigning them to spanning-tree groups, and enabling them. You can elect to have the switch create STP instances for you; we chose to manually create them instead.

Example 5-56 Create ISL hlthchk, Data and ISL VLANs on EN4093 flex_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 and Data VLANs (as shown in Example 5-57) so that we can verify connectivity between the various pieces of equipment when the configuration is verified. 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, with the last octet borrowed from the network diagram's Management address to quickly aid in the identification of which piece of equipment we are verifying connectivity to.

Example 5-57 Create IP interfaces and assign vlans and IP addresses

```
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 the ISL between EN4093 flex switches

Complete the following steps to configure the ISL between EN4093 flex switches:

1. Configure the eventual ISL that is shown in Example 5-58 between the EN4093 switches by configuring them to have a default (untagged) VLAN of 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 5-58 Initial ISL configuration on EN4093 flex_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
```

2. Configure the dedicated health check physical interface shown in Example 5-59 to be used for heartbeats between the EN4093 switches. We chose EXT4 as a dedicated interface and VLAN 4000 to serve as the health check for the ISL.

Example 5-59 Create vLAG health check on EN4093 flex_1

```
configure terminal
interface port ext4
    pvid 4000
    exit
```

3. Disable STP between the EN4093 switches and activate a vLAG between them so that they appear as a single entity to the upstream and downstream infrastructure, as shown in Example 5-60, and reference the LACP key that is configured in the previous step.

Example 5-60 Disable STP and activate ISL vLAG on EN4093 flex_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

Complete the following steps to configure downstream internal node ports:

1. Configure downstream node interfaces in Example 5-61 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 5-61 Downstream Internal node port configuration, on EN4093 flex_1, step [2]

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

2. For redundancy, we created two port-channels on each of the 14 nodes. Each port channel aggregates two ports, one from each EN4093 flex switch. Port channels 1 - 14 match the A internally labeled ports and port channel 15 - 28 match the B ports, as shown in Example 5-62.

Example 5-62 Node-facing port channel creation and vLAG activation on EN4093 flex_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 G8264 facing ports and Layer 2 failover

Complete the following steps to configure upstream G8264 facing ports and Layer 2 failover:

1. Setup the upstream ports to the G8264s that are shown in Example 5-63 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 5-63 Upstream G8264tor facing port configuration on EN4093 flex_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
```

2. Activate the vLAG feature for the upstream EN4093 ports so that the G8264s see the EN4093s as a single, virtualized entity, as shown in Example 5-64. Use adminkey 2000, which represents the LACP key bundling ports EXT15-22 together as one.

Example 5-64 Activating the upstream vLAG to the G8264 switches on EN4093 flex_1

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

3. Enable Layer 2 failover that is shown in Example 5-65, which effectively shuts down the links to the compute nodes should the uplinks for the EN4093 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, which in our case is connected to the other EN4093 switch in the Enterprise Chassis, which ensures that redundancy is maintained.

Example 5-65 Enabling L2 failover for the compute nodes on EN4093 flex_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 flex_2 on the other I/O module. The only difference between the EN4093 flex_1 switch and the EN4093 flex_2 switch is the vLAG health check peer address and the Data, and ISL hlthchk vlan ip addresses. To verify the EN4093 flex switch configuration, run the show commands that are described in 5.4.8, “Verification and show command output” on page 143.

5.4.5 G8264tor_1 configuration

Next is the configuration of the RackSwitch G8264.

General configuration

Complete the following steps to set up the general configuration:

1. Begin by creating the ISL Healthcheck, ISL data, and Data VLANs, as shown in Example 5-66. Give them descriptive names, assign them to spanning-tree groups, and enable them.

Example 5-66 Create 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 that is shown in Example 5-67. Interface ip 128 represents the management IP address that is referenced in the Network Topology diagram in Figure 5-4 on page 75. IP gateway 4 is the upstream router interface for our 172 management network.

Example 5-67 Create IP interfaces and assign 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 the ISL between G8264tor switches

Complete the following steps to configure the ISL between G8264t or switches:

1. Configure the ISL between the G8264tor switches, as shown in Example 5-68. 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 5-68 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
```

2. Disable STP between the G8264 switches and activate a vLAG between them so that they appear as a single entity to the upstream and downstream infrastructure, as shown in Example 5-69, which references the LACP key that was configured in the previous step.

Example 5-69 Disable STP and activate 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 EN4093 flex-facing ports

Complete the following steps to configure downstream EN4093 flex-facing ports:

1. Configure the downstream ports towards the EN4093 switches as shown in Example 5-70 to have a default (untagged) VLAN of 4092 (data vlan), with 802.1q tagging enabled and tag the PVID. Add the ability for all member ports to be on VLAN 4092 and use an LACP key of 2002 to bundle the ports together.

Example 5-70 Downstream EN4093 flex-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
```

2. Activate the vLAG for the downstream EN4093 flex-facing ports so that the 4093s see the G8264s as a single, virtualized entity, as shown in Example 5-71. Use adminkey 2002, which represents the LACP key bundling ports 25 - 28, and 37 - 40 together as one.

Example 5-71 Activate downstream EN4093 facing vLAG on G8264tor_1

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

Configuring upstream EX4500 facing ports

Complete the following steps to configure upstream EX4500 facing ports:

1. Configure the upstream ports to the EX4500 switches as shown in Example 5-72 with a default (untagged) VLAN of 4092 (data vlan), enable 802.1q tagging, tag the PVID and give a useful description on the interfaces.

Example 5-72 Upstream EX4500 facing port configuration on G8264tor_1

```
configure terminal
interface port 18,20
    name "Po to EX4500-1"
    pvid 4092
    tagging
    tag-pvid
    exit
interface port 22,24
    name "Po to EX4500-2"
    pvid 4092
    tagging
    tag-pvid
    exit
vlan 4092
    member 18,20,22,24
```

exit

2. Activate link aggregation groups by using static port-channeling, as shown in Example 5-73. We chose to use static port-channeling here to show that IBM System Networking equipment interoperates with an upstream Juniper infrastructure with LACP or static (no negotiation protocol) port-channeling.

Example 5-73 Create 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 hlthchk vlan ip addresses. To verify the G8264tor switch configuration, run the show commands that are described in 5.4.8, “Verification and show command output” on page 143.

5.4.6 EX4500_1 STP primary switch configuration

Here, we configure the Juniper EX4500_1 STP primary switch.

General configuration

Complete the following steps to set up the general configuration:

1. Assign an IP address for the management Port and configure an IP gateway for our 172 management network, as shown in Example 5-74.

Example 5-74 Management IP address and IP gateway configuration on EX4500_1

```
set interfaces me0 unit 0 family inet address 172.25.101.241/16
set routing-options static route 0.0.0.0/0 next-hop 172.25.1.1
```

2. Create vlan 4092 (data vlan). Enable VLAN Spanning Tree Protocol (VSTP), disable RSTP, and set the spanning tree priority for all VLAN to be half that of the EX4500_2 switch (8k). Because EX4500_1 switch has a lower spanning tree priority, it becomes the root bridge for L2 functionality, as shown in Example 5-75.

We use VSTP because of its compatibility with the Per-VLAN Rapid Spanning Tree Protocol (PVRST) that is running on the IBM NS switches. PVRST is required to configure vLAG in the IBM NS switches.

Example 5-75 Data VLAN configuration and spanning-tree configuration on EX4500_1

```
set vlans DATA vlan-id 4092
delete protocols rstp
set protocols vstp vlan all bridge-priority 8k
```

3. Create the Routed VLAN Interfaces (RVI) for our Data VLAN, as shown in Example 5-76, which is useful during verification of this scenario's implementation.

Example 5-76 Create IP address on vlan 4092 (Data vlan) on EX4500_1

```
set interfaces vlan unit 4092 family inet address 10.1.4.241/24
set vlans DATA 13-interface vlan.4092
```

Configuring switch-to-switch link between the EX4500 switches

Configure the logical and physical interfaces that comprise the switch-to-switch link between the EX4500_1 and EX4500_2 switches, as shown in Example 5-77.

Example 5-77 Switch-to-switch link logical and physical interface configuration, on EX4500_1

```
set chassis aggregated-devices ethernet device-count 1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members DATA
set interfaces xe-0/0/4 ether-options 802.3ad ae0
set interfaces xe-0/0/5 ether-options 802.3ad ae0
```

Configuring downstream G8264tor facing ports

For the EX4500 primary switch, configure the downstream logical and physical interfaces as shown in Example 5-78, with interfaces xe-0/0/0 and xe-0/0/1 to be bundled in static aggregation ae1, and interfaces xe-0/0/2 and xe-0/0/3 to be bundled in static aggregation ae2.

Example 5-78 Downstream G8264tor facing port configuration on EX4500_1

```
set chassis aggregated-devices ethernet device-count 3
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members DATA
set interfaces ae2 unit 0 family ethernet-switching port-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members DATA
set interfaces xe-0/0/0 ether-options 802.3ad ae1
set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/2 ether-options 802.3ad ae2
set interfaces xe-0/0/3 ether-options 802.3ad ae2
```

5.4.7 EX4500_2 STP secondary switch configuration

Next, we configure the Juniper EX4500_2 secondary switch.

General configuration

Complete the following steps to set up the general configuration:

1. Assign an IP address for the management Port and configure an IP gateway for our 172 management network, as shown in Example 5-79.

Example 5-79 Management IP address and IP gateway configuration on EX4500_2

```
set interfaces me0 unit 0 family inet address 172.25.101.242/16
set routing-options static route 0.0.0.0/0 next-hop 172.25.1.1
```

2. Create vlan 4092 (data vlan). Enable VLAN Spanning Tree Protocol (VSTP), disable RSTP and set the spanning-tree priority for all vlan to be twice that of EX4500_1 switch (16k). Because EX4500_2 switch has a higher spanning-tree priority than EX4500_1, it becomes the backup for the spanning tree protocol Layer 2 functionality, as shown in Example 5-80.

Example 5-80 Data VLAN configuration and spanning-tree configuration on EX4500_2

```
set vlans DATA vlan-id 4092
delete protocols rstp
set protocols vstp vlan all bridge-priority 16k
```

We use VSTP because of its compatibility with the Per-VLAN Rapid Spanning Tree Protocol (PVRST) that is running on the IBM NS switches. PVRST is required to configure vLAG in the IBM NS switches.

3. Create the Routed VLAN Interfaces (RVI) for our Data VLAN as shown in Example 5-81, which is useful during verification of this scenario's implementation.

Example 5-81 Create IP address on vlan 4092 (Data vlan) on EX4500_2

```
set interfaces vlan unit 4092 family inet address 10.1.4.242/24
set vlans DATA l3-interface vlan.4092
```

Configuring switch-to-switch link between EX4500 switches

Configure the logical and physical interfaces comprising the switch-to-switch link between the EX4500_1 and EX4500_2 switches as shown in Example 5-82.

Example 5-82 Switch-to-switch link logical and physical interface configuration on EX4500_2

```
set chassis aggregated-devices ethernet device-count 1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members DATA
set interfaces xe-0/0/4 ether-options 802.3ad ae0
set interfaces xe-0/0/5 ether-options 802.3ad ae0
```

Configuring downstream G8264tor facing ports

For the EX4500_2 switch, configure the downstream physical and logical interfaces in Example 5-83, with interfaces xe-0/0/0 and xe-0/0/1 to be bundled in static aggregation ae1, and interfaces xe-0/0/2 and xe-0/0/3 to be bundled in static aggregation ae2.

Example 5-83 Downstream G8264tor facing port configuration on EX4500_2, step [8]

```
set chassis aggregated-devices ethernet device-count 1
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members DATA
set interfaces ae2 unit 0 family ethernet-switching port-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members DATA
set interfaces xe-0/0/0 ether-options 802.3ad ae1
set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/2 ether-options 802.3ad ae2
set interfaces xe-0/0/3 ether-options 802.3ad ae2
```

5.4.8 Verification and show command output

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

We begin by showing helpful commands from the EN4093 switches, working our way up the Network Topology diagram all the way to the Juniper EX4500 pair.

EN4093 output

Here we list the output from the switch with hostname EN4093 flex_1. Similar or identical output exists for the switch with hostname EN4093 flex_2.

Show version

The command output in Example 5-84 shows information regarding the switch that we used and the associated code and firmware level at that time.

Example 5-84 EN4093 flex_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 5-85 shows the output regarding VLAN assignment for all the various ports on the switch.

Example 5-85 EN4093 flex_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-INTA14 INTB1-INTB14 EXT7-EXT10 EXT15-EXT22
4094	ISL	ena	dis	EXT7-EXT10
4095	Mgmt VLAN	ena	ena	EXTM MGT1

Show interface status

Because we only have one compute node in our chassis (in slot 1), this explains why all the other internal ports are listed as “down” from a link perspective in the output that is shown in Example 5-86.

Example 5-86 EN4093 flex_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 5-87 shows our physical topology and verifies that cables are plugged into the ports we specified in our Network Topology diagram and the configuration that is specified in Appendix A, “Fully Redundant Implementation with IBM vLAG and Juniper MC-LAG Technology” on page 293.

Example 5-87 EN4093 flex_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 5-88 shows the command output regarding the status of the ISL between the EN4093 switches and the ports that comprise the ISL.

Example 5-88 EN4093 flex_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 5-89 shows that the vLAG between the EN4093 switches and G8264 switches is up and operational as referenced by the LACP admin key of 2000. Our ISL between the EN4093 switches also is up.

Example 5-89 EN4093 flex_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
```

EN4093 flex_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, role election is performed. The switch with primary role controls the centralized operation. Role election is non-preemptive (that is, if a primary 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 priority is the same, the switch with smaller system MAC address is the vLAG primary switch. It is possible to configure vLAG priority to anything between 0 - 65535; priority was left at the default value of 0 in all examples.

show vlag adminkey 2000

The output in Example 5-90 shows that the vLAG is formed and enabled by using LACP reference key 2000.

Example 5-90 EN4093 flex_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 5-91 shows which ports are participating in an LACP aggregation and which reference keys are used on those specific interfaces.

Example 5-91 EN4093 flex_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

Failover output that shows which ports are monitored and which ports are shutdown should an issue be encountered are shown in Example 5-92. In our case, our upstream to G8264 links are monitored with LACP reference key 2000. Our control ports are the downstream internal I/O module ports that are used by the Compute Nodes.

Example 5-92 EN4093 flex_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, we issued ping commands to devices in the lab infrastructure on VLAN 4092 (Data VLAN) as shown in Example 5-93. IP address 10.4.1.10 represents a Compute Node with an operating system installed (flex_node1 on the Network Topology diagram in Appendix A, “Fully Redundant Implementation with IBM vLAG and Juniper MC-LAG Technology” on page 293).

Example 5-93 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 5 msec.
10.1.4.243: #5 ok, RTT 2 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 10.1.4.241 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

Here we list 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 5-94 shows information regarding the switch that we used and the associated code and firmware level at that time.

Example 5-94 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 5-95 shows the VLAN assignment for all of the various ports on the switch.

Example 5-95 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

Example 5-96 shows the full interface table, which lists port status, speed, description, and so on for the G8264tor_1 switch.

Example 5-96 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	Po to EX4500-1
19	19	1G/10G	full	no	no	down	19
20	20	10000	full	no	no	up	Po to EX4500-1
21	21	1G/10G	full	no	no	down	21
22	22	10000	full	no	no	up	Po to EX4500-2
23	23	1G/10G	full	no	no	down	23
24	24	10000	full	no	no	up	Po to EX4500-2
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 5-97 shows our physical topology and verifies that cables are plugged into the ports we specified in our Network Topology in diagram Figure 5-4 on page 75, and the configuration specified in Appendix A, “Fully Redundant Implementation with IBM vLAG and Juniper MC-LAG Technology” on page 293.

Example 5-97 G8264tor_1 show lldp remote-device output

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
11	1	08 17 f4 33 75 00	11	G8264TOR-2
12	2	08 17 f4 33 75 00	12	G8264TOR-2
20	3	28 c0 da 3c 64 40	507	EX4500-1
13	4	08 17 f4 33 75 00	13	G8264TOR-2
14	5	08 17 f4 33 75 00	14	G8264TOR-2
15	6	08 17 f4 33 75 00	15	G8264TOR-2
16	7	08 17 f4 33 75 00	16	G8264TOR-2
1	8	08 17 f4 33 75 00	1	G8264TOR-2
2	9	08 17 f4 33 75 00	2	G8264TOR-2
3	10	08 17 f4 33 75 00	3	G8264TOR-2
4	11	08 17 f4 33 75 00	4	G8264TOR-2
5	12	08 17 f4 33 75 00	5	G8264TOR-2
6	13	08 17 f4 33 75 00	6	G8264TOR-2
7	14	08 17 f4 33 75 00	7	G8264TOR-2
8	15	08 17 f4 33 75 00	8	G8264TOR-2
9	16	08 17 f4 33 75 00	9	G8264TOR-2
18	17	28 c0 da 3c 64 40	505	EX4500-1
10	18	08 17 f4 33 75 00	10	G8264TOR-2
24	19	80 71 1f d6 ad 40	512	EX4500-2
22	20	80 71 1f d6 ad 40	510	EX4500-2
MGT	21	fc cf 62 40 a6 00	22	BNT-AS-PM
26	22	6c ae 8b bf 6d 00	57	en4093flex_1
25	23	6c ae 8b bf 6d 00	58	en4093flex_1
28	24	6c ae 8b bf 6d 00	59	en4093flex_1
27	25	6c ae 8b bf 6d 00	60	en4093flex_1
37	26	6c ae 8b bf fe 00	57	en4093flex_2
39	27	6c ae 8b bf fe 00	58	en4093flex_2
38	28	6c ae 8b bf fe 00	59	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

Example 5-98 command output shows the status of the ISL between the G8264 switches and the ports that comprise the ISL.

Example 5-98 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 5-99 on page 155 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. Our ISL between the G8264 switches also is up.

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, role election is performed. The switch with primary role controls the centralized operation. Role election is non-preemptive (that is, if a primary 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 priority is the same, switch with smaller system MAC address is the vLAG primary switch. It is possible to configure vLAG priority to anything between 0 - 65535, priority was left at the default value of 0 in all examples.

Example 5-99 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 5-100 shows that the downstream vLAG towards the EN4093 switches is formed and enabled by using LACP reference key 2002.

Example 5-100 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 5-42 shows which ports are participating in an LACP aggregation and which reference keys are used on those specific interfaces.

Example 5-101 G8264tor_1 show lacp information state up output

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 5-102 lists output from the show spanning tree command on G8264tor_1. The blocked links are reflected in the Network Topology diagram in Figure 5-4 on page 75 for VLAN 4092 .

Example 5-102 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      18
```

	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	(pc67)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c1	P2P
26	(pc67)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c1	P2P
27	(pc67)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c1	P2P
28	(pc67)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c1	P2P
37	(pc67)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c1	P2P
38	(pc67)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c1	P2P
39	(pc67)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c1	P2P
40	(pc67)	128	200!+	FWD	DESG	8001-08:17:f4:33:9d:00	84c1	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         4

```

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 28:c0:da:3c:64:41    1105     16   2    20    15

```

```

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
           32894     2     20    15    300        54

```

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
1	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
2	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
3	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
4	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
5	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
6	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
7	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
8	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
9	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
10	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
11	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
12	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
13	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
14	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
15	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
16	(pc65)	128	115!+ FWD	ROOT	807e-08:17:f4:33:75:00	8440	P2P
18	(pc5)	128	990!+ DISC	ALTN	2ffc-28:c0:da:3c:64:41	8002	P2P
20	(pc5)	128	990!+ DISC	ALTN	2ffc-28:c0:da:3c:64:41	8002	P2P
22	(pc6)	128	990!+ DISC	ALTN	4ffc-80:71:1f:d6:ad:41	8003	P2P
24	(pc6)	128	990!+ DISC	ALTN	4ffc-80:71:1f:d6:ad:41	8003	P2P
25	(pc67)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c1	P2P
26	(pc67)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c1	P2P
27	(pc67)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c1	P2P
28	(pc67)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c1	P2P
37	(pc67)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c1	P2P
38	(pc67)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c1	P2P
39	(pc67)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c1	P2P
40	(pc67)	128	200!+ FWD	DESG	807e-08:17:f4:33:9d:00	84c1	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.

Show spanning tree on G8264tor_2

Example 5-103 lists output from the **show spanning tree** command on G8264tor_2. The blocked links are reflected in the Network Topology diagram in Figure 5-4 on page 75 for VLAN 4092.

Example 5-103 Output of show spanning tree command

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			52
	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	(pc67)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c1	P2P
26	(pc67)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c1	P2P
27	(pc67)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c1	P2P
28	(pc67)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c1	P2P
37	(pc67)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c1	P2P
38	(pc67)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c1	P2P
39	(pc67)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c1	P2P
40	(pc67)	128	200!+ FWD	DESG	8001-08:17:f4:33:75:00	84c1	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	3

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 28:c0:da:3c:64:41	990	22	2	20	15

Parameters:	Priority	Hello	MaxAge	FwdDel	Aging	Topology Change Counts
	32894	2	20	15	300	58

	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!+	FWD	DESG	807e-08:17:f4:33:75:00	8404	P2P
20	(pc5)	128	990!+	FWD	DESG	807e-08:17:f4:33:75:00	8404	P2P
22	(pc6)	128	990!+	FWD	ROOT	2ffc-28:c0:da:3c:64:41	8003	P2P
24	(pc6)	128	990!+	FWD	ROOT	2ffc-28:c0:da:3c:64:41	8003	P2P
25	(pc67)	128	200!+	FWD	DESG	807e-08:17:f4:33:75:00	84c1	P2P
26	(pc67)	128	200!+	FWD	DESG	807e-08:17:f4:33:75:00	84c1	P2P
27	(pc67)	128	200!+	FWD	DESG	807e-08:17:f4:33:75:00	84c1	P2P
28	(pc67)	128	200!+	FWD	DESG	807e-08:17:f4:33:75:00	84c1	P2P
37	(pc67)	128	200!+	FWD	DESG	807e-08:17:f4:33:75:00	84c1	P2P
38	(pc67)	128	200!+	FWD	DESG	807e-08:17:f4:33:75:00	84c1	P2P
39	(pc67)	128	200!+	FWD	DESG	807e-08:17:f4:33:75:00	84c1	P2P
40	(pc67)	128	200!+	FWD	DESG	807e-08:17:f4:33:75:00	84c1	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, we issued several ping commands to devices in the lab infrastructure on VLAN 4092 (Data VLAN) as shown in Example 5-104. IP address 10.4.1.10 represents a Compute Node with an operating system installed, flex_node1 on the Network Topology diagram.

Example 5-104 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 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.
```

```
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 10.1.4.241 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 0 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 2 msec.
10.1.4.241: #5 ok, RTT 1 msec.
Ping finished.
```

EX4500_1 output

Here we list output from the switch with hostname EX4500_1. Similar or identical output exists for the switch with hostname EX4500_2, unless otherwise noted.

show system software

Example 5-105 shows information regarding the switch that we used and the associated code and firmware level at that time.

Example 5-105 EX4500_1 show system software output

fpc0:

Information for fips-mode-powerpc:

Comment:

Junos OS FIPS mode utilities [12.1R3.5]

Information for jbase:

Comment:

Junos OS Base OS Software Suite [12.1R3.5]

Information for jcrypto-ex:

Comment:

Junos OS Crypto Software Suite [12.1R3.5]

Information for jdocs-ex:

Comment:

Junos OS Online Documentation [12.1R3.5]

Information for jkernel-ex:

Comment:

Junos OS Kernel Software Suite [12.1R3.5]

Information for jpfe-ex45x:

Comment:

Junos OS Packet Forwarding Engine Enterprise Software Release

Copyright (c) 1996-2012, Juniper Networks, Inc.

All rights reserved.

Junos OS Packet Forwarding Engine Enterprise Software Suite for EX45xx series
[12.1R3.5]

Information for jroute-ex:

Comment:

Junos OS Routing Software Suite [12.1R3.5]

Information for jswitch-ex:

Comment:

Junos OS Enterprise Software Suite [12.1R3.5]

Information for junos:
 Comment:
 Junos OS Base OS boot [12.1R3.5]

Information for jweb-ex:
 Comment:
 Junos OS Web Management [12.1R3.5]

show vlan

Example 5-106 shows the VLAN assignments for the used ports on the switch.

Example 5-106 EX4500_1 show vlan output

Name	Tag	Interfaces
DATA	4092	ae0.0*, ae1.0*, ae2.0*
default		None

show interface terse

Example 5-107 shows the full interface table, which lists port admin status, link status, and so on for the EX4500_1 switch.

Example 5-107 EX4500_1 show interface terse output

Interface	Admin	Link	Proto	Local	Remote
xe-0/0/0	up	up			
xe-0/0/0.0	up	up	aenet	--> ae1.0	
xe-0/0/1	up	up			
xe-0/0/1.0	up	up	aenet	--> ae1.0	
xe-0/0/2	up	up			
xe-0/0/2.0	up	up	aenet	--> ae2.0	
xe-0/0/3	up	up			
xe-0/0/3.0	up	up	aenet	--> ae2.0	
xe-0/0/4	up	up			
xe-0/0/4.0	up	up	aenet	--> ae0.0	
xe-0/0/5	up	up			
xe-0/0/5.0	up	up	aenet	--> ae0.0	
xe-0/0/8	up	down			
xe-0/0/9	up	down			
xe-0/0/10	up	down			
xe-0/0/11	up	down			
xe-0/0/12	up	down			
xe-0/0/13	up	down			
xe-0/0/14	up	down			
xe-0/0/15	up	down			
xe-0/0/16	up	down			
xe-0/0/17	up	down			
xe-0/0/38	up	down			
vcp-0	up	down			
vcp-0.32768	up	down			
vcp-1	up	down			
vcp-1.32768	up	down			

```

ae0                up    up
ae0.0              up    up    eth-switch
ae1                up    up
ae1.0              up    up    eth-switch
ae2                up    up
ae2.0              up    up    eth-switch
ae3                up    down
ae4                up    down
bme0               up    up
bme0.32768         up    up    inet    128.0.0.1/2
                                   128.0.0.16/2
                                   128.0.0.32/2
                                   tnp    0x10
dsc                up    up
gre                up    up
ipip               up    up
lo0                up    up
lsi                up    up
me0                up    up
me0.0              up    up    inet    172.25.101.241/16
mtun               up    up
pimd               up    up
pime               up    up
tap                up    up
vlan               up    up
vlan.4092          up    up    inet    10.1.4.241/24
vme                up    down

```

show lldp neighbors

Example 5-108 lists the LLDP information and serves as a means to verify our physical connectivity.

Example 5-108 EX4500_1 show lldp neighbors output

Local Interface	Parent Interface	Chassis Id	Port info	System Name
xe-0/0/2.0	ae2.0	08:17:f4:33:75:00	Ethernet22	G8264TOR-2
xe-0/0/3.0	ae2.0	08:17:f4:33:75:00	Ethernet24	G8264TOR-2
xe-0/0/0.0	ae1.0	08:17:f4:33:9d:00	Ethernet18	G8264TOR-1
xe-0/0/1.0	ae1.0	08:17:f4:33:9d:00	Ethernet20	G8264TOR-1
xe-0/0/4.0	ae0.0	80:71:1f:d6:ad:40	xe-0/0/4.0	EX4500-2
xe-0/0/5.0	ae0.0	80:71:1f:d6:ad:40	xe-0/0/5.0	EX4500-2
me0.0	-	fc:cf:62:40:a6:00	24	BNT-AS-PM

show spanning-tree interface on EX4500_1

Example 5-109 shows the spanning-tree output on the EX4500_1 switch. As indicated in the output, EX4500_1 is in a designated role and forwarding state from a spanning tree perspective on all three logical interfaces and is the root bridge for VLAN 4092.

Example 5-109 EX4500_1 show spanning-tree interface output

Spanning tree interface parameters for VLAN 4092

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae0.0	128:1	128:1	12284.28c0da3c6441	1000	FWD	DESG
ae1.0	128:2	128:2	12284.28c0da3c6441	1000	FWD	DESG
ae2.0	128:3	128:3	12284.28c0da3c6441	1000	FWD	DESG

show spanning-tree interface on EX4500_2

Example 5-110 shows the spanning tree output on the EX4500_2 switch. As indicated in the output, EX4500_2 is in root role and forwarding state for ae0, alternative role and blocking state for ae1, and designated role and forwarding state for ae2.

Example 5-110 EX4500_2 show spanning-tree output

Spanning tree interface parameters for VLAN 4092

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae0.0	128:1	128:1	12284.28c0da3c6441	1000	FWD	ROOT
ae1.0	128:2	128:1028	32894.0817f4337500	1000	BLK	ALT
ae2.0	128:3	128:3	20476.80711fd6ad41	1000	FWD	DESG

Ping output for equipment on VLAN 4092

To verify connectivity, we issued several ping commands to devices in the lab infrastructure on VLAN 4092 (Data VLAN), as shown in Example 5-111.

Example 5-111 Ping verification for equipment on VLAN 4092

```
admin@EX4500-1> 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=64 time=1.240 ms
64 bytes from 10.1.4.10: icmp_seq=1 ttl=64 time=0.743 ms
64 bytes from 10.1.4.10: icmp_seq=2 ttl=64 time=0.784 ms
64 bytes from 10.1.4.10: icmp_seq=3 ttl=64 time=0.742 ms
64 bytes from 10.1.4.10: icmp_seq=4 ttl=64 time=2.120 ms
64 bytes from 10.1.4.10: icmp_seq=5 ttl=64 time=0.749 ms
64 bytes from 10.1.4.10: icmp_seq=6 ttl=64 time=0.790 ms
64 bytes from 10.1.4.10: icmp_seq=7 ttl=64 time=0.749 ms
64 bytes from 10.1.4.10: icmp_seq=8 ttl=64 time=0.803 ms
64 bytes from 10.1.4.10: icmp_seq=9 ttl=64 time=0.783 ms
^C
--- 10.1.4.10 ping statistics ---
10 packets transmitted, 10 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.742/0.950/2.120/0.415 ms

admin@EX4500-1> 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=255 time=1.521 ms
```

```

64 bytes from 10.1.4.243: icmp_seq=1 ttl=255 time=2.441 ms
64 bytes from 10.1.4.243: icmp_seq=2 ttl=255 time=14.814 ms
64 bytes from 10.1.4.243: icmp_seq=3 ttl=255 time=7.499 ms
64 bytes from 10.1.4.243: icmp_seq=4 ttl=255 time=1.191 ms
^C
--- 10.1.4.243 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.191/5.493/14.814/5.187 ms

admin@EX4500-1> 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=254 time=1.872 ms
64 bytes from 10.1.4.239: icmp_seq=1 ttl=254 time=3.352 ms
64 bytes from 10.1.4.239: icmp_seq=2 ttl=254 time=1.301 ms
64 bytes from 10.1.4.239: icmp_seq=3 ttl=254 time=2.252 ms
64 bytes from 10.1.4.239: icmp_seq=4 ttl=254 time=2.251 ms
^C
--- 10.1.4.239 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.301/2.206/3.352/0.671 ms

```

5.4.9 Full configuration files

In this section, we show the configuration on all of the devices in the Network Topology diagram in Figure 5-4 on page 75.

EN4093 flex-1

Example 5-112 lists the configuration for the EN4093 flex-1 switch.

Example 5-112 EN4093 flex-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 h1thchk"
    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 h1thchk"
    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
!
ip gateway 1 address 10.1.4.241
ip gateway 1 enable
!
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 1ccca6b4f177661b42d268b55d2cfea55ca43558622
!
ntp trusted-key 8811
!
end

```

EN4093 flex_2

Example 5-113 lists the configuration for the EN4093 flex_2 switch.

Example 5-113 EN4093 flex_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 h1thchk"
        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
!
ip gateway 1 address 10.1.4.241
ip gateway 1 enable
!
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 4922dc04ddaa96d7fdd5a63571d5ce8f3ac7dfd1c4
!
ntp trusted-key 8811
!
end

```

G8264tor_1

Example 5-114 lists the configuration for the G8264tor_1 switch.

Example 5-114 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 "Po to EX4500-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 20
    name "Po to EX4500-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 22
    name "Po to EX4500-2"
    tagging

```

```

        tag-pvid
        pvid 4092
        exit
    !
interface port 24
    name "Po to EX4500-2"
    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 18
        lacp key 5
!
interface port 20
        lacp key 5
!
interface port 22
        lacp key 6
!
interface port 24
        lacp key 6
!
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 1
!      addr <default>
!      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 1 address 10.1.4.241
ip gateway 1 enable
!
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
!
!
end

```

G8264tor_2

Example 5-115 lists the configuration for the G8264tor_2 switch.

Example 5-115 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 EX4500-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 20
    name "Po5 to EX4500-2"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 22
    name "Po6 to EX4500-1"
    tagging
    tag-pvid
    pvid 4092
    exit
!
interface port 24
    name "Po6 to EX4500-1"
    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 h1thchk"
        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 18
        lacp key 5
!
interface port 20
        lacp key 5
!
interface port 22
        lacp key 6
!
interface port 24
        lacp key 6
!
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 1
!     addr <default>
!     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 1 address 10.1.4.241
!

```

```
ip gateway 4 address 172.25.1.1
ip gateway 4 enable
!
end
```

EX4500_1 switch

Example 5-116 lists the configuration of the EX4500_1 switch.

Example 5-116 EX4500_1 switch configuration

```
set version 12.1R3.5
set system host-name EX4500-1
set system root-authentication encrypted-password
"$1$bDC1FGB1$DjIV0w9ICEEJ4Q4Ivcgpm1"
set system services ssh
set system syslog user * any emergency
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set chassis aggregated-devices ethernet device-count 5
set interfaces xe-0/0/0 ether-options 802.3ad ae1
set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/2 ether-options 802.3ad ae2
set interfaces xe-0/0/3 ether-options 802.3ad ae2
set interfaces xe-0/0/4 ether-options 802.3ad ae0
set interfaces xe-0/0/5 ether-options 802.3ad ae0
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members DATA
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members DATA
set interfaces ae2 unit 0 family ethernet-switching port-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members DATA
set interfaces me0 unit 0 family inet address 172.25.101.241/16
set interfaces vlan unit 4092 family inet address 10.1.4.241/24
set routing-options static route 0.0.0.0/0 next-hop 172.25.1.1
set protocols igmp-snooping vlan all
set protocols dcbx interface all
set protocols vstp vlan all bridge-priority 8k
set protocols lldp interface all
set protocols lldp-med interface all
set ethernet-switching-options storm-control interface all
set vlans DATA vlan-id 4092
set vlans DATA l3-interface vlan.4092
```

EX4500_2 switch

Example 5-117 lists the configuration of the EX4500_2 switch.

Example 5-117 EX4500_2 switch configuration

```
set version 12.1R3.5
set system host-name EX4500-2
set system root-authentication encrypted-password
"$1$JI1vilgd$b5RdccEiW2Ac1xAXDU3b5/"
set system services ssh
set system syslog user * any emergency
```

```

set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set chassis aggregated-devices ethernet device-count 5
set interfaces xe-0/0/0 ether-options 802.3ad ae1
set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/2 ether-options 802.3ad ae2
set interfaces xe-0/0/3 ether-options 802.3ad ae2
set interfaces xe-0/0/4 ether-options 802.3ad ae0
set interfaces xe-0/0/5 ether-options 802.3ad ae0
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members DATA
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members DATA
set interfaces ae2 unit 0 family ethernet-switching port-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members DATA
set interfaces me0 unit 0 family inet address 172.25.101.242/16
set interfaces vlan unit 4092 family inet address 10.1.4.242/24
set routing-options static route 0.0.0.0/0 next-hop 172.25.1.1
set protocols igmp-snooping vlan all
set protocols dcbx interface all
set protocols vstp vlan all bridge-priority 16k
set protocols lldp interface all
set protocols lldp-med interface all
set ethernet-switching-options storm-control interface all
set vlans DATA vlan-id 4092
set vlans DATA l3-interface vlan.4092

```

5.5 Fully redundant with OSPF

This section describes the various aspects of the implementation scenario.

5.5.1 Topology and requirements

This implementation scenario uses the Layer 3 routing protocol OSPF to provide network connectivity to the G8264 switches. While this design is different from all the presented Layer 2 implementation scenarios, the end-goal of providing a fully redundant infrastructure to the compute nodes still applies. Customers with upstream Juniper equipment that prefer to limit the exposure of Layer 2 to their core or aggregation layer can choose to implement OSPF instead, but with some caveats.

This approach includes the following advantages:

- ▶ Limited Layer 2 exposure to network infrastructure equipment, which limits the ability of a mis-configuration that results 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, which frees up the upstream router to handle only inbound and outbound 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 includes the following disadvantages:

- ▶ Less flexibility in exposing compute nodes to VLANs that might exist on other switches, 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 might not be completely efficient from an address usage perspective.

Components used

Two of each of the following components were used:

- ▶ Juniper EX4500-40F
- ▶ IBM G8264 RackSwitch
- ▶ IBM Flex System Fabric EN4093 10Gb Scalable Switch

5.5.2 Network diagram and physical setup

Figure 5-6 shows the network topology for the fully redundant topology that uses OSPF.

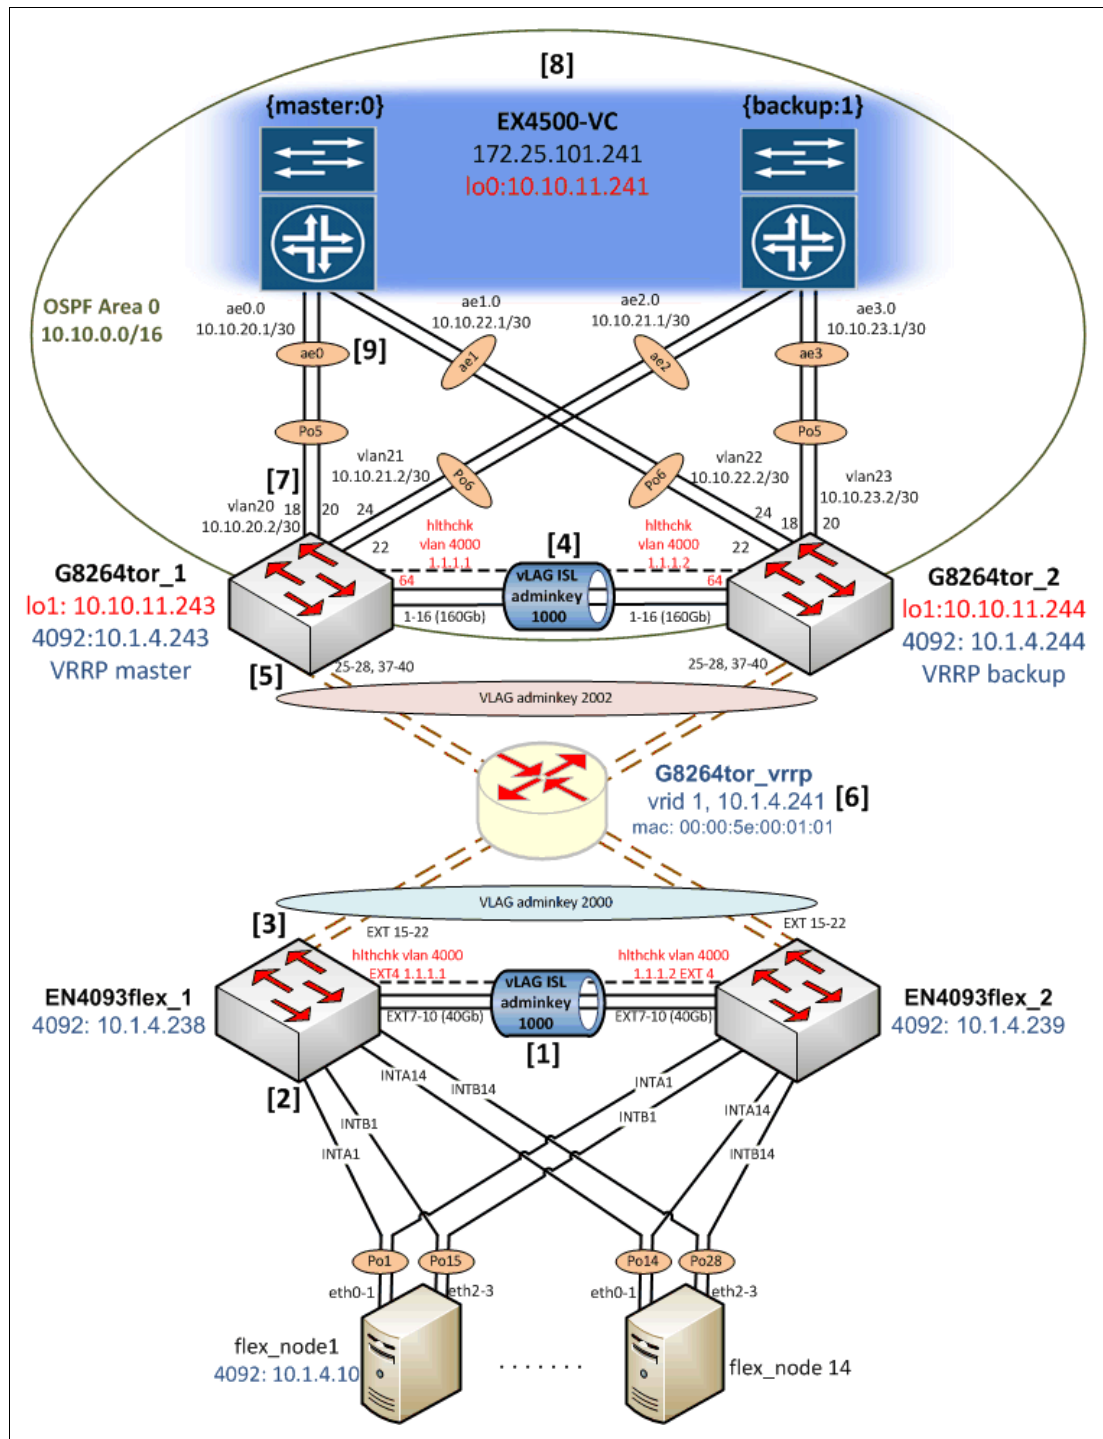


Figure 5-6 Network topology diagram for fully redundant topology that uses OSPF

Start by verifying the physical cabling between the EN4093 switches and G8264s. In our lab environment, we used four IBM QSFP+ DAC Break Out Cables from the EN4093 switches to the upstream G8264s. This requires that the EN4093 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 and 10 Gb SFP+ DAC cables were used for all other connections in the diagram.

5.5.3 EN4093 flex_1 configuration

We begin the implementation of this scenario on the IBM Flex System Fabric EN4093 switches, working our way northward in Figure 5-6 on page 191. Each step provides the commands necessary and are reflective of the numbering schema in the diagram to aid the user in what is configured.

General configuration

Complete the following steps to set up the configuration:

1. Create the ISL Healthcheck, ISL data, and Data VLANs, as shown in Example 5-118. Give them descriptive names, assign them to spanning tree groups, and enable them. You can elect to have the switch create STP instances for you; we chose to manually create them instead.

Example 5-118 Create ISL hlthchk, Data and ISL VLANs on EN4093 flex_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 and Data VLANs in Example 5-119 so that we can verify connectivity between the various pieces of equipment when the configuration is verified. 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, with the last octet borrowed from the network diagram's Management address to quickly aid in the identification of which piece of equipment we are verifying connectivity to.

Example 5-119 Create IP interfaces and assign vlans and IP addresses on EN4093 flex_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 the ISL between EN4093 flex switches

Complete the following steps to configure the ISL between EN4093 flex switches:

1. Configure the eventual ISL in Example 5-120 between the EN4093 flex switches by configuring them to have a default (untagged) VLAN of 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 5-120 Initial ISL configuration on EN4093 flex_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
```

2. Configure the dedicated health check physical interface in Example 5-59 to be used for heartbeats between the EN4093 switches. We chose EXT4 as a dedicated interface and VLAN 4000 to serve as the health check for the ISL.

Example 5-121 Create vLAG hlthchk van and interface on EN4093 flex_1

```
configure terminal
interface port ext4
    pvid 4000
    exit
```

3. Disable STP between the EN4093 switches and activate a vLAG between them so that they appear as a single entity to upstream and downstream infrastructure, as shown in Example 5-122, and reference the LACP key configured in the previous step.

Example 5-122 Disable STP and activate ISL vLAG on EN4093 flex_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

Complete the following steps to configure downstream internal node parts:

1. Configure downstream node interfaces in Example 5-123 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 5-123 Downstream internal node port configuration on EN4093 flex_1

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

2. For redundancy we created two port-channels on each of the 14 nodes. Each port channel aggregates two ports, one from each EN4093 flex switch. Port channels 1 - 14 match the A internally labeled ports and port channel 15 - 28 match the B ports, as shown in Example 5-124.

Example 5-124 Node-facing port channel creation and vLAG activation on EN4093 flex_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

Complete the following steps to configure upstream G8264tor facing ports on Layer 2 failover:

1. Set up the upstream G8264tor facing ports in Example 5-125 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 5-125 Upstream G8264tor facing port configuration on EN4093 flex_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
```

2. Activate the vLAG feature for the upstream EN4093 ports so that the G8264s see the EN4093s as a single, virtualized entity, as shown in Example 5-126. Use adminkey 2000, which represents the LACP key bundling ports EXT15-22 together as one.

Example 5-126 Activating the upstream vLAG on G8264tor facing ports on EN4093 flex_1

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

3. Enable Layer 2 failover in Example 5-127, which effectively shuts down the links to the compute nodes should the uplinks for the EN4093 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, which in our case is connected to the other EN4093 switch in the Enterprise Chassis and ensures that redundancy is maintained.

Example 5-127 Enabling L2 failover for the compute nodes on EN4093 flex_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 flex_2 on the other I/O module. The only difference between the EN4093 flex_1 switch and the EN4093 flex_2 switch is the vLAG health check peer address and the Data, and ISL hlthchk vlan ip addresses. To verify EN4093 flex switch configuration, run the show commands that are described in 5.5.7, “Verification and show command output” on page 207.

5.5.4 G8264tor_1 configuration

Next is the configuration of the switch named G8264tor_1. While the G8264 switches are mostly similar from a configuration standpoint, differences exist that warrant more explanation in this particular use case.

General configuration

Complete the following steps to set up the general configuration:

1. Create the Point-to-Point L3 VLANs (20 and 21), ISL Healthcheck, ISL data, and Data VLANs, as shown in Example 5-128. Give them descriptive names, assign them to spanning tree groups, and enable them.

Example 5-128 Create 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 in Example 5-129. 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 our 172 management network. Loopback 1 is created to use as the router-id when the eventual OSPF adjacencies are built.

Example 5-129 Create IP interfaces and assign 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 EX4500-VC switch that is used to build the OSPF adjacencies. We chose to consistently use /30 networks with the .1 address on the upstream EX4500-VCr, and the .2 address on the G8264s, as shown in Example 5-130.

Example 5-130 Create IP interfaces and assign VLANs and IP addresses

```
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 the ISL between G8264tor switches

Complete the following steps to configure the ISL between G8264tor switches:

1. Configure the ISL between the G8264tor switches in Example 5-131. 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 5-131 Initial ISL configuration on G8264tor_1, step [4]

```
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
```

2. 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 5-132 on page 199, which references the LACP key that was configured in the previous step.

Example 5-132 Disable STP and activate 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 EN4093 flex-facing ports

Complete the following steps to configure downstream EN4093 flex-facing ports:

1. Configure the downstream EN4093 flex-facing ports in Example 5-133 to have a default (untagged) VLAN of 4092 (data vlan), with 802.1q tagging enabled and tag the PVID. Add the ability for all member ports to be on VLAN 4092.

Example 5-133 Downstream EN4093 flex-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
```

2. Activate the vLAG for the downstream EN4093 ports so that the 4093s see the G8264s as a single, virtualized entity, as shown in Example 5-134. Use adminkey 2002, which represents the LACP key bundling ports 25-28, and 37-40 together as one.

Example 5-134 Activate the downstream EN4093 flex-facing vLAG on G8264tor_1

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

Configuring Virtual Router Redundancy Protocol

In this scenario, the G8264tor switches function at the boundary 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, we 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 unavailable G8264tor_2 can seamlessly take over the functionality of the 10.1.4.241 gateway router, as shown in Example 5-135 on page 200.

Example 5-135 VRRP configuration on G8264tor_1, step [6]

```
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 EX4500-VC facing ports

Complete the following steps to configure upstream EX4500-VC facing ports:

1. Configure the Layer 3 upstream EX4500-VC facing ports in Example 5-136 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 5-136 Upstream EX4500-VC facing port configuration on G8264tor_1

```
configure terminal
interface port 18,20
  name "Po to EX4500-VC"
  pvid 20
  exit
interface port 22,24
  name "Po to EX4500-VC"
  pvid 21
  exit
```

2. Activate link aggregation groups by using static port-channeling, as shown in Example 5-137. We chose to use static port-channeling here to show that IBM System Networking equipment interoperates with an upstream Juniper infrastructure with LACP or static (no negotiation protocol) port-channeling.

Example 5-137 Create 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

Complete the following steps to set up the OSPF configuration:

1. Set the router-id parameter to the loopback 1 IP address, as shown in Example 5-138. By default, OSPF uses the lowest-configured IP address on the device in all OSPF neighbor advertisements. It is good practice to 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 5-138 Set the router-id on G8264tor_1

```
configure terminal
ip router-id 10.10.11.243
```

2. Create an OSPF instance and advertise OSPF routes through IP interfaces 20, 21, and 92 in Example 5-139 for area 0. OSPF begins to build its associated adjacency matrixes after this step is complete.

Example 5-139 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, we configure the switch that is named G8264tor_2.

General configuration

Complete the following steps to set up the general configuration:

1. Create the Point-to-Point L3 VLANs (22 and 23), ISL Healthcheck, ISL data, and Data VLANs, as shown in Example 5-140. Give them descriptive names, assign them to spanning tree groups, and enable them.

Example 5-140 Create vlan 22, vlan 23, 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 5-141. 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 our 172 management network. Loopback 1 is created to use as the router-id when the eventual OSPF adjacencies are built.

Example 5-141 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 EX4500-VC equipment that is used to build the OSPF adjacencies. We chose to consistently use /30 networks with the .1 address on the upstream EX4500-VC, and the .2 address on the G8264s, as shown in Example 5-142 on page 203.

Example 5-142 Create IP interfaces and assign vlans for vlan 22 and vlan 23 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 the ISL between G8264tor switches

Complete the following step to configure the ISL between G864tor switches:

1. Configure the ISL between the G8264tor switches in Example 5-143. 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 5-143 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
```

2. 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 5-144, which references the LACP key that was configured in the previous step.

Example 5-144 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 flex-facing ports

Complete the following steps to configure the downstream EN4093 flex-facing ports:

1. Configure the downstream EN4093 flex-facing ports in Example 5-133 to have a default (untagged) VLAN of 4092 (data vlan), with 802.1q tagging enabled and tag the PVID. Add the ability for all member ports to be on VLAN 4092.

Example 5-145 Downstream EN4093 flex-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
```

2. Activate the vLAG for the downstream EN4093 ports so that the 4093s see the G8264s as a single, virtualized entity, as shown in Example 5-134. Use adminkey 2002, which represents the LACP key bundling ports 25-28, and 37-40 together as one.

Example 5-146 Activate the downstream EN4093 flex-facing vLAG on G8264tor_2

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

Configuring VRRP

Configure VRRP on G8264tor_2. It functions as the secondary VRRP gateway router for the 10.1.4.241 IP address, as shown in Example 5-147.

Example 5-147 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 EX4500-VC facing ports

Complete the following steps to configure upstream EX4500-VC facing ports:

1. Setup the Layer 3 upstream ports to the EX4500-VC in Example 5-148 on page 205 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 5-148 Upstream EX4500-VC facing configuration on G8264tor_2

```
configure terminal
interface port 18,20
    name "Po5 to EX4500-VC"
    pvid 23
    exit
interface port 22,24
    name "Po6 to EX4500-VC"
    pvid 22
    exit
```

2. Activate link aggregation groups by using static port-channeling, as shown in Example 5-149.

Example 5-149 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

Complete the following steps to set up the OSPF configuration:

1. Set the router-id parameter to be the loopback 1 IP address, as shown in Example 5-150.

Example 5-150 Set router-id on G8264tor_2

```
configure terminal
ip router-id 10.10.11.244
```

2. Create an OSPF instance and enable IP interfaces 23, 23, and 92 for area 0, as shown in Example 5-151.

Example 5-151 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 EX4500-VC switch configuration

We now configure the EX4500-VC switch.

Configuring VC on Juniper switches

Complete the following steps to configure VC on Juniper switches:

1. Interconnect the member switches that are using the dedicated Virtual Chassis Ports (VCPs) that are on the rear side of the switch. Power on only the switch that you plan to use as the master switch (for management purposes, we recommend arranging the switches in member ID sequence from top to bottom).
2. Configure mastership priority for the master and disable the split and merge features (recommended for a two-member VC), as shown in Example 5-152.

Example 5-152 Configuring mastership priority

```
set virtual-chassis member 0 mastership-priority 255
set no-split-detection
```

3. Power on the second EX4500 switch. This switch has the backup role and now both switches work as a single switch.

Important: If you do not edit the VC configuration file, a nonprovisioned configuration is generated by default. The mastership priority value for each member switch is 128. The master role is selected by default.

General configuration

Complete the following steps to set up the general configuration:

1. Assign IP addresses for the management Port and configure an IP gateway for our 172 management network, as shown in Example 5-153.

Example 5-153 Management IP address and IP gateway configuration on EX4500-VC

```
set interfaces vme unit 0 family inet address 172.25.101.241/16
set routing-options static route 0.0.0.0/0 next-hop 172.25.1.1
```

2. An important difference with this scenario is the absence of vlan 4092 (data vlan) on the EX4500-VC switch. As Layer 3 connectivity is used between EX4500-VC and G8264tor switches, the data VLAN does not extend up to the EX4500-VC switch. A loopback interface (as shown in Example 5-154) is created as a server network that is used on the EX4500-VC switch only. This is used to demonstrate connectivity to compute nodes on vlan 4092.

Example 5-154 Loopback interface configuration on EX4500-VC

```
set interfaces lo0 unit 0 family inet address 10.10.11.241/32
```

Configuring downstream G8264tor facing ports

Configure the downstream physical and logical interfaces in Example 5-155 on page 207, with interfaces xe-0/0/0 and xe-0/0/1 to be bundled in static aggregation ae0, interfaces xe-0/0/2 and xe-0/0/3 to be bundled in static aggregation ae1, interfaces xe-1/0/0 and xe-1/0/1 to be bundled in static aggregation ae2, and interfaces xe-1/0/2 and xe-1/0/3. To each logical interface, assign the IP address that is shown in Figure 5-6 on page 191.

Example 5-155 Downstream G8264tor facing interface configuration on EX4500-VC

```
set chassis aggregated-devices ethernet device-count 4
set interfaces xe-0/0/0 description "Connection to 8264_1 Port 18"
set interfaces xe-0/0/0 ether-options 802.3ad ae0
set interfaces xe-0/0/1 description "Connection to 8264_1 Port 20"
set interfaces xe-0/0/1 ether-options 802.3ad ae0
set interfaces xe-0/0/2 description "Connection to 8264_2 Port 22"
set interfaces xe-0/0/2 ether-options 802.3ad ae1
set interfaces xe-0/0/3 description "Connection to 8264_2 Port 24"
set interfaces xe-0/0/3 ether-options 802.3ad ae1
set interfaces xe-1/0/0 description "Connection to 8264_2 Port 18"
set interfaces xe-1/0/0 ether-options 802.3ad ae3
set interfaces xe-1/0/1 description "Connection to 8264_2 Port 20"
set interfaces xe-1/0/1 ether-options 802.3ad ae3
set interfaces xe-1/0/2 description "Connection to 8264_1 Port 22"
set interfaces xe-1/0/2 ether-options 802.3ad ae2
set interfaces xe-1/0/3 description "Connection to 8264_1 Port 24"
set interfaces xe-1/0/3 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family inet address 10.10.20.1/30
set interfaces ae1 unit 0 family inet address 10.10.22.1/30
set interfaces ae2 unit 0 family inet address 10.10.21.1/30
set interfaces ae3 unit 0 family inet address 10.10.23.1/30
```

OSPF configuration

Create the backbone area (area 0) for your OSPF network and add the appropriate interfaces to the area, as shown in Example 5-156. Set the router-id parameter to be the loopback 0 IP address.

Example 5-156 OSPF instance configuration on EX4500-VC

```
set protocols ospf area 0.0.0.0 interface ae0.0
set protocols ospf area 0.0.0.0 interface ae1.0
set protocols ospf area 0.0.0.0 interface ae2.0
set protocols ospf area 0.0.0.0 interface ae3.0
set protocols ospf area 0.0.0.0 interface lo0.0
set routing-options router-id 10.10.11.241
```

5.5.7 Verification and show command output

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

We begin by showing helpful commands from the EN4093 switches, working our way up the Network Topology diagram all the way to the Juniper EX4500-VC switch.

EN4093 output

Here we list output from the switch with hostname EN4093 flex_1. Similar or identical output exists for the switch with hostname EN4093 flex_2.

Show version

The command output in Example 5-157 shows information regarding the switch that we used and the associated code and firmware level at the time.

Example 5-157 EN4093 flex_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 5-158 shows output regarding VLAN assignment for all the various ports on the switch.

Example 5-158 EN4093 flex_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-INTA14 INTB1-INTB14 EXT7-EXT10 EXT15-EXT22
4094	ISL	ena	dis	EXT7-EXT10
4095	Mgmt VLAN	ena	ena	EXTM MGT1

Show interface status

Because we only have one compute node in our chassis (in slot 1), this explains why all the other internal ports are listed as down from a link perspective in the output that is shown in Example 5-159.

Example 5-159 EN4093 flex_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 hlthchk
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 5-160 shows our physical topology and verifies that cables are plugged into the ports we specified in both our Network Topology diagram in Figure 5-4 on page 75 and the configuration that is specified in 5.5.8, “Full configuration files”.

Example 5-160 EN4093 flex_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 5-161 shows command output regarding the status of the ISL between the EN4093 switches and the ports that comprise the ISL.

Example 5-161 EN4093 flex_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 5-162 shows that the vLAG between the EN4093 switches and G8264 switches is up and operational as referenced by the LACP admin key of 2000. Our ISL between the EN4093 switches also is up.

EN4093 flex_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, role election is performed. The switch with primary role controls the centralized operation. Role election is non-preemptive (that is, if a primary 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 priority is the same, switch with smaller system MAC address is the vLAG primary switch. It is possible to configure vLAG priority to anything between 0 - 65535, priority was left at the default value of 0 in all examples.

Example 5-162 EN4093 flex_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 5-163 output shows that the vLAG is formed and enabled by using LACP reference key 2000.

Example 5-163 EN4093 flex_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 5-164 shows which ports are participating in an LACP aggregation and which reference keys are used on those specific interfaces.

Example 5-164 EN4093 flex_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

Failover output showing which ports are monitored and which ports are shutdown should an issue be encountered is shown in Example 5-165 on page 213. In our case, our upstream to G8264 links are monitored with LACP reference key 2000. Our control ports are the downstream internal I/O module ports that are used by the Compute Nodes.

Example 5-165 EN4093 flex_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, it is possible to run the **show arp** command on the EN4093 flex switches. Here we can see that the VRRP IP gateway address 10.1.4.241 is present in the ARP table. The MAC address that is 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, which is 01 in this case, as shown in Example 5-166.

Example 5-166 EN4093 flex_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, we issued ping commands to devices in the lab infrastructure on VLAN 4092 (Data VLAN) in Example 5-167. IP address 10.4.1.10 represents a Compute Node with an operating system installed, flex_node1 on the Network Topology diagram in Figure 5-4 on page 75.

Example 5-167 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 10.1.4.244 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.
```

G8264 output

Here we list output from the switch with hostname G8264tor_1 and note specific differences on G8264tor_2, where applicable.

Show version

Example 5-168 shows information regarding the switch that we used and the associated code and firmware level at that time.

Example 5-168 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 5-169 shows VLAN assignments for all the various ports on G8264tor_1.

Example 5-169 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 5-170 shows VLAN assignments for all the various ports on G8264tor_2.

Example 5-170 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 h1thchk	ena	64
4092	DATA	ena	1-16 25-28 37-40
4094	ISL	ena	1-16
4095	Mgmt VLAN	ena	MGT

Show interface status

All the other internal ports are listed as down from a link perspective in the output that is shown in Example 5-171 because we have only one compute node in our chassis (in slot 1).

Example 5-171 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	Po to EX4500-VC
19	19	1G/10G	full	no	no	down	19
20	20	10000	full	no	no	up	Po to EX4500-VC
21	21	1G/10G	full	no	no	down	21
22	22	10000	full	no	no	up	Po to EX4500-VC
23	23	1G/10G	full	no	no	down	23
24	24	10000	full	no	no	up	Po to EX4500-VC
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 5-172 shows our physical topology and verifies that cables are plugged into the ports we specified in our Network Topology diagram in Figure 5-4 on page 75, and the configuration specified in Appendix A, “Fully Redundant Implementation with IBM vLAG and Juniper MC-LAG Technology” on page 293.

Example 5-172 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	5	08 17 f4 33 75 00	5	G8264TOR-2
6	6	08 17 f4 33 75 00	6	G8264TOR-2
7	7	08 17 f4 33 75 00	7	G8264TOR-2
MGT	8	fc cf 62 40 a6 00	22	BNT-AS-PM
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
15	16	08 17 f4 33 75 00	15	G8264TOR-2
16	17	08 17 f4 33 75 00	16	G8264TOR-2
64	18	08 17 f4 33 75 00	64	G8264TOR-2
18	19	80 71 1f d6 ad 40	505	EX4500-VC
20	20	80 71 1f d6 ad 40	507	EX4500-VC
22	21	80 71 1f d6 ad 40	618	EX4500-VC
24	22	80 71 1f d6 ad 40	619	EX4500-VC
25	23	6c ae 8b bf 6d 00	58	en4093flex_1
26	24	6c ae 8b bf 6d 00	57	en4093flex_1
27	25	6c ae 8b bf 6d 00	60	en4093flex_1
28	26	6c ae 8b bf 6d 00	59	en4093flex_1

37	27	6c ae 8b bf fe 00	57	en4093f1ex_2
38	28	6c ae 8b bf fe 00	59	en4093f1ex_2
39	29	6c ae 8b bf fe 00	58	en4093f1ex_2
40	30	6c ae 8b bf fe 00	60	en4093f1ex_2

show vlag isl

The command output in Example 5-173 shows the status of the ISL between the G8264 switches and the ports that comprise the ISL.

Example 5-173 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 5-174 shows that the downstream vLAG between the G8264 and EN4093 switches is up and operational, as referenced by the LACP admin key of 2002. Our ISL between the G8264 switches also is up.

Example 5-174 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
```

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, role election is performed. The switch with primary role controls the centralized operation. Role election is non-preemptive (that is, if a primary 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 priority is the same, switch with smaller system MAC address is the vLAG primary switch. It is possible to configure vLAG priority to anything between 0 - 65535; priority was left at the default value of 0 in all examples.

show vlag adminkey 2002

The output in Example 5-175 shows that the downstream vLAG towards the EN4093 switches is formed and enabled by using LACP reference key 2002.

Example 5-175 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 5-176 shows which ports are participating in an LACP aggregation and which reference keys are used on those specific interfaces.

Example 5-176 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 5-177 lists output from the **show ip ospf neighbor** command, which shows that OSPF is enabled and displaying associated neighbor information. We can use this information to verify our Network Topology diagram in Figure 5-4 on page 77.

Example 5-177 G8264tor_1 show ip ospf neighbor output

Intf	NeighborID	Prio	State	Address
----	-----	----	-----	-----
20	10.10.11.241	128	Full	10.10.20.1
21	10.10.11.241	128	Full	10.10.21.1
92	10.10.11.244	1	Full	10.1.4.244

Show ip ospf neighbor on G8264tor_2

Example 5-178 lists output from the **show ip ospf neighbor** command on the second G8264tor_2 switch.

Example 5-178 G8264tor_2 show ip ospf neighbor output

Intf	NeighborID	Prio	State	Address
----	-----	----	-----	-----
22	10.10.11.241	128	Full	10.10.22.1
23	10.10.11.241	128	Full	10.10.23.1
92	10.10.11.243	1	Full	10.1.4.243

Show ip ospf routes on G8264tor_1

Example 5-179 on page 222 lists output from the **show ip ospf routes** command, which shows learned routes identified via neighboring interfaces. Here you can see the learned route 10.10.11.241/32, which is the route that we configured in the EX4500-VC switch to simulate a server network.

Example 5-179 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
* 10.10.11.241/32 via 10.10.21.1
* 10.10.11.241/32 via 10.10.20.1
* 10.10.22.0/30 via 10.1.4.244
* 10.10.22.0/30 via 10.10.21.1
* 10.10.22.0/30 via 10.10.20.1
* 10.10.23.0/30 via 10.1.4.244
* 10.10.23.0/30 via 10.10.21.1
* 10.10.23.0/30 via 10.10.20.1
```

Show ip ospf routes on G8264tor_2

Example 5-180 lists output from the **show ip ospf routes** command on the 2nd G8264tor_2 switch. You can see all the learned route and a route to 10.10.11.241/32, which is the route we configure in the EX4500-VC switch to simulate a server network.

Example 5-180 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
* 10.10.11.241/32 via 10.10.23.1
* 10.10.11.241/32 via 10.10.22.1
* 10.10.20.0/30 via 10.1.4.243
* 10.10.20.0/30 via 10.10.23.1
* 10.10.20.0/30 via 10.10.22.1
* 10.10.21.0/30 via 10.1.4.243
* 10.10.21.0/30 via 10.10.23.1
* 10.10.21.0/30 via 10.10.22.1
```

show ip ospf interface on G8264tor_1

OSPF interface-related information for G8264tor_1 is shown in Example 5-181, which shows which interfaces participate in the OSPF process.

Example 5-181 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) 10.10.11.241, 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 2, Authentication type none
```

Ip Address 10.10.21.2, Area 0.0.0.0, Admin Status UP
Router ID 10.10.11.243, State BackupDR, Priority 1
Designated Router (ID) 10.10.11.241, Ip Address 10.10.21.1
Backup Designated Router (ID) 10.10.11.243, Ip Address 10.10.21.2
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 shown in Example 5-182 for the 2nd G8264 switch.

Example 5-182 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) 10.10.11.241, 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 2, Authentication type none

Ip Address 10.10.23.2, Area 0.0.0.0, Admin Status UP
Router ID 10.10.11.244, State BackupDR, Priority 1
Designated Router (ID) 10.10.11.241, Ip Address 10.10.23.1
Backup Designated Router (ID) 10.10.11.244, Ip Address 10.10.23.2
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 5-183 confirms that G8264tor_1 is the master for Virtual IP Address (VIP) 10.1.4.241.

Example 5-183 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 5-184 confirms that G8264tor_2 is the backup.

Example 5-184 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, we issued several ping commands to devices in the lab infrastructure on VLAN 4092 (Data VLAN) in Example 5-185. IP address 10.4.1.10 represents a Compute Node with an operating system installed, flex_node1 on the Network Topology diagram in Figure 5-4 on page 77.

Example 5-185 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 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.

G8264TOR-1#ping 10.1.4.241 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 7 msec.

10.1.4.241: #2 ok, RTT 4 msec.

10.1.4.241: #3 ok, RTT 4 msec.

10.1.4.241: #4 ok, RTT 13 msec.

10.1.4.241: #5 ok, RTT 4 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 1 msec.

10.1.4.238: #2 ok, RTT 0 msec.

10.1.4.238: #3 ok, RTT 0 msec.

10.1.4.238: #4 ok, RTT 4 msec.

10.1.4.238: #5 ok, RTT 0 msec.

Ping finished.

```
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 0 msec.
10.1.4.239: #2 ok, RTT 7 msec.
10.1.4.239: #3 ok, RTT 2 msec.
10.1.4.239: #4 ok, RTT 0 msec.
10.1.4.239: #5 ok, RTT 11 msec.
Ping finished.
```

EX4500-VC output

Here we list output from the Juniper switch with hostname EX4500-VC.

show system software

Example 5-186 shows information regarding the switch that we used and the associated code and firmware level at that time.

Example 5-186 EX4500-VC show system software output

fpc0:

Information for fips-mode-powerpc:

Comment:

Junos OS FIPS mode utilities [12.1R3.5]

Information for jbase:

Comment:

Junos OS Base OS Software Suite [12.1R3.5]

Information for jcrypto-ex:

Comment:

Junos OS Crypto Software Suite [12.1R3.5]

Information for jdocs-ex:

Comment:

Junos OS Online Documentation [12.1R3.5]

Information for jkernel-ex:

Comment:

Junos OS Kernel Software Suite [12.1R3.5]

Information for jpfe-ex45x:

Comment:

Junos OS Packet Forwarding Engine Enterprise Software Release

Copyright (c) 1996-2012, Juniper Networks, Inc.

All rights reserved.

Junos OS Packet Forwarding Engine Enterprise Software Suite for EX45xx series
[12.1R3.5]

Information for jroute-ex:

Comment:

Junos OS Routing Software Suite [12.1R3.5]

Information for jswitch-ex:

Comment:

Junos OS Enterprise Software Suite [12.1R3.5]

Information for junos:

Comment:

Junos OS Base OS boot [12.1R3.5]

Information for jweb-ex:

Comment:

Junos OS Web Management [12.1R3.5]

fpc1:

Information for fips-mode-powerpc:

Comment:

Junos OS FIPS mode utilities [12.1R3.5]

Information for jbase:

Comment:

Junos OS Base OS Software Suite [12.1R3.5]

Information for jcrypto-ex:

Comment:

Junos OS Crypto Software Suite [12.1R3.5]

Information for jdocs-ex:

Comment:

Junos OS Online Documentation [12.1R3.5]

Information for jkernel-ex:

Comment:

Junos OS Kernel Software Suite [12.1R3.5]

Information for jpfe-ex45x:

Comment:

Junos OS Packet Forwarding Engine Enterprise Software Release
Copyright (c) 1996-2012, Juniper Networks, Inc.

All rights reserved.

Junos OS Packet Forwarding Engine Enterprise Software Suite for EX45xx series
[12.1R3.5]

Information for jroute-ex:

Comment:

Junos OS Routing Software Suite [12.1R3.5]

Information for jswitch-ex:

Comment:

Junos OS Enterprise Software Suite [12.1R3.5]

Information for junos:

Comment:

Junos OS Base OS boot [12.1R3.5]

Information for jweb-ex:

Comment:

Junos OS Web Management [12.1R3.5]

show interface terse

Example 5-187 shows the full interface table, which lists administrative status, link status, and so on for the EX4500-VC switch.

Example 5-187 EX4500-VC show interface terse output

Interface	Admin	Link	Proto	Local	Remote
xe-0/0/0	up	up			
xe-0/0/0.0	up	up	aenet	--> ae0.0	
xe-0/0/1	up	up			
xe-0/0/1.0	up	up	aenet	--> ae0.0	
xe-0/0/2	up	up			
xe-0/0/2.0	up	up	aenet	--> ae1.0	
xe-0/0/3	up	up			
xe-0/0/3.0	up	up	aenet	--> ae1.0	
xe-0/0/4	up	up			
xe-0/0/5	up	up			
xe-0/0/8	up	down			
xe-0/0/9	up	down			
xe-0/0/10	up	down			
xe-0/0/11	up	down			
xe-0/0/12	up	down			
xe-0/0/13	up	down			
xe-0/0/14	up	down			
xe-0/0/15	up	down			
xe-0/0/16	up	down			
xe-0/0/17	up	down			
xe-0/0/38	up	down			
xe-1/0/0	up	up			
xe-1/0/0.0	up	up	aenet	--> ae3.0	

xe-1/0/1	up	up		
xe-1/0/1.0	up	up	aenet	--> ae3.0
xe-1/0/2	up	up		
xe-1/0/2.0	up	up	aenet	--> ae2.0
xe-1/0/3	up	up		
xe-1/0/3.0	up	up	aenet	--> ae2.0
xe-1/0/4	up	up		
xe-1/0/5	up	up		
xe-1/0/8	up	down		
xe-1/0/9	up	down		
xe-1/0/10	up	down		
xe-1/0/11	up	down		
xe-1/0/12	up	down		
xe-1/0/13	up	down		
xe-1/0/14	up	down		
xe-1/0/15	up	down		
xe-1/0/16	up	down		
xe-1/0/17	up	down		
xe-1/0/38	up	down		
vcp-0	up	up		
vcp-0.32768	up	up		
vcp-1	up	up		
vcp-1.32768	up	up		
ae0	up	up		
ae0.0	up	up	inet	10.10.20.1/30
ae1	up	up		
ae1.0	up	up	inet	10.10.22.1/30
ae2	up	up		
ae2.0	up	up	inet	10.10.21.1/30
ae3	up	up		
ae3.0	up	up	inet	10.10.23.1/30
ae4	up	down		
bme0	up	up		
bme0.32768	up	up	inet	128.0.0.1/2 128.0.0.16/2 128.0.0.32/2
			tnp	0x10
bme0.32770	down	up	eth-switch	
dsc	up	up		
gre	up	up		
ipip	up	up		
lo0	up	up		
lo0.0	up	up	inet	10.10.30.1 --> 0/0
lsi	up	up		
me0	up	up		
me0.0	up	up	eth-switch	
mtun	up	up		
pimd	up	up		
pime	up	up		
tap	up	up		
vlan	up	up		
vme	up	up		
vme.0	up	up	inet	172.25.101.241/16

show lldp neighbors on EX4500-VC

Example 5-188 lists the LLDP information and serves as a means to verify our physical connectivity.

Example 5-188 EX4500-VC show lldp neighbors output

Local Interface	Parent Interface	Chassis Id	Port info	System Name
xe-1/0/0.0	ae3.0	08:17:f4:33:75:00	Ethernet18	G8264TOR-2
xe-1/0/1.0	ae3.0	08:17:f4:33:75:00	Ethernet20	G8264TOR-2
xe-0/0/2.0	ae1.0	08:17:f4:33:75:00	Ethernet22	G8264TOR-2
xe-0/0/3.0	ae1.0	08:17:f4:33:75:00	Ethernet24	G8264TOR-2
xe-0/0/0.0	ae0.0	08:17:f4:33:9d:00	Ethernet18	G8264TOR-1
xe-0/0/1.0	ae0.0	08:17:f4:33:9d:00	Ethernet20	G8264TOR-1
xe-1/0/2.0	ae2.0	08:17:f4:33:9d:00	Ethernet22	G8264TOR-1
xe-1/0/3.0	ae2.0	08:17:f4:33:9d:00	Ethernet24	G8264TOR-1
vme.0	-	fc:cf:62:40:a6:00	24	BNT-AS-PM

show ospf interface on EX4500-VC

Example 5-189 shows the interfaces on which OSPF is configured on the EX4500-VC switch.

Example 5-189 EX4500-VC show ospf interface output

Interface	State	Area	DR ID	BDR ID	Nbrs
ae0.0	DR	0.0.0.0	10.10.11.241	10.10.11.243	1
ae1.0	DR	0.0.0.0	10.10.11.241	10.10.11.244	1
ae2.0	DR	0.0.0.0	10.10.11.241	10.10.11.243	1
ae3.0	DR	0.0.0.0	10.10.11.241	10.10.11.244	1
lo0.0	DR	0.0.0.0	10.10.11.241	0.0.0.0	0

show ospf neighbor on EX4500-VC

Example 5-190 shows the OSPF neighbor data from the perspective of the EX4500-VC switch.

Example 5-190 EX4500-VC show ospf neighbor output

Address	Interface	State	ID	Pri	Dead
10.10.20.2	ae0.0	Full	10.10.11.243	1	34
10.10.22.2	ae1.0	Full	10.10.11.244	1	38
10.10.21.2	ae2.0	Full	10.10.11.243	1	34
10.10.23.2	ae3.0	Full	10.10.11.244	1	39

show route protocol ospf for EX4500-VC

Example 5-191 on page 230 list the routes that were learned by OSPF for EX4500-VC. For network 10.1.4.0/24, it has learned four different routes, but only one is active in the forwarding table, (the route marked with >). If you want to balance traffic through the different learned route, you can enable the load balance feature.

Example 5-191 EX4500-VC show route protocol ospf output

inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```
10.1.4.0/24      *[OSPF/10] 00:33:57, metric 2
                  > to 10.10.20.2 via ae0.0
                  to 10.10.21.2 via ae2.0
                  to 10.10.23.2 via ae3.0
                  to 10.10.22.2 via ae1.0
224.0.0.5/32    *[OSPF/10] 16:48:48, metric 1
                  MultiRecv
```

show route forwarding-table for EX4500-VC

Example 5-192 shows the forwarding table, including the network-layer prefixes and their next hops for EX4500-VC.

Example 5-192 EX4500-VC show route forwarding-table output

Routing table: default.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	user	1	fc:cf:62:10:b2:0	ucst	336	4	vme.0
default	perm	0		rjct	36	1	
0.0.0.0/32	perm	0		dscd	34	1	
10.1.4.0/24	user	0	10.10.20.2	ucst	1316	3	ae0.0
10.10.20.0/30	intf	0		rslv	1315	1	ae0.0
10.10.20.0/32	dest	0	10.10.20.0	recv	1313	1	ae0.0
10.10.20.1/32	intf	0	10.10.20.1	loc1	1314	2	
10.10.20.1/32	dest	0	10.10.20.1	loc1	1314	2	
10.10.20.2/32	dest	0	8:17:f4:33:9d:0	ucst	1316	3	ae0.0
10.10.20.3/32	dest	0	10.10.20.3	bcst	1307	1	ae0.0
10.10.21.0/30	intf	0		rslv	1354	1	ae2.0
10.10.21.0/32	dest	0	10.10.21.0	recv	1352	1	ae2.0
10.10.21.1/32	intf	0	10.10.21.1	loc1	1353	2	
10.10.21.1/32	dest	0	10.10.21.1	loc1	1353	2	
10.10.21.2/32	dest	0	8:17:f4:33:9d:0	ucst	1334	1	ae2.0
10.10.21.3/32	dest	0	10.10.21.3	bcst	1351	1	ae2.0
10.10.22.0/30	intf	0		rslv	1350	1	ae1.0
10.10.22.0/32	dest	0	10.10.22.0	recv	1348	1	ae1.0
10.10.22.1/32	intf	0	10.10.22.1	loc1	1349	2	
10.10.22.1/32	dest	0	10.10.22.1	loc1	1349	2	
10.10.22.2/32	dest	0	8:17:f4:33:75:0	ucst	1328	1	ae1.0
10.10.22.3/32	dest	0	10.10.22.3	bcst	1347	1	ae1.0
10.10.23.0/30	intf	0		rslv	1345	1	ae3.0
10.10.23.0/32	dest	0	10.10.23.0	recv	1343	1	ae3.0
10.10.23.1/32	intf	0	10.10.23.1	loc1	1344	2	
10.10.23.1/32	dest	0	10.10.23.1	loc1	1344	2	
10.10.23.2/32	dest	0	8:17:f4:33:75:0	ucst	1317	1	ae3.0
10.10.23.3/32	dest	0	10.10.23.3	bcst	1342	1	ae3.0
10.10.30.1/32	intf	0	10.10.30.1	loc1	1335	1	
172.25.0.0/16	intf	0		rslv	331	1	vme.0
172.25.0.0/32	dest	0	172.25.0.0	recv	324	1	vme.0
172.25.1.1/32	dest	0	fc:cf:62:10:b2:0	ucst	336	4	vme.0
172.25.101.241/32	intf	0	172.25.101.241	loc1	328	2	
172.25.101.241/32	dest	0	172.25.101.241	loc1	328	2	

172.25.255.255/32	dest	0	172.25.255.255	bcst	321	1 vme.0
224.0.0.0/4	perm	1		mdsc	35	1
224.0.0.1/32	perm	0	224.0.0.1	mcst	31	4
224.0.0.5/32	user	1	224.0.0.5	mcst	31	4
224.0.0.22/32	user	0	224.0.0.22	mcst	31	4
255.255.255.255/32	perm	0		bcst	32	1
...						
...						
...						

Ping output for equipment on VLAN 4092

To verify connectivity, we issued several ping commands to devices in the lab infrastructure on VLAN 4092 (Data VLAN) in Example 5-193. Included is our Compute Node with an assigned IP address of 10.1.4.10.

Example 5-193 Ping verification for equipment on VLAN 4092

```

root@EX4500-VC> 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=3.405 ms
64 bytes from 10.1.4.10: icmp_seq=1 ttl=63 time=0.909 ms
64 bytes from 10.1.4.10: icmp_seq=2 ttl=63 time=0.911 ms
64 bytes from 10.1.4.10: icmp_seq=3 ttl=63 time=0.923 ms
64 bytes from 10.1.4.10: icmp_seq=4 ttl=63 time=0.916 ms
^C
--- 10.1.4.10 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.909/1.413/3.405/0.996 ms

{master:0}
root@EX4500-VC>

{master:0}
root@EX4500-VC>
root@EX4500-VC> 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=3.012 ms
64 bytes from 10.1.4.238: icmp_seq=1 ttl=254 time=1.782 ms
64 bytes from 10.1.4.238: icmp_seq=2 ttl=254 time=1.223 ms
64 bytes from 10.1.4.238: icmp_seq=3 ttl=254 time=1.124 ms
64 bytes from 10.1.4.238: icmp_seq=4 ttl=254 time=1.180 ms
^C
--- 10.1.4.238 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.124/1.664/3.012/0.714 ms

{master:0}
root@EX4500-VC>
root@EX4500-VC> 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=254 time=2.623 ms
64 bytes from 10.1.4.239: icmp_seq=1 ttl=254 time=7.734 ms
64 bytes from 10.1.4.239: icmp_seq=2 ttl=254 time=1.838 ms
64 bytes from 10.1.4.239: icmp_seq=3 ttl=254 time=1.855 ms
^C

```

```
--- 10.1.4.239 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.838/3.513/7.734/2.458 ms
```

5.5.8 Full configuration files

In this section, we show the configuration on all of the devices in the Network Topology diagram in Figure 5-4 on page 77.

EN4093 flex-1

Example 5-194 shows the configuration for the EN4093 flex-1 switch.

Example 5-194 EN4093 flex-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 h1thchk"
    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
!
ip gateway 1 address 10.1.4.241

```

```

ip gateway 1 enable
!
!
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 4ccca6b4f177661b42d268b55d2cfea55ca43558622
!
ntp trusted-key 8811
!
end

```

EN4093 flex_2

Example 5-195 lists the configuration for the EN4093 flex_2 switch.

Example 5-195 EN4093 flex_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 h1thchk"
    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
!
ip gateway 1 address 10.1.4.241
ip gateway 1 enable
!
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 44922dc04ddaa96d7fdd5a63571d5ce8f3ac7dfd1c4
!
ntp trusted-key 8811
!
end

```

G8264tor_1

Example 5-196 lists the configuration for the G8264tor_1 switch.

Example 5-196 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 18
    lacp key 5
    !
interface port 20
    lacp key 5

```

```

!
!
interface port 22
    lacp key 6
!
interface port 24
    lacp key 6
!
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 5-197 lists the configuration for the G8264tor_2 switch.

Example 5-197 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 18
    lacp key 5
    !
interface port 20
    lacp key 5
    !
interface port 22
    lacp key 6
    !
interface port 24
    lacp key 6
    !
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

```

EX4500-VC switch

Example 5-198 lists the configuration of the EX4500-VC switch.

Example 5-198 EX4500-VC switch configuration

```


set version 12.1R3.5
set system host-name EX4500-VC
set system root-authentication encrypted-password
"$1$16b7xXmg$rEkslG/xYqFKDv4y2rHCt/"
set system services ftp
set system services ssh
set system services netconf ssh
set system syslog user * any emergency
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set chassis aggregated-devices ethernet device-count 5
set interfaces xe-0/0/0 description "Connection to 8264_1 Port 18"
set interfaces xe-0/0/0 ether-options 802.3ad ae0
set interfaces xe-0/0/1 description "Connection to 8264_1 Port 20"
set interfaces xe-0/0/1 ether-options 802.3ad ae0
set interfaces xe-0/0/2 description "Connection to 8264_2 Port 22"
set interfaces xe-0/0/2 ether-options 802.3ad ae1
set interfaces xe-0/0/3 description "Connection to 8264_2 Port 24"
set interfaces xe-0/0/3 ether-options 802.3ad ae1
set interfaces xe-1/0/0 description "Connection to 8264_2 Port 18"
set interfaces xe-1/0/0 ether-options 802.3ad ae3

```

```

set interfaces xe-1/0/1 description "Connection to 8264_2 Port 20"
set interfaces xe-1/0/1 ether-options 802.3ad ae3
set interfaces xe-1/0/2 description "Connection to 8264_1 Port 22"
set interfaces xe-1/0/2 ether-options 802.3ad ae2
set interfaces xe-1/0/3 description "Connection to 8264_1 Port 24"
set interfaces xe-1/0/3 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family inet address 10.10.20.1/30
set interfaces ae1 unit 0 family inet address 10.10.22.1/30
set interfaces ae2 unit 0 family inet address 10.10.21.1/30
set interfaces ae3 unit 0 family inet address 10.10.23.1/30
set interfaces lo0 unit 0 family inet address 10.10.11.241/32
set interfaces vme unit 0 family inet address 172.25.101.241/16
set routing-options static route 0.0.0.0/0 next-hop 172.25.1.1
set routing-options router-id 10.10.11.241
set routing-options forwarding-table export load-balance
set protocols ospf area 0.0.0.0 interface ae0.0
set protocols ospf area 0.0.0.0 interface ae1.0
set protocols ospf area 0.0.0.0 interface ae2.0
set protocols ospf area 0.0.0.0 interface ae3.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols igmp-snooping vlan all
set protocols lldp interface all
set protocols lldp-med interface all
set policy-options policy-statement load-balance then load-balance per-packet
set ethernet-switching-options storm-control interface all
set virtual-chassis member 0 mastership-priority 255

```



IBM PureFlex Systems Troubleshooting and Maintenance

In this chapter, we describe the troubleshooting and maintenance steps that are used on IBM PureFlex Systems switches, with emphasis on EN4093 switch.

This chapter includes the following topics:

- ▶ Troubleshooting
- ▶ Configuration management
- ▶ Firmware management
- ▶ Logging and reporting

6.1 Troubleshooting

In this section, we show the basic troubleshooting tools and techniques. We describe 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. We recommend that you download and use the EN4093 documentation, which is available from the IBM Flex System Fabric EN4093 10Gb Scalable Switch Information Center at this website:

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, as shown in Figure 6-1.




Figure 6-1 System status LEDs: (left to right) OK, Identify and Error

The system status LEDs 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 yellow Error LED is lit, it indicates a critical alert.
 - When both LEDs are off, this indicates the switch is off.
- Identify (blue):

You can use this LED to identify the location of switch in chassis.

Complete the following steps to use Chassis Management Module (CMM) web interface to change the state of this LED:

- i. Click **Chassis Management** → **I/O Modules** in CMM Web GUI.

The window that is shown in Figure 6-2 opens.

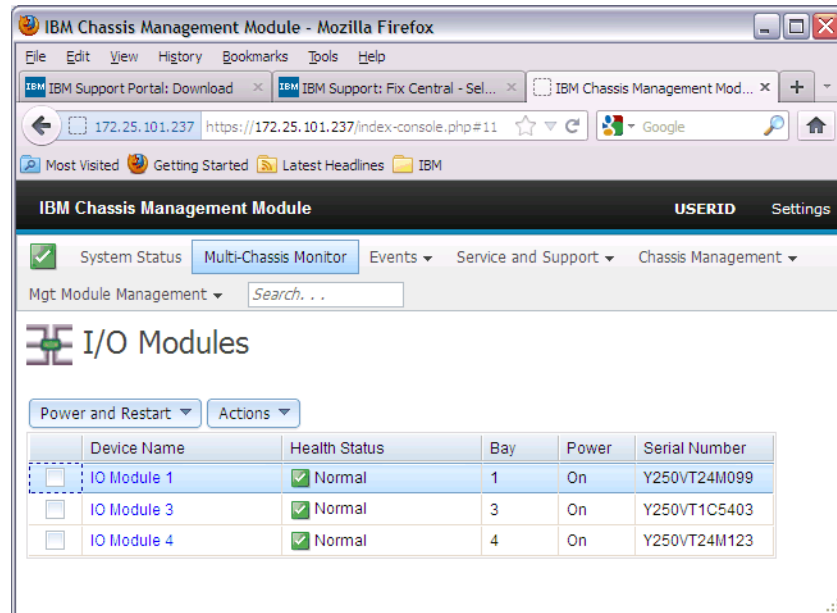


Figure 6-2 Select I/O module

- ii. Click the I/O module that you want to identify. In our case, we click **IO Module 1**. The window that is shown in Figure 6-3 opens.

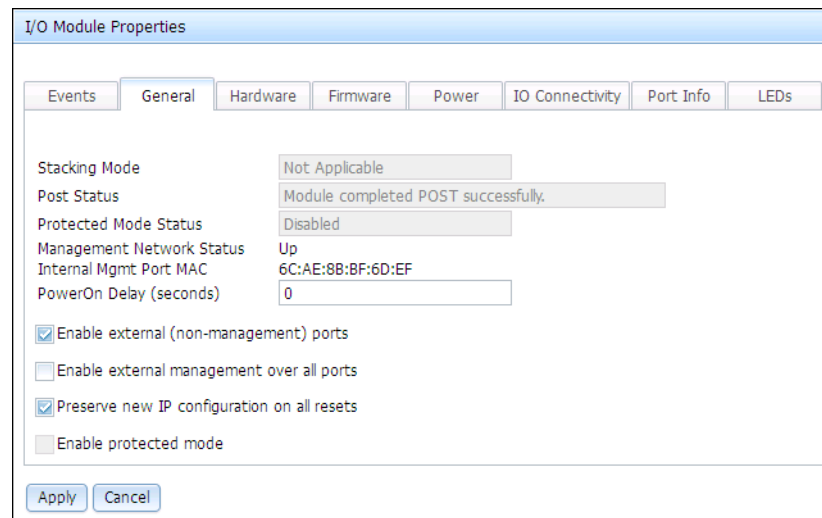


Figure 6-3 I/O module properties

- iii. Click the **LEDs** tab to open the window that is shown in Figure 6-4.

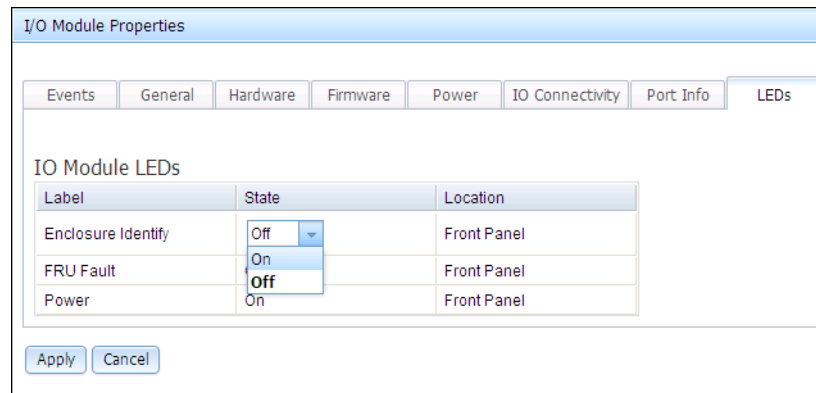


Figure 6-4 Toggle the Identify LED state

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, as shown in Figure 6-5 and Figure 6-6 on page 261.




Figure 6-5 SFP+ port LEDs




Figure 6-6 QSFP+ port LEDs

The Link and Tx/Rx LEDs indicate the following conditions:

- 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 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 perform a serial upgrade of the switch firmware, as described in 6.3.3, “Recovering from a failed firmware upgrade” on page 280.

6.1.2 Connectivity troubleshooting

In this section, we describe how to troubleshoot the IP connectivity in a network that is built on IBM System Networking switches. IBM switches include a set of simple tools that can be helpful for troubleshooting IP connectivity issues.

Ping

The **ping** command is a simple tool that is based on a request-response mechanism that is used 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 Internet Control Message Protocol (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 also is 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 verify connectivity to a remote network node and to track the responses from intermediate nodes. This task 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 likely to not be used on a remote node with a TTL of 1. After the packet reaches the intermediate router, the TTL is decremented, and the ICMP time-exceeded message is 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. We use OSPF in our network, which selects this path as the shortest one.

For test purposes, we shut down the direct link between the switch and target system and run the **traceroute** command again. The output is shown in Example 6-3.

Example 6-3 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  10.0.104.1      0 ms
 2  10.0.100.1      1 ms
Trace host responded.
```

Now we can see that to reach the destination, the switch uses 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 on page 264, selected traffic for ports 1 and 2 is monitored by port 3. In the example, 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.

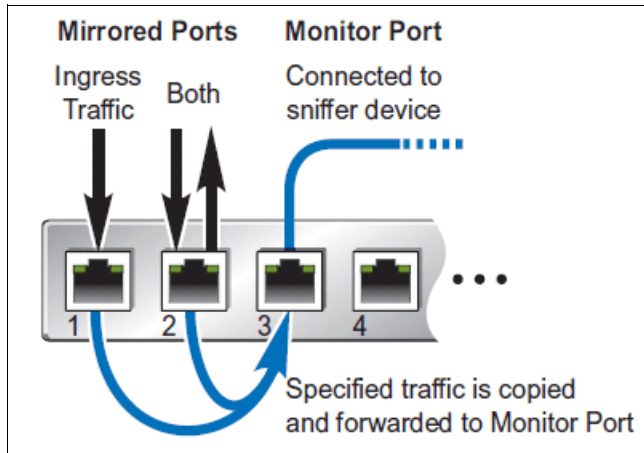


Figure 6-7 Mirroring ports

The composition of monitored packets in the EN4093 (based on the configuration of the ports) works in the following manner:

- Packets mirrored at port egress are mirrored before VLAN tag processing and might have a different PVID than packets that egress the port toward their actual network destination.
- Packets mirrored at port ingress are not modified.

In Example 6-4, we show 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 we show in Example 6-5, ingress and egress traffic on ports EXT1 - EXT4 are 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 CMM Web GUI, you can 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:
```

By 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 boot image using serial interface” on page 275.
- ▶ 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, as described in “Resetting with no terminal access to the switch” on page 273.
- ▶ Download new firmware to the switch. This option can be helpful if you must recover the switch after a failed firmware upgrade. We show an example of firmware recovery in 6.3.3, “Recovering from a failed firmware upgrade” on page 280.

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 the following files:

- ▶ `startup-config` is the configuration the switch uses when it is reloaded.
- ▶ `running-config` is the configuration that reflects all the changes you made from the command-line interface (CLI). It is stored in memory and is lost after the switch is reloaded.

6.2.2 Configuration blocks

The switch stores its configuration in one of the following configuration blocks:

- ▶ `active-config` is stored in the active configuration block.
- ▶ `backup-config` is stored in the backup configuration block.

When you save 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.

There also is a factory configuration block. This block holds the factory default configuration, with which you can restore the switch to factory defaults, if needed.

This setup has the flexibility you must manage the configuration of the switch and perform a possible 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 by using ISCLI

You can manage the configuration files by using the following 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 by 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 285.

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 tasks 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 also can use the SNMP MIBs in Table 6-1 to perform other functions, such as upgrading the switch firmware and saving the switch memory 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 also is 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. Start 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. Start the transfer. To save a running configuration file, use transfer action 4.

```
Set agTransferAction.0 "4"
```

Other tasks: Saving a switch memory dump

SNMP MIBs are not only useful to save and load switch configuration, you can also perform other tasks, such as saving a switch memory dump. To save a switch memory 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. Start 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 the switch for use in a different scenario is redeployed, or when you are troubleshooting a configuration issue. To reset the switch to factory defaults, you must perform one of the following procedures.

Resetting EN4093 to factory defaults via CMM

Complete the following steps to reset EN4093 to factory defaults via CMM:

1. Point your web browser to CMM IP address and log in, as shown in Figure 6-8.

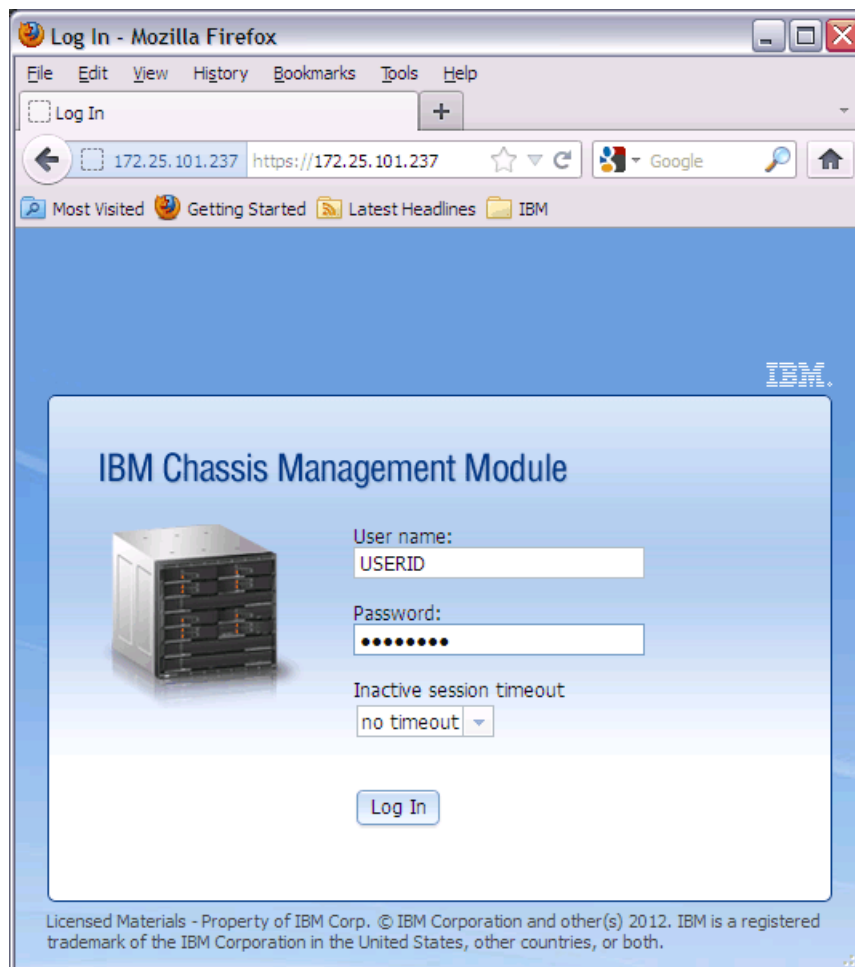


Figure 6-8 Log in to CMM

2. After successful login, CMM GUI displays, as shown in Figure 6-9.

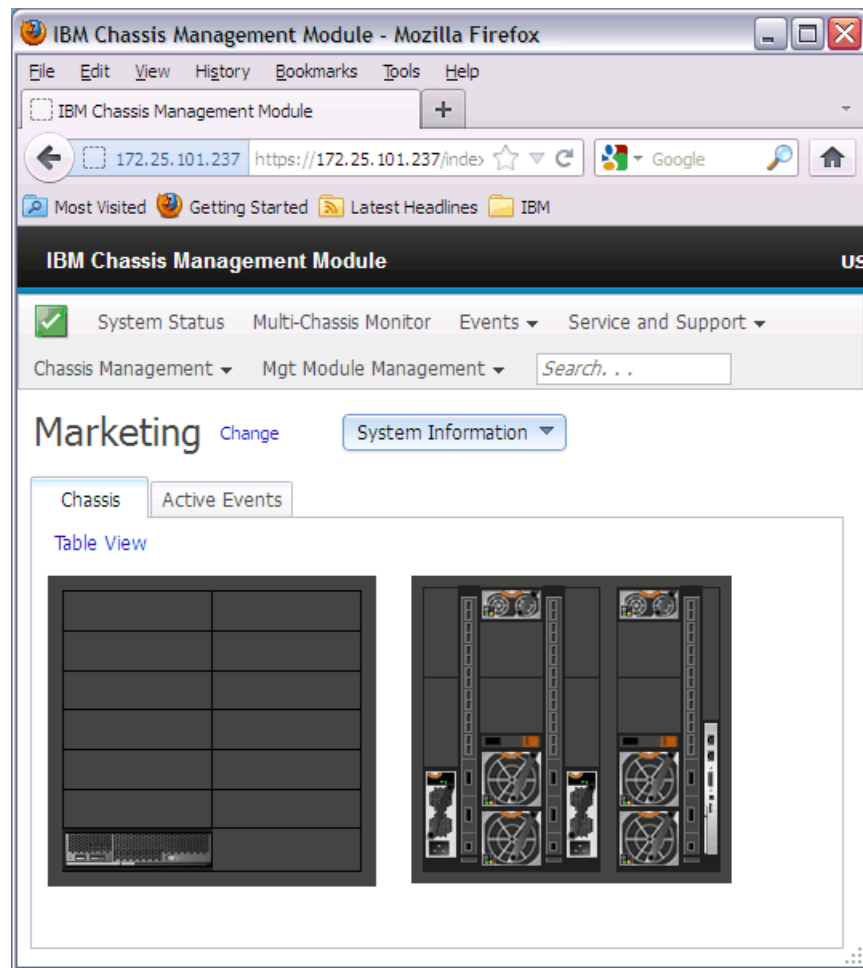


Figure 6-9 CMM GUI

3. Select **Chassis Management** → **I/O Modules**, as shown in Figure 6-10.




Figure 6-10 Select I/O Modules management

- As shown in Figure 6-11, select the I/O module that must be reset to factory defaults and click **Actions** → **Restore Factory Defaults**.

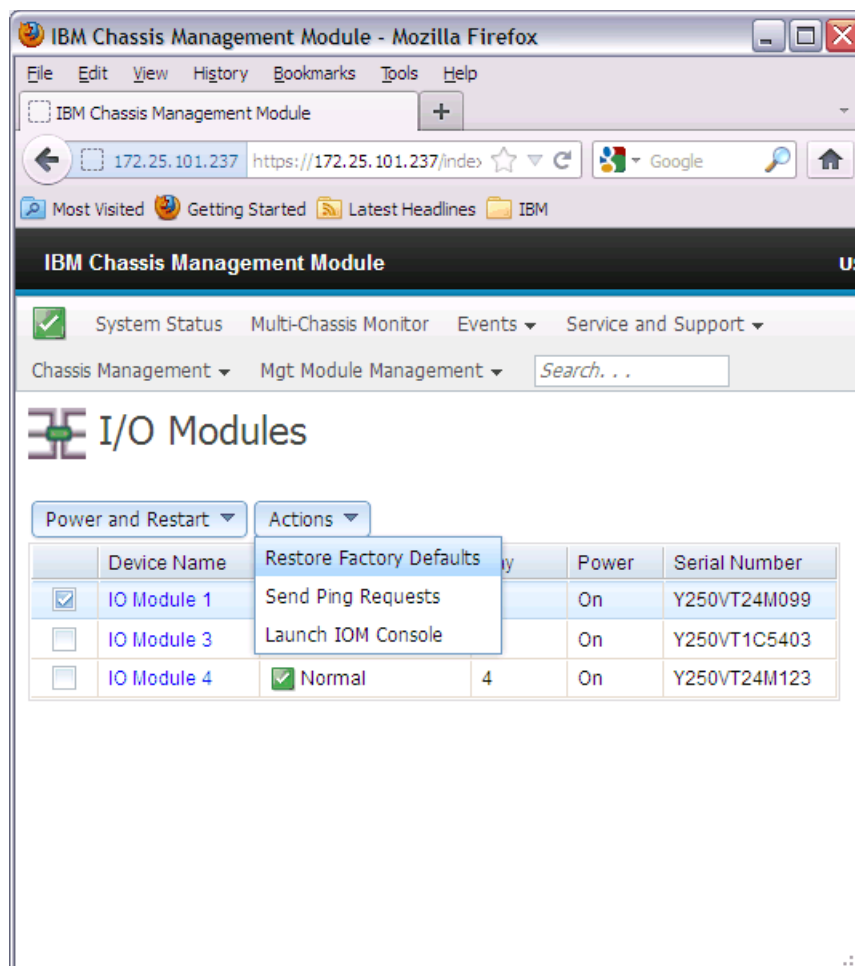


Figure 6-11 I/O Module 1 restore factory defaults

Resetting with terminal access to the switch

If you have terminal access to the switch, and you want to reset the switch to factory defaults, use the ISCLI command **boot configuration-block factory** and then reload the switch, as shown in Example 6-7).

Example 6-7 Reset to factory defaults using ISCLI

```
compass-2(config)#boot configuration-block factory
Next boot will use factory default config block instead of active.
```

```
compass-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

Complete the following steps to reset the switch to factory defaults and have no terminal access by using the serial console port:

1. Connect the management system to 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. You must restart the switch by powering it off and back on, or by restarting it in CMM Web interface.
3. You can 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 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, as shown in Example 6-9.

Example 6-9 Change 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 shown in Example 6-10, enter f to use the factory defaults configuration block.

Example 6-10 Use factory defaults configuration block

```
Enter configuration block: a, b or f (active, backup or factory): f
```

6. You see the initial menu again. Enter 6 to exit and reset the switch with the default configuration, as shown in Example 6-11.

Example 6-11 Exit from 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 recover a password, you must set the switch to the factory default by using one of the procedures described in 6.2.4, “Resetting to factory defaults” on page 269.

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 the **password** command in configuration mode, as shown in the following example:

```
Switch(config)#password
```

6.3 Firmware management

The switch firmware is the executable code that runs on the switch. The device comes preinstalled with certain firmware level. As new firmware versions are released, we recommend upgrading the code that runs on your switch. You can find the latest version of firmware supported for your switch at this website:

<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 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

```
compass-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, you can see that the system has the following OS images:

- ▶ image1: Version 7.2.2.2
- ▶ image2: Version 7.3.1

The boot image version is 7.3.1. But, the switch is set to boot from OS image1, which is at version 7.2.2.2. We want to make sure that the switch uses the same version for boot image and OS image. To boot from OS image2, run the **boot image image2** command, as shown in Example 6-13.

Example 6-13 Change to boot from image2

```
compass-2(config)#boot image image2
Next boot will use switch software image2 instead of image1.
```

Changing boot image using serial interface

Complete the following steps to use the serial connection and Boot Management Menu to change the boot image:

1. Connect serial cable to the switch serial management port and the management system, then start terminal emulation utility on the management system.
2. 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

3. When the system shows Memory Test, press Shift + B. The Boot Management Menu appears, 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
```

4. As shown in Example 6-14, we select the menu option 1 to change boot image from image1 to image2.

6.3.2 Upgrading the firmware with ISCLI

In this section, we show how to upgrade firmware of FlexSystem embedded switch EN4093. The latest firmware version at the time of this writing is 7.3.1.0. This code level is available on IBM Fix Central at this website:

<http://www.ibm.com/support/entry/portal/docdisplay?lnodocid=migr-5090394>

Complete the following steps to upgrade the firmware with ISCLI:

1. We must download the code update package (from IBM Fix Central or from <http://www.ibm.com/support/entry/portal/docdisplay?lnodocid=migr-5090394>) and unpack it. The update package contains the following 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, we renamed these files to the following names:
 - Boot image file *7310boot.img*
 - OS image file *7310os.img*
2. We put the two files onto an FTP or SFTP server. In our example, we use the CMM built-in TFTP server. Figure 6-12 on page 277 shows the two files on CMM TFTP server.




Figure 6-12 Firmware v7.3.1.0 image files on CMM TFTP server

3. We are now ready to download the image files to EN4093. First, we must log in to EN4093 as administrator, as shown in Example 6-15. When prompted to select CLI mode, we choose **iscli**.

Example 6-15 Log in to EN4093

```
log in as: admin
Using keyboard-interactive authentication.
Enter password:
```

IBM Flex System Fabric EN4093 10Gb Scalable Switch.

```
Select Command Line Interface mode (ibmnos-cli/iscli): iscli
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. We enable privileged EXEC mode (by using the **enable** command) and download the boot image file. As shown in Example 6-16, we use the **copy tftp boot-image** command to download the boot image file.

Example 6-16 Enable privileged EXEC mode and download boot image

```
compass-2>enable
```

```
Enable privilege granted.
compass-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. As shown in Example 6-17, we download the OS image file into image2 and set switch to boot from image2 by using the **copy tftp image2** command.

Example 6-17 Download OS image file

```

compass-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 compass-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. We must reboot the switch to activate the new code, as shown in Example 6-18.

Example 6-18 Reboot the switch

```
compass-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, you can use the **show boot** command 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

```
compass-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
```

The EN4093 firmware is now upgraded.

6.3.3 Recovering from a failed firmware upgrade

While it is extremely unlikely, the firmware upgrade process can fail. If this situation occurs, you can still recover the EN4093 switch. Connect a PC that is running a terminal emulation utility to the serial port of your switch while the switch is off, and 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 also might be useful when you boot the switch and the boot and OS versions are not equal.

Next, 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 on page 281.

Example 6-20

```
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:
```

Select 4 for Xmodem download of boot image. Change the serial connection speed as shown in the following example:

Switch baudrate to 115200 bps and press ENTER ...

Change the settings of your terminal to meet the 115200 bps requirement and press Enter. The system switches to download accept mode. You see a series of C characters on the screen that prompt you when the switch is ready. Start an Xmodem terminal to push the boot code you want to restore into the switch. Select the boot code for your system, and the switch starts the download. You should see a screen similar to what is shown in 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

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. The Boot Management Menu opens again. Select option 3 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 the speed is changed to 115200 bps, press Enter to continue the download. Select the OS image 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 screen similar to the one that is shown 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] :
```

You are prompted to select the image space in the switch that you want to upgrade. After selecting the OS image bank, you see a screen similar to the one that is shown 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
```

When this process is done, you are prompted to reconfigure your terminal to 9600 bps speed again, as shown in the following example:

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 includes 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, which is included in text form as a prefix to the log message. One of the following eight different prefixes is used, depending on the condition of which the administrator is being notified:

- ▶ 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 of the following information should be logged by syslog:

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 format:

```
[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.

In Example 6-25 on page 285, we set a configuration 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 also can use the **logging host** command to specify the interface used for logging. The following options are available:

- ▶ data-port
- ▶ extm-port
- ▶ mgt-port

For example, to send the logs to a second destination from a data port, run the command that is 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 the **logging console** command. You can select the severity level of messages to be logged by using 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. 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, it is a best practice that you disable these functions before the switch is connected to the network.

SNMP versions 1 and 2

To access the SNMP agent on the EN4093, the read and write community strings on the SNMP manager should be configured 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 should 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 that are sent out by the SNMP agent on the switch, configure the trap host on the switch by running the following commands:

```
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 SNMP, which 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.

By 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 is adminmd5). The authentication that is used is MD5.
- ▶ User 2 name is adminsha (password is adminsha). The authentication that is 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 that are 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 with the views the group may access by running the commands that are 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 that are 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

In this section we describe 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 choose to send the traps with 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

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, which is available at this website:

<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 the collection of Ethernet statistics as outlined in the RMON statistics MIB, which refers 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.

Example configuration

Complete the following steps to set up the example configuration:

1. Enable RMON on a port. To enable RMON on a port, run the following **interface** and **rmon** commands:
 - en4093flex_1(config)# interface port 1
 - en4093flex_1(config-if)# rmon
2. To view the RMON statistics, run the **interface** command, run the **rmon** command, and then run the **show** command to show the interface, as shown in Example 6-36.

Example 6-36 View of 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
etherStatsPkts1024to1518Octets: 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 that is gathered during discrete 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) that are requested by the user for each History Group. Granted buckets are the number of buckets that are granted by the system and 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.

History MIB Object ID

The type of data that can be sampled must be of an Index object type, as described in RFC 1213 and RFC 1573, which are available at the following websites:

- ▶ <http://www.ietf.org/rfc/rfc1213.txt>
- ▶ <http://www.ietf.org/rfc/rfc1573.txt>

The most common data type for the History sample is shown in the following example:

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 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, the switch reports general port statistics and port Ethernet statistics at the end of each configured polling interval.

sFlow network sampling

In addition to statistical counters, IBM System Networking switches can be configured to collect periodic samples of the traffic data that is 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 to disable (the default) 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 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 performed before ACLs are processed. For ports configured with 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 usage to reach maximum limits. Use larger rate values for ports that experience heavy traffic.



Fully Redundant Implementation with IBM vLAG and Juniper MC-LAG Technology

This implementation scenario incorporates switch virtualization features that allow a downstream switch the ability to be connected to two upstream, virtualized switches through the means of aggregated links. Inter-Switch Links (ISLs) between the same or similar products on the aggregation or access-layer provide a loop-free design that is 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 and internally prevent implicit loops. This design is recommended for customers that want to incorporate a best practice implementation on a Juniper network that uses next generation networking features such as Juniper Virtual Chassis (VC) or Multi-chassis Link Aggregation Groups (MC-LAG) technology.

This approach has the following advantages:

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

This appendix includes the following topics:

- ▶ Components used
- ▶ Network topology and physical setup
- ▶ Switches configuration
- ▶ Verification and show command output
- ▶ Full configuration files

A.1 Components used

Two of each of the following components are used in this scenario:

- ▶ Juniper QFX3500
- ▶ IBM G8264 RackSwitch
- ▶ IBM Flex System Fabric EN4093 10Gb Scalable Switch

A.2 Network topology and physical setup

Figure A-1 on page 295 shows the network topology for the fully redundant scenario with virtualization technology, Juniper MC-LAG and IBM vLAG.

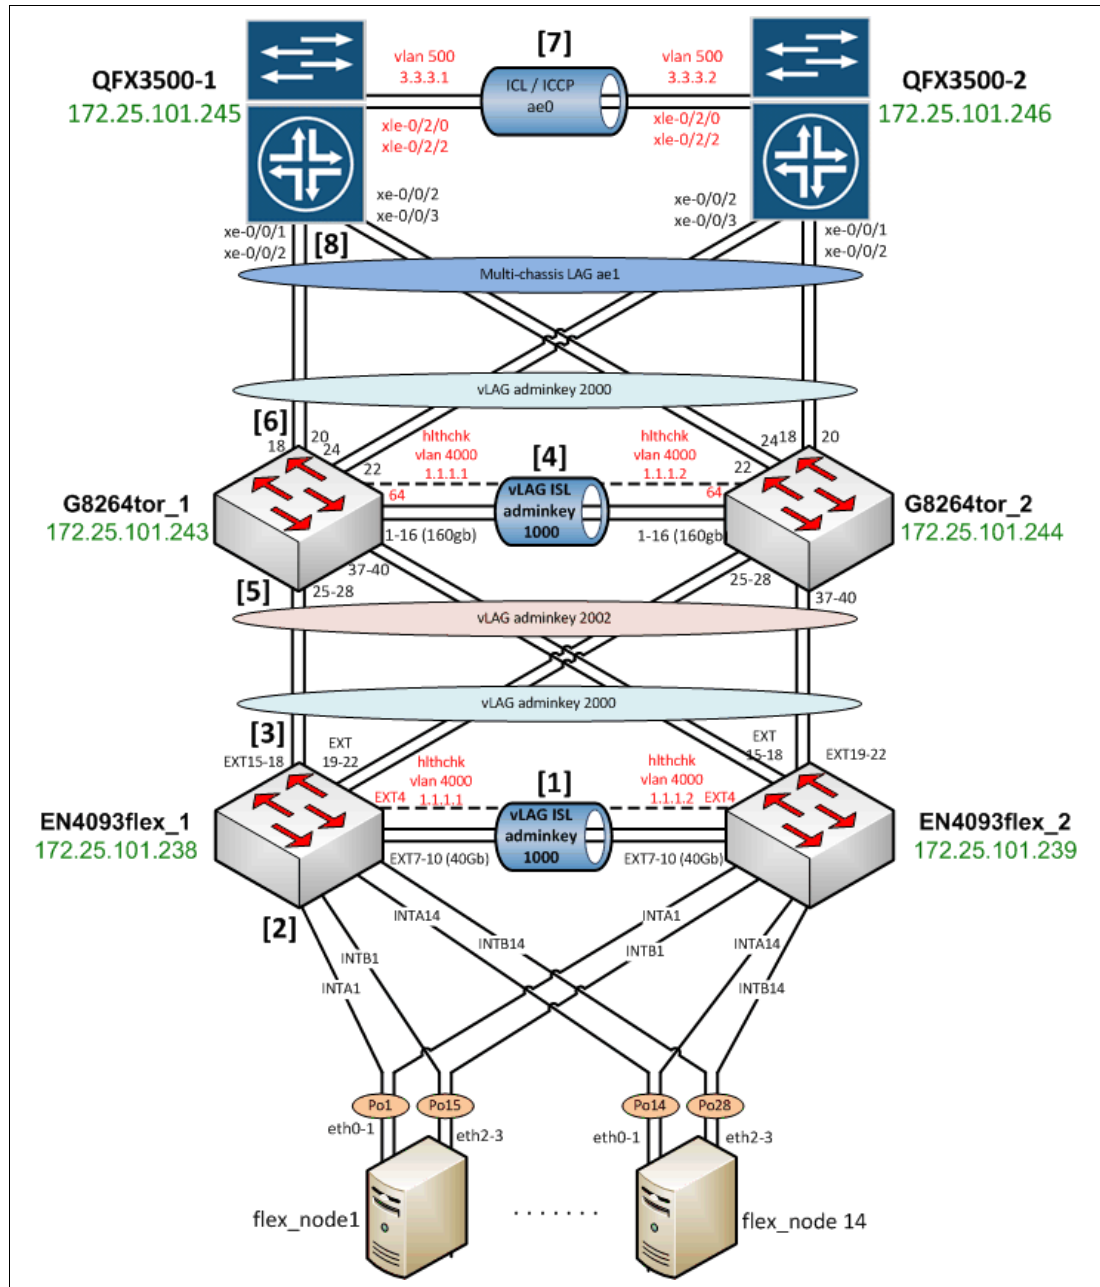


Figure A-1 Network Topology diagram for fully redundant scenario

Start by verifying the physical cabling between the EN4093 switches and G8264s. In our lab environment, we used four IBM QSFP+ DAC Break Out Cables from the EN4093 switches to the upstream G8264s. This requires that the EN4093 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 160Gb ISL between the G8264 switches and 10Gb SFP+ DAC cables were used for all other connections in the diagram.

A.3 Switches configuration

As we did in the other use cases, we begin the implementation of this scenario on the IBM Flex System Fabric EN4093 switches, working our way northward on the diagram in Figure A-1 on page 295. Each step provides the commands that are necessary and are reflective of the numbering schema in the diagram in Figure A-1 on page 295 to aid the user in what is configured.

A.3.1 IBM EN4093flex and G8264tor switches configuration

From a configuration standpoint, the configuration of the EN4093 flex and G8264tor switches (steps 1 - 6 in “Juniper QFX3500-1 switch configuration” on page 296) is exactly the same as the configuration that is described in 5.3, “Fully redundant with Virtualized Chassis technology” on page 74.

For more information about configuration details, see 5.3.3, “EN4093 flex_1 configuration” on page 76 and 5.3.4, “G8264tor_1 configuration” on page 81.

A.3.2 Juniper QFX3500-1 switch configuration

Here we configure the Juniper QFX3500-1 core switch.

General configuration

Complete the following steps to set up the general configuration:

1. Assign IP addresses for the management Port and configure an IP gateway for our 172 management network, as shown in Example A-1.

Example: A-1 Management IP address and IP gateway configuration on QFX3500-1

```
set interfaces me4 unit 0 family inet address 172.25.101.245/16
set routing-options static route 0.0.0.0/0 next-hop 172.25.1.1
```

2. Create the DATA (4092) VLAN and create the Routed VLAN Interfaces (RVI) for our Data VLAN, as shown in Example A-2.

Example: A-2 Data VLAN and associated RVI configuration on QFX3500-1

```
set vlans DATA vlan-id 4092
set vlans DATA l3-interface vlan.4092
set interfaces vlan unit 4092 family inet address 10.1.4.245/24
```

Configuring the Interchassis Link on Juniper switches

Complete the following steps to configure the Interchassis Link (ICL) on Juniper switches:

1. Configure VLAN 500 and the associated the Routed VLAN Interface that is used by the Internet Chassis Control Protocol (ICCP), as shown in Example A-3 on page 297. Then, configure the logical and physical interface that is used in the Interchassis Link.

Example: A-3 ICL VLAN and interface configuration on QFX3500-1

```
set vlans VLAN500 vlan-id 500
set vlans VLAN500 l3-interface vlan.500
set interfaces vlan unit 500 family inet address 3.3.3.1/30

set chassis aggregated-devices ethernet device-count 1

set interfaces xle-0/2/0 ether-options 802.3ad ae0
set interfaces xle-0/2/2 ether-options 802.3ad ae0
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members VLAN500
```

2. For MC-LAG, you must configure ICCP to exchange information between two MC-LAG peers. Example A-4 shows how to enable ICCP in the QFX3500-1 switch. The IP 3.3.3.1 acts as the source address in the QFX3500-1 switch and the IP 3.3.3.2 is the address of the QFX3500-2 peer switch. Then, configure the ae0 interface as the multi-chassis protection link between the QFX switches.

Example: A-4 ICCP configuration on QFX3500-1

```
set protocols iccp local-ip-addr 3.3.3.1
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip
172.25.101.246
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval
minimum-interval 60

set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

Configuring downstream G8264tor facing ports

Configure the downstream logical and physical interfaces to all be on the same aggregated interface by using LACP aggregation (as shown in Example A-5) and add them to the MC-LAG.

Example: A-5 Downstream G8264tor facing interfaces on QFX3500-1

```
set chassis aggregated-devices ethernet device-count 2

set interfaces xe-0/0/0 ether-options 802.3ad ae1
set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/2 ether-options 802.3ad ae1
set interfaces xe-0/0/3 ether-options 802.3ad ae1

set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members DATA
```

A.3.3 Juniper QFX3500-2 switch configuration

Here we configure the Juniper QFX3500-2 core switch.

General configuration

Complete the following steps to set up the general configuration:

1. Assign IP addresses for the management Port and configure an IP gateway for our 172 management network, as shown in Example A-6.

Example: A-6 Management IP address and IP gateway configuration on QFX3500-2

```
set interfaces me4 unit 0 family inet address 172.25.101.246/16
set routing-options static route 0.0.0.0/0 next-hop 172.25.1.1
```

2. Create the DATA (4092) VLAN and create the Routed VLAN Interfaces (RVI) for our Data VLAN, as shown in Example A-7.

Example: A-7 Data VLAN and associated RVI configuration on QFX3500-2

```
set vlans DATA vlan-id 4092
set vlans DATA l3-interface vlan.4092
set interfaces vlan unit 4092 family inet address 10.1.4.246/24
```

Configuring the ICL on Juniper switches

Complete the following steps to set up the ICL on Juniper switches:

1. Configure VLAN 500 and the associated the Routed VLAN Interface that is used by the ICCP, as shown in Example A-8. Then, configure the logical and physical interface use in the Interchassis Link.

Example: A-8 ICL VLAN and interface configuration on QFX3500-2

```
set vlans VLAN500 vlan-id 500
set vlans VLAN500 l3-interface vlan.500
set interfaces vlan unit 500 family inet address 3.3.3.2/30

set chassis aggregated-devices ethernet device-count 1

set interfaces xle-0/2/0 ether-options 802.3ad ae0
set interfaces xle-0/2/2 ether-options 802.3ad ae0
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members VLAN500
```

2. For MC-LAG, you must configure ICCP to exchange information between two MC-LAG peers. Example A-9 on page 299 shows how to enable ICCP in the QFX3500-2 switch. The IP 3.3.3.2 acts as the source address in the QFX3500-2 switch and the IP 3.3.3.1 is the address of the QFX3500-1 peer switch. Then configure the ae0 interface as the multi-chassis protection link between the QFX switches.

Example: A-9 ICCP configuration on QFX3500-2

```
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip
172.25.101.245
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 60
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval
minimum-interval 60

set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
```

Configuring downstream G8264tor facing ports

Configure the downstream logical and physical interfaces to all be on the same aggregated interface by using LACP aggregation (as shown in Example A-10) and add them to the MC-LAG.

Example: A-10 Downstream G8264tor facing interfaces on QFX3500-2

```
set chassis aggregated-devices ethernet device-count 2

set interfaces xe-0/0/0 ether-options 802.3ad ae1
set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/2 ether-options 802.3ad ae1
set interfaces xe-0/0/3 ether-options 802.3ad ae1

set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members DATA
```

A.4 Verification and show command output

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

A.4.1 QFX3500 output

Here we list output from the switch with host name QFX3500-1. Similar or identical output exists for the switch with host name QFX3500-2.

show system software

The output in Example A-11 shows information regarding the switch that we used and the associated code and firmware level at that time.

Example: A-11 Juniper QFX3500-1 show system software output

Information for jbase:

Comment:

Junos OS Base OS Software Suite [12.2X50-D20.4]

Information for jcrypto:

Comment:

Junos OS Crypto Software Suite [12.2X50-D20.4]

Information for jdocs:

Comment:

Junos OS Online Documentation [12.2X50-D20.4]

Information for jkernel:

Comment:

Junos OS Kernel Software Suite [12.2X50-D20.4]

Information for jpfe:

Comment:

Junos OS Packet Forwarding Engine Support (QFX) [12.2X50-D20.4]

Information for jroute:

Comment:

Junos OS Routing Software Suite [12.2X50-D20.4]

Information for jswitch:

Comment:

Junos OS Enterprise Software Suite [12.2X50-D20.4]

Information for junos:

Comment:

Junos OS Base OS boot [12.2X50-D20.4]

Information for jweb:

Comment:

Junos OS Web Management [12.2X50-D20.4]

show vlan

Example A-12 shows the VLAN assignments for all the various ports on the switch.

Example: A-12 QFX3500-1 show VLAN output

Name	Tag	Interfaces
DATA	4092	ae0.0*, ae1.0*
VLAN500	500	ae0.0*
default		None
native-vlan	1	ae0.0*, ae1.0*

show interfaces terse

Example A-13 shows the full interface table, listing port status, link status, IP address, and so on.

Example: A-13 QFX-1 show interfaces terse output

Interface	Admin	Link	Proto	Local	Remote
xe-0/0/0	up	up			
xe-0/0/0.0	up	up	aenet	--> ae1.0	
xe-0/0/1	up	up			
xe-0/0/1.0	up	up	aenet	--> ae1.0	
xe-0/0/2	up	up			
xe-0/0/2.0	up	up	aenet	--> ae1.0	
xe-0/0/3	up	up			
xe-0/0/3.0	up	up	aenet	--> ae1.0	
xe-0/0/42	up	down			
xe-0/0/43	up	down			
xe-0/0/44	up	down			
xe-0/0/45	up	down			
xe-0/0/46	up	down			
xe-0/0/47	up	down			
x1e-0/2/0	up	up			
x1e-0/2/0.0	up	up	aenet	--> ae0.0	
x1e-0/2/2	up	up			
x1e-0/2/2.0	up	up	aenet	--> ae0.0	
ae0	up	up			
ae0.0	up	up	eth-switch		
ae1	up	up			
ae1.0	up	up	eth-switch		
ae2	up	down			
ae3	up	down			
ae4	up	down			
bme0	up	up			
bme0.0	up	up	inet	128.0.0.1/2 128.0.0.16/2 128.0.32.0/2	
			tnp	0x10	
dcfabric	up	up			
dcfabric.0	up	up	inet mpls eth-switch		

```

dsc                up    up
gre                up    up
pip               up    up
lo0               up    up
lo0.16384         up    up    inet    127.0.0.1        --> 0/0
lo0.16385         up    up    inet
lsi               up    up
me0               up    down
me1               up    down
me4               up    up
me4.0             up    up    inet    172.25.101.245/16
mtun              up    up
pimd              up    up
pime              up    up
tap               up    up
vlan              up    up
vlan.500          up    up    inet    3.3.3.1/30
vlan.4092         up    up    inet    10.1.4.245/24

```

show lldp neighbors

Example A-14 shows the LLDP information and serves as a means to verify our physical connectivity.

Example: A-14 QFX3500-1 show LLDP neighbors output

Local Interface	Parent Interface	Chassis Id	Port info	System Name
xe-0/0/2.0	ae1.0	08:17:f4:33:75:00	Ethernet22	G8264TOR-2
xe-0/0/3.0	ae1.0	08:17:f4:33:75:00	Ethernet24	G8264TOR-2
xe-0/0/0.0	ae1.0	08:17:f4:33:9d:00	Ethernet18	G8264TOR-1
xe-0/0/1.0	ae1.0	08:17:f4:33:9d:00	Ethernet20	G8264TOR-1
xle-0/2/0.0	ae0.0	f8:c0:01:3b:c4:00	xle-0/2/0.0	QFX3500-2
xle-0/2/2.0	ae0.0	f8:c0:01:3b:c4:00	xle-0/2/2.0	QFX3500-2

show iccp

The output in Example A-15 shows the status of the ICCP protocol, which is the protocol that synchronizes state, ARP, and the redundancy group between the two independent chassis.

Example: A-15 QFX3500-1 show iccp output

```

Redundancy Group Information for peer 3.3.3.2
  TCP Connection      : Established
  Liveliness Detection : Up

```

```
Backup liveness peer status: Up
```

```
Client Application: MCSNOOPD
```

```
Client Application: lacpd
```

```
Client Application: eswd
```

show interfaces mc-ae

Example A-16 shows if the ICCP protocol is working at the MC-LAG level.

Example: A-16 Juniper QFX3500-1 show interfaces MC-AE output

```
Member Link          : ae1
Current State Machine's State: mcae active state
Local Status         : active
Local State          : up
Peer Status          : active
Peer State           : up
  Logical Interface   : ae1.0
  Topology Type       : bridge
  Local State         : up
  Peer State          : up
  Peer Ip/MCP/State   : 3.3.3.1 ae0.0 up
```

show lacp interfaces

The command output in Example A-17 shows which ports are participating in an LACP aggregation and in which aggregated interface.

Example: A-17 QFX3500-1 show lacp interfaces output

Aggregated interface: ae0

Aggregated interface: ae1

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/0/2	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
xe-0/0/2	Partner	No	No	Yes	Yes	Yes	Yes	Slow	Active
xe-0/0/3	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
xe-0/0/3	Partner	No	No	Yes	Yes	Yes	Yes	Slow	Active
xe-0/0/0	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
xe-0/0/0	Partner	No	No	Yes	Yes	Yes	Yes	Slow	Active
xe-0/0/1	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
xe-0/0/1	Partner	No	No	Yes	Yes	Yes	Yes	Slow	Active

LACP protocol:	Receive State	Transmit State	Mux State
xe-0/0/2	Current	Slow periodic	Collecting distributing
xe-0/0/3	Current	Slow periodic	Collecting distributing
xe-0/0/0	Current	Slow periodic	Collecting distributing
xe-0/0/1	Current	Slow periodic	Collecting distributing

Ping output for equipment on VLAN 4092

To verify connectivity, we issued several ping commands to devices in the lab infrastructure on VLAN 4092 (Data VLAN), as shown in Example A-18.

Example: A-18 Ping verification for equipment on VLAN 4092

```
root@QFX3500-1> 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=64 time=4.022 ms
64 bytes from 10.1.4.10: icmp_seq=1 ttl=64 time=1.095 ms
64 bytes from 10.1.4.10: icmp_seq=2 ttl=64 time=1.218 ms
64 bytes from 10.1.4.10: icmp_seq=3 ttl=64 time=1.243 ms
64 bytes from 10.1.4.10: icmp_seq=4 ttl=64 time=1.095 ms
^C
```

```

--- 10.1.4.10 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.095/1.735/4.022/1.145 ms

root@QFX3500-1> 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=1 ttl=64 time=1.095 ms
64 bytes from 10.1.4.238: icmp_seq=2 ttl=64 time=1.218 ms
64 bytes from 10.1.4.238: icmp_seq=3 ttl=64 time=1.243 ms
64 bytes from 10.1.4.238: icmp_seq=4 ttl=64 time=1.095 ms
^C
--- 10.1.4.238 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.095/1.162/1.243/1.105 ms

root@QFX3500-1> 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=64 time=2.343 ms
64 bytes from 10.1.4.243: icmp_seq=1 ttl=64 time=1.082 ms
64 bytes from 10.1.4.243: icmp_seq=2 ttl=64 time=1.227 ms
64 bytes from 10.1.4.243: icmp_seq=3 ttl=64 time=1.240 ms
64 bytes from 10.1.4.243: icmp_seq=4 ttl=64 time=1.095 ms
^C
--- 10.1.4.243 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.082/1.397/2.343/1.122 ms

```

A.4.2 Full configuration files

In this section, we describe the full configuration of the Juniper QFX3500 switches. The full configuration of the IBM System Networking switches (EN4093 and G8264) are exactly the same as the configuration that is shown in 5.3, “Fully redundant with Virtualized Chassis technology” on page 74.

QFX3500-1

Example A-19 lists the full configuration for the Juniper QFX3500-1 switch.

Example: A-19 Juniper QFX3500-1 configuration file

```

set version 12.2X50.4
set system host-name QFX3500-1
set system root-authentication encrypted-password
"$1$aA9cqdbF$Ean2XEiJTW3VTJY7k/zWb/"
set system services ftp
set system services ssh
set system syslog user * any emergency
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set chassis aggregated-devices ethernet device-count 5
set chassis fpc 0 pic 2 xle port-range 0 3
set interfaces xe-0/0/0 ether-options 802.3ad ae1
set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/2 ether-options 802.3ad ae1

```

```

set interfaces xe-0/0/3 ether-options 802.3ad ae1
set interfaces xle-0/2/0 ether-options 802.3ad ae0
set interfaces xle-0/2/2 ether-options 802.3ad ae0
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members VLAN500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-aether mc-aether-id 3
set interfaces ae1 aggregated-ether-options mc-aether chassis-id 0
set interfaces ae1 aggregated-ether-options mc-aether mode active-active
set interfaces ae1 aggregated-ether-options mc-aether status-control active
set interfaces ae1 aggregated-ether-options mc-aether init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members DATA
set interfaces me4 unit 0 family inet address 172.25.101.245/16
set interfaces vlan unit 500 family inet address 3.3.3.1/30
set interfaces vlan unit 4092 family inet address 10.1.4.245/24
set routing-options static route 0.0.0.0/0 next-hop 172.25.1.1
set protocols iccp local-ip-addr 3.3.3.1
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip
172.25.101.246
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval
minimum-interval 60
set protocols igmp-snooping vlan all
set protocols dcbx interface all
deactivate protocols rstp
set protocols lldp interface all
set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
set vlans DATA vlan-id 4092
set vlans DATA l3-interface vlan.4092
set vlans VLAN500 vlan-id 500
set vlans VLAN500 l3-interface vlan.500
set ethernet-switching-options storm-control interface all

```

QFX3500-2

Example A-20 lists the full configuration for the Juniper QFX3500-2 switch.

Example: A-20 Juniper QFX3500-2 configuration file

```

set version 12.2X50.4
set system host-name QFX3500-2
set system root-authentication encrypted-password
"$1$Tf3jjLSa$GfUi4Hcu5SH194SVSyUj0/"
set system services ftp
set system services ssh
set system syslog user * any emergency
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set chassis aggregated-devices ethernet device-count 5
set chassis fpc 0 pic 2 xle port-range 0 3
set interfaces xe-0/0/0 ether-options 802.3ad ae1


```

```

set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/2 ether-options 802.3ad ae1
set interfaces xe-0/0/3 ether-options 802.3ad ae1
set interfaces xle-0/2/0 ether-options 802.3ad ae0
set interfaces xle-0/2/2 ether-options 802.3ad ae0
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members VLAN500
set interfaces ae1 aggregated-ether-options lACP active
set interfaces ae1 aggregated-ether-options lACP system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lACP admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members DATA
set interfaces me4 unit 0 family inet address 172.25.101.246/16
set interfaces vlan unit 500 family inet address 3.3.3.2/30
set interfaces vlan unit 4092 family inet address 10.1.4.246/24
set routing-options static route 0.0.0.0/0 next-hop 172.25.1.1
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip
172.25.101.245
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 60
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval
minimum-interval 60
set protocols igmp-snooping vlan all
set protocols dcbx interface all
deactivate protocols rstp
set protocols lldp interface all
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
set vlans DATA vlan-id 4092
set vlans DATA l3-interface vlan.4092
set vlans VLAN500 vlan-id 500
set vlans VLAN500 l3-interface vlan.500
set ethernet-switching-options storm-control interface all

```

This setup is now complete.



Junos Operating System to IBM Networking Operating System Command Comparison

This appendix shows a command comparison between Junos operating system (OS) and the IBM Networking OS on certain selected features. Although not an exhaustive list of available commands on either platform, common commands are shown here to aid the user in quickly implementing IBM System Networking products.

IBM switches offer two different command-line interfaces (CLIs): IBM Networking OS-CLI, and IBM Industry standard CLI (isCLI). The isCLI is designed to be familiar to network professionals who are accustomed to other vendors' CLI. As of IBM System Networking version 7.6.1.0, the configuration command syntax is updated to match industry standards even closer when ports and trunking, virtual local area networks (VLANs), and Spanning Tree Protocol (STP) are configured. The existing (pre-7.6.1.0) and the updated command syntax are recognized in version 7.6.1.0 and above.

This appendix includes the following topics:

- ▶ Host name and banner
- ▶ Management
- ▶ SSH and Telnet
- ▶ Local authentication
- ▶ Syslog
- ▶ Network Time Protocol
- ▶ Simple Network Management Protocol
- ▶ Link Layer Discovery Protocol
- ▶ Spanning Tree Protocol
- ▶ Interface speed and duplex
- ▶ VLAN
- ▶ VLAN Tagging (802.1q)
- ▶ Trunking and link aggregation: Static
- ▶ Trunking and link aggregation: Link Aggregation Control Protocol
- ▶ External authentication

- ▶ Bridge Protocol Data Unit Guard
- ▶ Dynamic Host Configuration Protocol snooping
- ▶ Port mirroring
- ▶ Open Shortest Path First configuration

B.1 Host name and banner

This section includes information about the host name and banner.

Juniper Junos OS

To change the host name and create a login message or banner, use the commands that are shown in Example B-1.

Example: B-1 Junos OS host name and login message configuration

```
admin@EX4500# set system host-name EX4500-VC
admin@EX4500-VC# set system login message "IBM's internal use only. Use is subject
to audit at any time by IBM management."
```

IBM Networking OS

Configure the host name and a login message on IBM isCLI as shown in Example B-2.

Example: B-2 IBM Networking OS host name and login message configuration

```
switch(config)#hostname "ToR-G8264"
ToR-G8264(config)#system notice addline "Login Banner Message"
```

B.2 Management

This section includes information about configuring a dedicated network management port IP address and gateway.

B.2.1 Juniper Junos OS

Example B-3 shows how to configure a dedicated network management port IP address and gateway.

Example: B-3 Management IP address and gateway

```
admin@EX4500-VC# set interfaces me0 unit 0 family inet address <ip
address>/<prefix>

admin@EX4500-VC# set routing-options static route 0.0.0.0/0 next-hop <ip address>
```

B.2.2 IBM Networking OS

In IBM Networking OS interface ip 128 is used as the dedicated network management port. Example B-4 shows how to configure the management port IP address and gateway.

Example: B-4 Management IP address and gateway

```
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
```

B.3 SSH and Telnet

This section includes information about Secure Shell (SSH) and Telnet.

B.3.1 Juniper Junos OS

SSHv1 or SSHv2 and Telnet can be configured in Junos OS, as shown in Example B-5.

Example: B-5 SSH and Telnet configuration

```
admin@EX4500-VC# set system services ssh protocol-version v2
admin@EX4500-VC# set system services telnet
```

B.3.2 IBM Networking OS

SSH is disabled by default. After the SSH server on the device is enabled, all applicable keys are generated automatically for the user. Telnet is enabled by default.

Example B-6 shows how to enable SSH and disable Telnet.

Example: B-6 Enabling SSH and disabling Telnet

```
switch(config)#ssh enable
switch(config)#no access telnet enable
```

B.4 Local authentication

This section includes information about local authentication.

B.4.1 Juniper Junos OS

There is one user name that is defined by default and it is the root user. Example B-7 shows how to set the root password and create a new user.

Example: B-7 Root password and create a new user account

```
admin@EX4500-VC# set system root-authentication plain-text-password
New password:<password>
Retype new password:<password>

admin@EX4500-VC# set system login user admin class super-user authentication
plain-text-password
New password:<password>
Retype new password:<password>
```

The root account cannot be disabled and is always enabled.

B.4.2 IBM Networking OS

There are three user names that are defined on the system as factory defaults, as shown in Table B-1.

Table B-1 IBM ISCLI factory default user names

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 shown in Example B-8.

Example: B-8 Change admin password and create 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
```

B.5 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.

B.5.1 Juniper Junos OS

Example B-9 shows how to configure the syslog level and how to configure a syslog server in Junos OS.

Example: B-9 JNOS syslog server configuration

```
admin@EX4500-VC# set system syslog host <ip address> any error
```

B.5.2 IBM Networking OS

Example B-10 shows how to configure the syslog level and server in IBM isCLI.

Example: B-10 IBM Networking OS syslog configuration

```
switch(config)#logging host 1 address <ip address>  
switch(config)#logging host 1 severity 3
```

B.6 Network Time Protocol

This section includes information about Network Time Protocol (NTP).

B.6.1 Juniper Junos OS

To configure time zone and NTP server in Junos OS, use the commands described in Example B-11.

Example: B-11 Junos OS time zone and NTP configuration

```
admin@EX4500-VC# set system time-zone America/New_York  
admin@EX4500-VC# set system ntp server <ip address> prefer
```

B.6.2 IBM Networking OS

To configure time zone and NTP server in IBM Networking OS, use the commands that are shown in Example B-12.

Example: B-12 IBM Networking OS time zone and NTP configuration

```
switch(config)#system timezone 145
switch(config)#system daylight
switch(config)#ntp enable
switch(config)#ntp primary-server <ip address> mgt-port
```

B.7 Simple Network Management Protocol

Simple Network Management Protocol (SNMP) community strings, SNMP v3 users, SNMP views, SNMP traps, and SNMP target servers can be configured by using the examples described in this section. SNMP v1 and v2 also can be configured.

B.7.1 Juniper Junos OS

Example B-13 shows how to configure various SNMP v3 access on Junos OS.

Example: B-13 Junos OS SNMPv3 configuration

```
admin@EX4500-VC# set snmp v3 usm local-engine user secure authentication-md5
authentication-password <authentication-password>
admin@EX4500-VC# set snmp v3 vacm security-to-group security-model usm
security-name secure group AdminAuthPriv

admin@EX4500-VC# set snmp v3 vacm access group AdminAuthPriv
default-context-prefix security-model usm security-level privacy read-view ALLVIEW
admin@EX4500-VC# set snmp v3 vacm access group AdminAuthPriv
default-context-prefix security-model usm security-level privacy write-view
ALLVIEW

admin@EX4500-VC# set snmp view ALLVIEW oid internet include
```

B.7.2 IBM Networking OS

Example B-14 shows how to configure various SNMP v3 access on IBM Networking OS with an encrypted MD5 password that uses Data Encryption Standard (DES) encryption.

Example: B-14 IBM Networking OS SNMPv3 configuration

```
snmp-server user 4 name "secure"
snmp-server user 4 authentication-protocol md5 authentication-password
"<encrypted>"
snmp-server user 4 privacy-protocol des privacy-password "<encrypted>"
!
snmp-server group 4 user-name secure
snmp-server group 4 group-name "admingrp"
!
snmp-server access 4 level authPriv
```

```
snmp-server access 4 read-view "iso"  
snmp-server access 4 write-view "iso"  
!  
snmp-server version v3only
```

B.8 Link Layer Discovery Protocol

To enable Link Layer Discovery Protocol (LLDP) globally and on a per interface basis, use the commands that are described in this section.

B.8.1 Juniper Junos OS

By default, LLDP is enabled in Junos OS for all interfaces. Example B-15 shows how to configure LLDP and how to disable LLDP in a particular interface.

Example: B-15 Junos OS LLDP configuration

```
admin@EX4500-VC# set protocols lldp interface all  
admin@EX4500-VC# set protocols lldp interface xe-0/0/0.0 disable
```

B.8.2 IBM Networking OS

By default, LLDP in Junos OS is enabled for all interfaces. Example B-16 shows how to configure LLDP and how to disable LLDP in a particular interface.

Example: B-16 LLDP configuration in IBM isCLI

```
switch(config)#lldp enable  
switch(config)#interface port 1  
switch(config-if)#no lldp admin-status
```

B.9 Spanning Tree Protocol

Junos OS and IBM Networking OS can run different versions of Spanning Tree Protocol (STP).

B.9.1 Juniper Junos OS

Rapid Spanning Tree Protocol (RSTP) is enabled by default in Junos OS. To configure a different STP mode, first disable RSTP and enable another mode, such as STP or VLAN STP (VSTP), as shown in Example B-17. VSTP runs a separate instance of spanning tree for each VLAN and is compatible with Per-VLAN Rapid Spanning Tree (PVRST).

Example: B-17 Junos OS STP configuration

```
admin@EX4500-VC# delete protocols rstp  
admin@EX4500-VC# set protocols vstp vlan all
```

B.9.2 IBM Networking OS

Spanning tree mode by default in IBM isCLI is PVRST, By using PVRST, each VLAN runs a separate instance of spanning tree. To configure other STP modes, see Example B-18.

Example: B-18 Configure different STP modes on IBM Networking OS

```
switch(config)#spanning-tree mode <disable/mst/pvrst/rstp>
```

B.10 Interface speed and duplex

Set the interface speed and duplex by using the commands that are described in this section.

B.10.1 Juniper Junos OS

Example B-19 shows how to set interface speed and duplex in Junos OS.

Example: B-19 Speed and duplex interface configuration

```
admin@EX4500-VC# set interface ge-0/0/0 ether-options link-mode full-duplex
admin@EX4500-VC# set interface ge-0/0/0 ether-options speed 100m
```

B.10.2 IBM Networking OS

Example B-20 shows how to configure duplex and port speed in IBM Networking OS.

Example: B-20 Speed and duplex interface configuration

```
switch(config)#interface port 12
switch(config-if)#speed 100
switch(config-if)#duplex full
```

B.11 VLAN

To add new VLANs, use following commands.

B.11.1 Juniper Junos OS

Example B-21 shows how to add a new VLAN in Junos OS.

Example: B-21 VLAN configuration

```
admin@EX4500-VC# set vlan <vlan-name> vlan-id <vlan-id-number>
```

B.11.2 IBM Networking OS existing syntax

Example B-22 shows how to configure new VLANs in IBN Networking OS switches by using existing syntax.

Example: B-22 VLAN configuration in IBM Networking OS existing syntax

```
switch(config)# vlan <vlan>
switch(config-vlan)# enable
switch(config-vlan)# member <portX-portY>
```

B.11.3 IBM Networking OS isCLI syntax

Example B-23 shows how to configure new VLANs in IBN Networking OS switches by using isCLI syntax.

Example: B-23 VLAN configuration in IBM Networking OS isCLI syntax

```
switch(config)# vlan <vlan>
switch(config-vlan)# no shutdown
```

B.12 VLAN Tagging (802.1q)

VLAN tagging is supported with 802.1q protocol in Junos OS and IBM Networking OS.

B.12.1 Juniper Junos OS

Example B-24 shows how to configure 802.1q VLAN tagging in Juniper switches. The example also shows the configuration of a native-vlan-id if it is needed.

Example: B-24 Port configuration for VLAN tagging

```
admin@EX4200# set interfaces ge-0/0/<port#>.0 family ethernet-switching port-mode trunk
admin@EX4200# set interfaces ge-0/0/<port#>.0 family ethernet-switching vlan members <VLAN#1>,<VLAN#2>
admin@EX4200-VC# set interfaces ge-0/0/<port#>.0 family ethernet-switching native-vlan-id <VLAN>
```

B.12.2 IBM Networking OS existing syntax

In IBM Networking OS, enable tagging on the port. Use a pvid=1 if no systems on these ports must have a native VLAN defined. Untagged frames on interfaces with tagging enabled must be given a Port VLAN identifier (PVID) if the end station device cannot tag or recognize 802.1q frames. Similar functionality is enabled on Juniper equipment with native-vlan-id. Example B-25 shows the port configuration for VLAN tagging.

Example: B-25 Port configuration for VLAN tagging

```
swtich(config)#interface port <portX-portY>
swtich(config-if)#tagging
swtich(config-if)#pvid 1
```

B.12.3 IBM Networking OS isCLI syntax

As of version 7.6 and above, isCLI configuration command syntax also supports more industry-standard commands, as shown in Example B-26.

Example: B-26 Port configuration for VLAN tagging (isCLI syntax)

```
switch(config)# interface port <portX-portY>
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan <VLAN>
switch(config-if)# switchport trunk native vlan <VLAN>
switch(config-if)# exit
```

B.13 Trunking and link aggregation: Static

To create a port channel between a Juniper switch and IBM switch, use the commands that are described in this section.

B.13.1 Juniper Junos OS

Use the link aggregation feature to aggregate one or more links to form a virtual link or link aggregation group (LAG). Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability. Example B-27 shows the configuration of LAG in Juniper switches.

Example: B-27 Junos OS link aggregation configuration

```
admin@EX4200# set chassis aggregated-devices ethernet device-count 1

admin@EX4200# set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
admin@EX4200# set interfaces ae0 unit 0 family ethernet-switching vlan members
<VLAN#1>
admin@EX4200# set interfaces ge-0/0/4 ether-options 802.3ad ae0
admin@EX4200# set interfaces ge-0/0/5 ether-options 802.3ad ae0
```

B.13.2 IBM Networking OS existing syntax

Example B-28 shows how to configure a port channel in IBM Networking OS existing syntax.

Example: B-28 Port channel configuration in IBM Networking OS existing syntax

```
switch(config)#interface port 18,20
switch(config-if)#tagging
switch(config-if)#pvid 1

switch(config)#portchannel 5 port 18
switch(config)#portchannel 5 port 20
switch(config)#portchannel 5 enable
```

B.13.3 IBM Networking OS isCLI syntax

As of version 7.6 and above, isCLI configuration command syntax also supports more industry-standard commands. Example B-29 shows how to configure a port channel in IBM isCLI syntax.

Example: B-29 Port channel configuration in IBM isCLI syntax

```
swtich(config)#interface port 18,20
swtich(config-if)#switchport mode trunk
swtich(config-if)#switchport trunk native vlan <VLAN>

swtich(config)#portchannel 5 port 18
swtich(config)#portchannel 5 port 20
swtich(config)#portchannel 5 enable
```

B.14 Trunking and link aggregation: Link Aggregation Control Protocol

Link Aggregation Control Protocol (LACP) (IEEE 802.3ad) is one method of bundling several physical interfaces to form one logical interface.

B.14.1 Juniper Junos OS

Example B-30 shows how to configure an LACP in Juniper switches. The configuration is almost identical to creating a standard link aggregation except the LACP active command is different.

Example: B-30 Configure LACP on Juniper switches

```
admin@EX4200# set chassis aggregated-devices ethernet device-count 1

admin@EX4200# set interfaces ae0 aggregated-ether-options lacp active
admin@EX4200# set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
admin@EX4200# set interfaces ae0 unit 0 family ethernet-switching vlan members
<VLAN#1>
admin@EX4200# set interfaces ge-0/0/4 ether-options 802.3ad ae0
admin@EX4200# set interfaces ge-0/0/5 ether-options 802.3ad ae0
```

B.14.2 IBM Networking OS existing syntax

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 LACP groups, as shown in Example B-31 on page 318.

Example: B-31 Configure LACP on IBM switches by using existing syntax

```
switch(config)#interface port 18,20
switch(config-if)#tagging
switch(config-if)#pvid 1
```

```
switch(config)# interface port 18,20
switch(config-if)# lacp key 2000
switch(config-if)# lacp mode active
```

B.14.3 IBM Networking OS isCLI syntax

As of version 7.6 and above, isCLI configuration command syntax also supports more industry-standard commands.

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 LACP groups, as shown in Example B-32.

Example: B-32 Configure LACP on IBM switches using isCLI syntax

```
switch(config)#interface port 18,20
switch(config-if)#switchport mode trunk
switch(config-if)#switchport trunk native vlan <VLAN>
```

```
switch(config)# interface port 18,20
switch(config-if)# lacp key 2000
switch(config-if)# lacp mode active
```

B.15 External authentication

Remote authentication dial-in user service (RADIUS) is an authentication protocol commonly used among a team of network engineers.

B.15.1 Juniper Junos OS

Example B-33 shows a common RADIUS implementation on Junos OS.

Example: B-33 Juniper External Authentication commands

```
admin@EX4200# set system radius-server <ip address> secret SECRET
admin@EX4200# set system authentication-order [ radius password ]
```

B.15.2 IBM Networking OS

Similar functionality can be implemented in IBM Networking OS, as shown in Example B-34.

Example: B-34 IBM External Authentication commands

```
switch(config)#radius-server primary-host <ip address> key SECRET
switch(config)#radius-server enable
```

B.16 Bridge Protocol Data Unit Guard

Bridge Protocol Data Unit (BPDU) Guard often is used to prevent general users from plugging management network equipment into the overall network infrastructure.

B.16.1 Juniper Junos OS

Example B-35 shows a common implementation on Junos OS.

Example: B-35 Juniper BPDU Guard commands

```
admin@EX4200# set ethernet-switching-options bpdu-block ge-0/0/1
admin@EX4200# set ethernet-switching-options bpdu-block disable-timeout 1800
```

B.16.2 IBM Networking OS

Similar functionality can be implemented in IBM Networking OS, as shown in Example B-36.

Example: B-36 IBM Networking OS BPDU Guard commands

```
switch(config)#interface port 1
switch(config-if)#bpdu-guard
```

B.17 Dynamic Host Configuration Protocol snooping

Dynamic Host Configuration Protocol (DHCP) snooping is a DHCP security feature that provides network security by filtering untrusted DHCP messages and building and maintaining a DHCP snooping binding database, which 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.

B.17.1 Juniper Junos OS

Configuration for Junos OS DHCP snooping is shown in Example B-37.

Example: B-37 Junos OS DHCP snooping configuration

```
admin@EX4500-VC# set ethernet-switching-options secure-access-port vlan DATA
examine-dhcp
admin@EX4500-VC# set ethernet-switching-options secure-access-port interface
<interface> dhcp-trusted
```

B.17.2 IBM Networking OS

Example B-38 shows the configuration of DHCP snooping in IBM Networking OS.

Example: B-38 Enabling DHCP snooping in IBM Networking OS

```
switch(config)#ip dhcp snooping
switch(config)#ip dhcp snooping vlan <vlan>
switch(config)#interface port <port>
switch(config-if)#ip dhcp snooping trust
```

B.18 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.

B.18.1 Juniper Junos OS

Example B-39 shows how to configure port mirroring in Junos OS. In the example, ingress and egress traffic for port ge-0/0/0.0 is copied to port ge-0/0/2.0

Example: B-39 Junos O port mirroring

```
admin@EX4200# set ethernet-switching-options analyzer MIRROR input ingress
interface ge-0/0/0.0
admin@EX4200# set ethernet-switching-options analyzer MIRROR input egress
interface ge-0/0/0.0
admin@EX4200# set ethernet-switching-options analyzer MIRROR output interface
ge-0/0/2.0
```

B.18.2 IBM Networking OS

IBM System Networking switches support a many-to-one mirroring model. As shown in Example B-40, traffic for ports 1 is monitored by port 3. In the example, ingress traffic and egress traffic on port 1 are copied and forwarded to the monitor port 3, but you also can configure only ingress or egress traffic.

Example: B-40 IBM Networking OS port mirroring

```
switch(config)#port-mirroring monitor-port 1 mirroring-port 3 both
```

B.19 Open Shortest Path First configuration

Some of the more basic Open Shortest Path First (OSPF) commands for Junos OS and IBM isCLI are described in this section.

B.19.1 Juniper Junos OS

Example B-41 shows basic OSPF configuration commands on Junos OS.

Example: B-41 Configuration commands for OSPF on Juniper Junos OS

```
admin@EX4200# set interface lo0 unit 0 family inet address <ip address>/32
```

```
admin@EX4200# set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
```

```
admin@EX4200# set protocols ospf area 0.0.0.0 interface lo0.0
```

```
admin@EX4200# set routing-options router-id <loopback ip address>
```

B.19.2 IBM isCLI

Example B-42 shows basic OSPF configuration commands on IBM isCLI.

Example: B-42 Configuration commands for OSPF on IBM isCLI

```
switch(config)#interface loopback 1
switch(config)#ip address <ip address> <subnet mask>
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
```



Easy Connect

IBM Easy Connect is a simple configuration mode implemented on IBM System Networking Ethernet and Converged switches that enables easy integration of IBM Flex/PureSystems with existing Juniper 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 and supports your existing and future network.

This appendix includes the following topics:

- ▶ Introduction to IBM Easy Connect
- ▶ Single Mode
- ▶ Storage Mode
- ▶ Easy Connect Multi-Chassis Mode
- ▶ Customer examples with diagrams
- ▶ Easy Connect limitations

C.1 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 comprising PureSystems compute, storage, system management and networking resources, and then manage this scalable resource with the simplicity of a single network node.

The following IBM System Networking Ethernet switches support the Easy Connect feature:

- ▶ IBM Flex System Fabric EN4093/EN0493R and Virtual Fabric 10 Gb Scalable Switches
- ▶ IBM Flex System Fabric CN4093 10 Gb Converged Scalable Switch
- ▶ IBM System Networking RackSwitch G8264CS
- ▶ IBM RackSwitch G8264 or G8124E
- ▶ IBM RackSwitch G8214 (not in FCoE mode)

Easy Connect mode provides transparent PureSystems connectivity to your existing Juniper 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 coming 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 with the main difference being that intra-chassis switching is supported.

The Spanning Tree Protocol (STP) is disabled on the supported IBM System Networking switch in all Easy Connect modes, which eliminates the data center administrator’s spanning tree concerns. This loop-free topology requires no other configuration after it is set up, and helps to provide economical bandwidth use with prioritized pipes and network virtualization for Intel and Power Compute Nodes.

C.2 Single Mode

Easy Connect Single Mode allows the IBM Flex System EN4093/R switch to act transparently to the upstream network. Customers that use Active/Passive NIC teaming with no NIC bonding (LACP or static PortChannel) on the Compute Nodes are well-suited with Single Mode, as show in Figure C-1 on page 325.




Figure C-1 IBM Easy Connect Single Mode diagram

Single Mode includes the following important distinctions:

- ▶ All local Layer-2 traffic that is pointing to the same I/O Bay in the Enterprise Chassis remains within the same chassis.
- ▶ Because the CN4093 or EN4093/EN4093R I/O modules are not connected with a Virtual Link Aggregation Group (vLAG), traffic that is destined for Compute Nodes by using different I/O Bays within the same Enterprise Chassis must travel to the upstream switch and then back down.
- ▶ Each Enterprise Chassis appears as two separate devices to the upstream network when two I/O modules are used.

C.2.1 Implementation

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 by using Telnet or Secure Shell (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 6-37. Enable the CLI prompt in the last step if the Flex System Manager (FSM) is used in the environment.

Example 6-37 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, use the commands that are shown in Example 6-38 after it is connected via Telnet or SSH.

Example 6-38 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, use the commands that are shown in Example 6-39 to enable Easy Connect Single Mode.

Example 6-39 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
```

Easy Connect Single Mode is now implemented.

Important: The IBM Virtual Fabric Switch Module (VFSM) for the IBM BladeCenter H or HT chassis is supported with Easy Connect Single and Storage Mode or Modes. Configuration steps are identical. This also can be done in a System x environment with rack servers by using the G8124, G8264, or G8264CS.

Now that Easy Connect Single Mode is enabled, you might want to configure spanning tree Bridge Protocol Data Unit (BPDU) Guard and Edge on the upstream switch for more protection.

C.3 Storage Mode

Easy Connect Storage Mode allows the IBM Flex System EN4093/R switch to act transparently to the upstream network when Fibre Channel over Ethernet (FCoE) traffic is run. Storage Mode is nearly identical to Single Mode from a configuration standpoint, the only difference is that Converged Enhanced Ethernet (CEE) must be enabled for FCoE to function. Storage Mode is shown in Figure C-2 on page 327.




Figure C-2 BM Easy Connect Storage Mode diagram

The distinctions that are listed for Single Mode are the same for Storage Mode.

C.3.1 Implementation

To configure the CN4093 or EN4093/EN4093R I/O modules for Easy Connect Storage Mode, first perform steps 1, 2, and 3 as described in C.2.1, “Implementation” on page 325.

After the steps are completed, use the commands that are shown in Example 6-40 to implement Storage Mode. The only difference is highlighted in bold text.

Example 6-40 Implementing Easy Connect Storage Mode

```
spanning-tree mode disable
portchannel 1 port ext1-ext10 enable
vnic enable
  vnic vnicgroup 1
  vlan 4091
  port intal-intal4
  portchannel 1
  enable
  failover
  exit
cee enable
write memory
```

Easy Connect Storage Mode is now implemented.

Important: The same considerations that are listed for Single Mode and next steps also apply to Storage Mode, with the exception that IBM Networking OS 7.6 and earlier does not support FCoE traffic over multiple aggregated links, by using LACP or static PortChannels.

C.4 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 a single entity to the upstream network when connected with a Virtual Link Aggregation Group (vLAG). Customers that use Active/Active NIC teaming with Link Aggregation Control Protocol (LACP, or IEEE 802.3ad), or Static IP Hash on the Compute Node are best-suited with Multi-Chassis Mode, as shown in Figure C-3.

Multiple chassis: Alternatively, this might be multiple chassis' connected to a pair of G8264s at the top-of-rack going out to clients existing network.

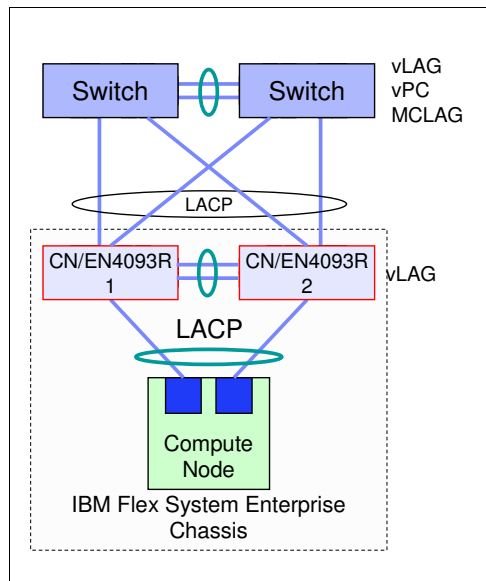


Figure C-3 IBM Easy Connect Multi-Chassis Mode diagram

Multi-Chassis Mode includes the following important distinctions:

- ▶ Because the CN4093 or EN4093/EN4093R I/O modules are connected with a vLAG inter-switch link (ISL), all Layer 2 traffic that is destined for Compute Nodes that use the same or different I/O Bays within the same Enterprise Chassis never leave the chassis.
- ▶ Each Enterprise Chassis appears as a single device to the upstream network when two I/O modules are used.
- ▶ All operating systems (AIX, Linux, Windows, VMware, VIO, and so on) within the IBM Flex System Enterprise Chassis must use VLAN tags.

Exception: If the Flex System Manager (FSM) is used, the client must enable the top-of-rack port Native VLAN ID with the VLAN on which the FSM is configured.

- ▶ 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 traffic separation or use IBM Virtual Fabric Mode, each vNIC Group requires its own uplink/PortChannel.

Important: IBM Flex System POWER Nodes support pNIC mode only as of this writing.

- Multi-Chassis Mode allows for the eventual implementation of IBM Virtual Fabric Mode.

C.4.1 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. For more information, see Example 6-38 on page 326.
2. Disable the SPT globally.
3. Configure all the internal (INT) and external (EXT) CN4093 or EN4093/R ports by using the tagpvid-ingress keyword that 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 by using 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.

Easy Connect Multi-Chassis Mode is now implemented on the CN/EN4093/R.

Important: The IBM Virtual Fabric Switch Module (VFSM) for the IBM BladeCenter H or HT chassis does not work in Multi-Chassis Mode because it does not support vLAG as of this writing.

A sample script to enable Easy Connect Multi-Chassis Mode on the CN/EN4093/R I/O module is shown in Example 6-41.

Example 6-41 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

```

C.4.2 Implementation with G8264

If the client is using a pair of IBM RackSwitch G8264 switches in the overall topology (as shown in Figure C-4) such as in a pre-racked, pre-cabled IBM PureFlex System Express, Standard, or Enterprise rack configuration, the following section describes how Easy Connect can be used.

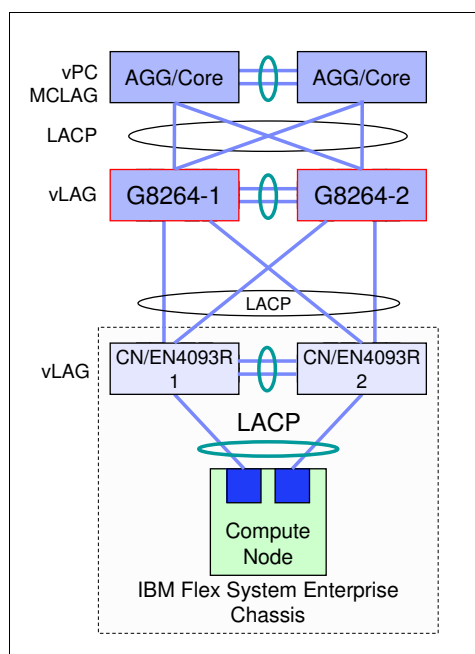


Figure C-4 IBM Easy Connect Multi-Chassis Mode with RackSwitch G8264

To configure the RackSwitch G8264 for Easy Connect Multi-Chassis Mode, complete the following steps:

1. Restore the factory default configuration to the G8264. For more information, see Example 6-38 on page 326.
2. Disable the SPT globally.
3. Configure all the upstream and downstream G8264 ports by using the tagpvid-ingress keyword that use VLAN 4091 as the PVID.
4. Enable 802.1Q VLAN tagging on the ports that are used as the vLAG Peer Link between the G8264s that are using 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.

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 6-42.

Example 6-42 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

```

Now that Easy Connect Multi-Chassis Mode is enabled, you might want to configure spanning tree BPDUs Guard and Edge on the upstream switch for more protection.

C.5 Customer examples with diagrams

This section describes common implementation scenarios with Easy Connect for various industries that purchased IBM PureFlex System hardware. Requirements are listed as dictated by the customer, and a network diagram to fit those requirements is proposed.

C.5.1 Telecommunications customer

This customer requires the following specifications:

- ▶ No STP or any other protocols are seen by the network.
- ▶ The EN4093/R I/O modules in the IBM Flex System Enterprise Chassis must be completely transparent devices that require no management by any group after the initial setup.

Figure C-5 shows how Easy Connect satisfies all of the telecommunications customer requirements.

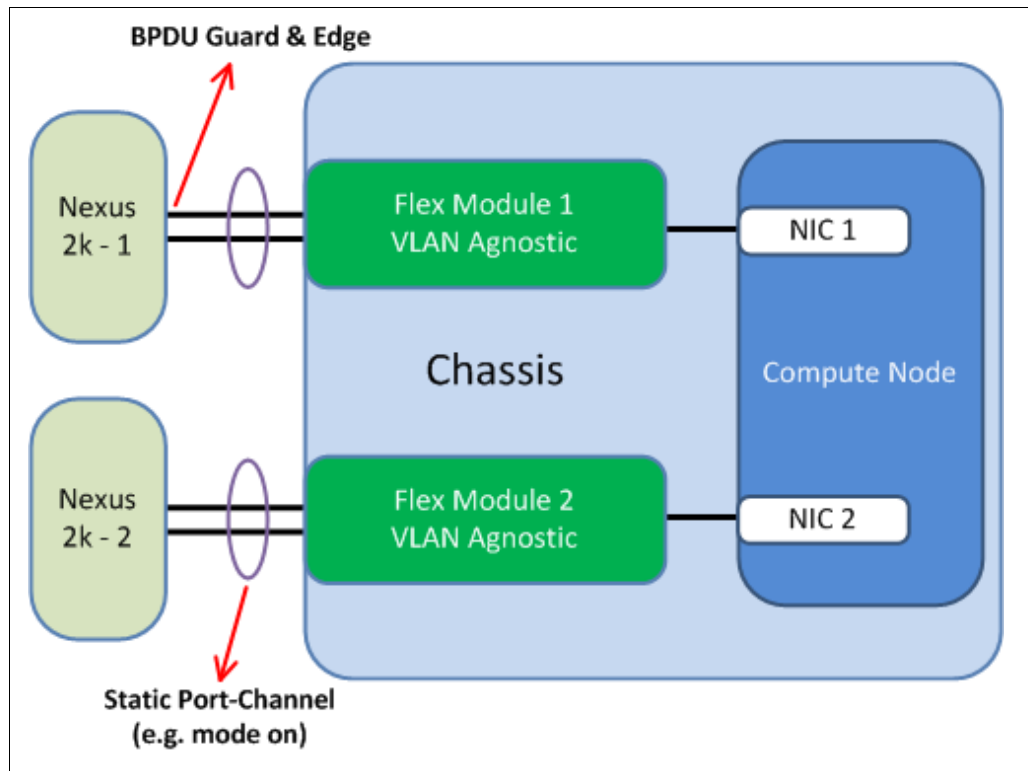


Figure C-5 Telecommunications customer network diagram

C.5.2 State government customer

This customer requires the following specifications:

- ▶ A local area network (LAN) on Motherboard (LoM) in Virtual Fabric Mode must be used so bandwidth can be adjusted for each vNIC as required.
- ▶ A dedicated uplink vPC PortChannel from each EN4093/R must be used for each vNIC Group for separation of traffic.
- ▶ The EN4093/R I/O modules in the IBM Flex System Enterprise Chassis must be completely transparent devices that require no management by any group after the initial setup.

Figure C-6 shows how Easy Connect satisfies all of the state government customer requirements.

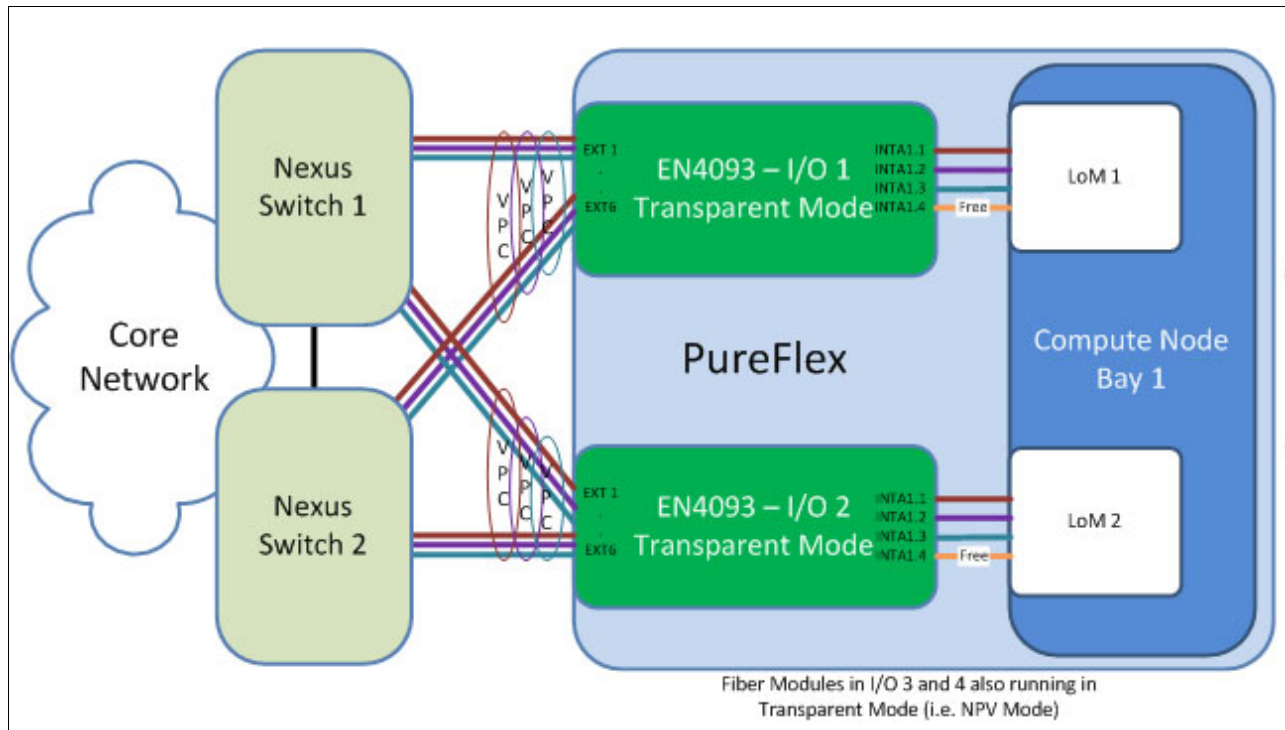


Figure C-6 State Government Customer network diagram

C.5.3 Medical center customer

This customer requires the following specifications:

- ▶ Separation of and dedicated Fibre Channel and Ethernet from each compute node and IBM Flex System Enterprise Chassis.
- ▶ Total hardware redundancy, including NIC and ASIC on each compute node that uses the CN4054 mezzanine adapter.
- ▶ Transparency on Ethernet (Easy Connect) and Fibre Channel (NPV).

Figure C-7 on page 334 shows how Easy Connect satisfies all of the medical center customer requirements.

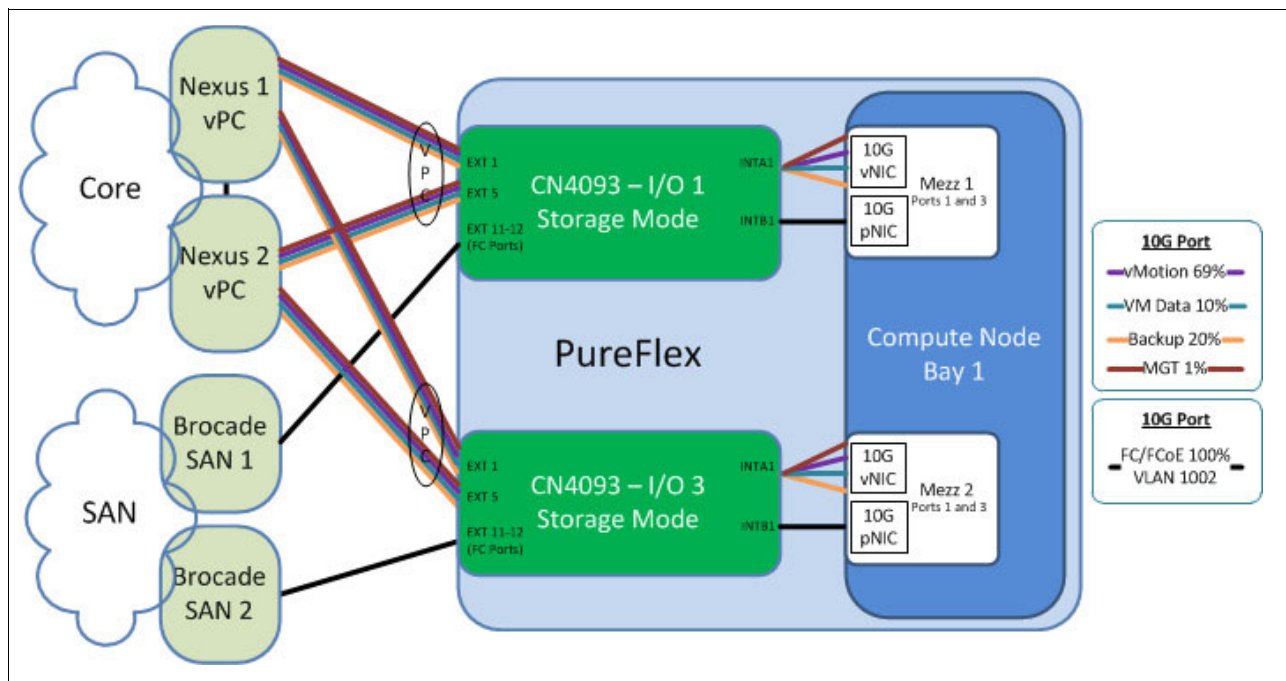


Figure C-7 Medical center customer network diagram

C.6 Easy Connect limitations

When the switch 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 as each vNIC group is a single broadcast domain.

Related publications

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

IBM Redbooks intellectual property

The following IBM Redbooks publications provide more information about the topic in this document. Note that some publications that are referenced in this list might be available in softcopy only:

- ▶ *Moving to IBM PureFlex System x86-to-x86 Migration*, REDP-4887
- ▶ *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984
- ▶ *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984
- ▶ *Implementing Systems Management of IBM PureFlex System*, SG24-8060
- ▶ *IBM System Networking RackSwitch G8264/G8264T*, TIPS0815

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft and other materials, at the following website:

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

Online resources

The following websites also are 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/>
- ▶ Juniper Networks
<http://www.juniper.net>

Help from IBM

IBM Support and downloads:

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

IBM Global Services

<http://www.ibm.com/services>



IBM Flex System and PureFlex System Network Implementation with Juniper Networks

(0.5" spine)

0.475" <-> 0.873"

250 <-> 459 pages



IBM Flex System and PureFlex System Network Implementation with Juniper Networks

Learn how to implement the IBM Flex System and PureFlex System

Learn how to connect to Juniper Networks

Learn troubleshooting techniques

To meet today's complex and ever-changing business demands, you need a solid foundation of server, storage, networking and software resources that is simple to deploy and can 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, which is a part of the IBM PureSystems family of expert integrated systems, combines advanced IBM hardware and software along with patterns of expertise and integrates them into three optimized configurations that are simple to acquire and deploy so you can achieve faster time to value.

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

In this IBM Redbooks publication, we use EX4500 core switches to demonstrate interoperability with the System Networking switches (RackSwitch G8264 top of rack switch and the Flex system fabric EN4093 10Gb scalable switch).

We also describe a redundant environment using QFX3500 switches running IBM Virtual-Link Aggregation Group (MC-LAG/vLAG) and Juniper Multi- Chassis-Link Aggregation Group.

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