


Deploying IBM Flex System into a Cisco Network

Learn how to integrate IBM Flex System into your network

See real life Layer 2 configurations with Flex System switches

Find out how easy it is to connect network devices



Christoph Raisch
Bernd Albrecht
Peter Demharter
Stephan Fleck
Joachim Gross
Ruediger Rissmann
Werner Sponer
Arwed Tschoeke
Pietro Volante



International Technical Support Organization

Deploying IBM Flex System into a Cisco Network

February 2013

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

First Edition (February 2013)

This edition applies to the following switches and firmware levels:

- ▶ IBM Flex System EN2092 1Gb Ethernet Scalable Switch: Version 7.2.2.2
- ▶ IBM RackSwitch G8264: Version 7.2.2.0
- ▶ Cisco Nexus 5000: Version 5.1(3)N2(1)
- ▶ Cisco Catalyst 6500: Version 12.2.33-SXH8a

© 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
Foreword	ix
Preface	xi
The team who wrote this paper	xi
Now you can become a published author, too!	xiii
Comments welcome.	xiv
Stay connected to IBM Redbooks	xiv
Chapter 1. Introduction	1
1.1 Networking	2
1.2 IBM PureSystems	2
1.3 Switch configuration	3
1.4 How to use this paper	4
Chapter 2. Layer 2 Network protocols and technologies	5
2.1 Basic frame forwarding concepts	6
2.2 Virtual local area network	7
2.3 Spanning tree	8
2.3.1 Spanning Tree Protocol: IEEE 802.1D	9
2.3.2 Rapid Spanning Tree Protocol: IEEE 802.1w	10
2.3.3 Multi-instance Spanning Tree Protocol: IEEE 802.1s	10
2.3.4 Per VLAN Rapid Spanning Tree	11
2.4 Link aggregation	11
2.4.1 Link Aggregation Control Protocol	12
2.4.2 Virtual Link Aggregation Groups	12
2.4.3 Cisco Virtual Port Channel	13
2.4.4 Link Layer Discovery Protocol: 802.1AB	13
Chapter 3. IBM RackSwitch G8264 connectivity	15
3.1 Prerequisites	16
3.2 Use Case 1: PVRST	16
3.2.1 Verifying the topology by using lldp	18
3.2.2 Verifying trunks	19
3.2.3 Verifying PVRST spanning tree configurations	19
3.2.4 Show running-config of all switches in Use Case 1	25
3.3 Use Case 2: Link aggregation and PVRST	30
3.3.1 Verifying the topology that is used by using lldp	32
3.3.2 Verifying trunks	33
3.3.3 Verifying link aggregation by using lacp	34
3.3.4 Verifying PVRST spanning tree configuration	34
3.3.5 Show running-config of all switches in Use Case 2	41
3.4 Use Case 3: Link aggregation and MST	47
3.4.1 Verifying the topology that was used by using lldp	49
3.4.2 Verifying trunks	50
3.4.3 Verifying link aggregation by using lacp	51
3.4.4 Verifying MST spanning tree configuration	51

3.4.5 Show running-config of all switches in Use Case 3	55
3.5 Use Case 4: Link aggregation, MSTP and VLAG	60
3.5.1 Verifying the topology by using lldp	62
3.5.2 Verify interface status	62
3.5.3 Verifying trunks	64
3.5.4 Verify spanning tree	65
3.5.5 Verify virtual link aggregation	66
3.5.6 Show running-config of all switches in Use Case 4	70
3.6 Use Case 5: Link aggregation and VLAG without STP	78
3.6.1 Verifying the topology by using lldp	79
3.6.2 Verify interface status	79
3.6.3 Verifying trunks	81
3.6.4 Verify virtual link aggregation	82
3.6.5 Show running-config of all switches in Use Case 5	86
Chapter 4. Cisco Nexus 5000 connectivity	95
4.1 Prerequisites	96
4.2 Use Case 1: PVRST	96
4.2.1 Verifying the topology that is used by using lldp	96
4.2.2 Verifying trunks	98
4.2.3 Show running-config of all switches in Use Case 1	105
4.3 Use Case 2: PVRST with LACP Channeling	110
4.3.1 Verifying the topology used by using lldp	110
4.3.2 Verifying trunks	112
4.3.3 Verifying PVRST spanning tree configuration	115
4.3.4 Bridge priority field in the show spanning tree output	118
4.3.5 Show running-config of all switches in Use Case 2	122
4.4 Use Case 3: MST with LACP Channeling	127
4.4.1 Verifying the topology used by using lldp	128
4.4.2 Verifying trunks	128
4.4.3 Verifying MST spanning tree configuration	132
4.4.4 Show running-config of all switches in Use Case 3	135
4.5 Use Case 4: MST with LACP Channeling and vPC	141
4.5.1 Configuring vPC on STR	141
4.5.2 Configuring MST on the STR	143
4.5.3 Configuring vPC on VIE	143
4.5.4 Configuring MST on VIE	144
4.5.5 Reviewing the Flex System switch configuration	145
4.5.6 Configuring MST on the Flex System switch	146
4.5.7 Logical view	146
4.5.8 Verifying the configuration	147
4.5.9 Verifying the vPC configuration on VIE	150
4.6 Use Case 5: LACP Channeling and vPC without spanning tree	162
4.6.1 Configuring vPC on STR	163
4.6.2 Configuring vPC on VIE	166
4.6.3 Disabling STP on the Flex System switch	169
Chapter 5. Cisco Catalyst 6500 switch connectivity	171
5.1 Use Case 1: LACP channeling and vPC without spanning tree	172
5.1.1 Catalyst 6500 switch configuration	172
5.1.2 Flex System switch configuration	174
Appendix A. Troubleshooting	177
Basic troubleshooting for connectivity problems	178

Approach.	178
Layer 2 troubleshooting commands	179
Baseline documentation.	181
Additional useful information for baseline documentation.	183
Firmware update of IBM Flex System network switches	183
Update the switch by using the web-based GUI	183
Using SSHv2 or Telnet	187
Abbreviations and acronyms	191
Related publications	193
IBM Redbooks	193
Other publications	193
Online resources	194
Help from IBM	194

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:

IBM Flex System™
IBM®
PureApplication™
PureFlex™

PureSystems™
RackSwitch™
Redbooks®
Redpaper™

Redbooks (logo) ®
System x®
zEnterprise®

The following terms are trademarks of other companies:

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.

Foreword

This IBM® Redpaper™ publication was initiated and authored by members of the Technical Expert Council, Central Region (TEC CR) workgroup “Workload optimized networks” that was founded November 2011. When IBM moved back into the networking market with the acquisition of BNT, this move positioned IBM to capture a significant share of an emerging market for converged fabrics.

The initial idea of the workgroup was that Ethernet will become pervasive for all aspects of networking and storage in the next couple of years, which requires users to rethink how connectivity aspects become an integrated part of any computing solution. The workgroup established an expert community to bring development expertise, networking background, and customer and market insights together. Business sponsor of the TEC workgroup is Erich Baier, IBM Vice President, who is responsible for Modular Systems and Networking Development.

The TEC CR is the local affiliate for Germany, Switzerland, and Austria of the IBM Academy of Technology (AoT). The mission of the TEC CR is to strengthen the technical leadership in the local markets through promoting communication among experts and by consulting the executive management of IBM. It identifies and pursues technical opportunities that are relevant to the business of IBM, and aims to advance the technology base of IBM and its application in market-leading products, solutions, and services.

A major finding from the collaboration in the workgroup was that with the announcement of IBM PureSystems, many clients will have to integrate IBM Flex System into a typical Cisco dominated customer network. However, the documentation that is needed to complete this integration was not readily available. In close collaboration with the development labs, the group took initiative to close the gap and wrote this paper.

This paper is a good example of a collaborative effort of technical experts and leaders from different organizations that results in a holistic view of the relevant steps that are needed to make a solution successful in the market. As a chairman of the TEC CR, I would like to thank the authors of the paper for this initiative.

Thomas Harrer
Chairman, Technical Expert Council, Central Region (TEC CR)
Member IBM Academy of Technology

Preface

This IBM® Redpaper™ publication provides information about how to integrate an IBM Flex System into an existing customer network. It focuses on interoperability and seamless integration from the network perspective.

The paper describes the complete configuration of the most common scenarios. It guides you through several setups, and shows in detail how to configure the network switches and verify the functionality and proper operation.

This paper can help you to easily configure and monitor your Layer 2 setup. Typical, well-established Layer 2 Network setups use combinations of Spanning Tree Protocol, VLANs, and link aggregation.

The scenarios that are described in this paper include the use of the following switching products:

- ▶ Cisco Nexus 5000 (including vPC)
- ▶ Cisco Catalyst 6500
- ▶ IBM RackSwitch (including VLAG)
- ▶ IBM Flex System Ethernet Scalable Switch (including VLAG)

We describe the use of these switches with each of the following Spanning Tree Protocol (STP) configurations:

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

The paper is for network administrators who are familiar with Cisco network products. It uses the industry standard command-line interface (isCLI) as the management interface. It is assumed that the reader is familiar with Cisco products and the use of isCLI.

The team who wrote this paper

This paper was produced by a team of specialists from around the world.

Christoph Raisch is a Senior Technical Staff Member at IBM Germany Research & Development, Boeblingen. He has 15 years of experience in defining and implementing firmware architectures in the areas of Fibre Channel, InfiniBand, PCI Express, Ethernet, and FCoE for different IBM platforms. He received a Dipl.-Ing. degree in Electrical Engineering from the University of Stuttgart. He works on future technologies for IBM networking switches.

Bernd Albrecht is an IT Specialist in IBM Germany specializing in IBM PureSystems and Storage. He has 21 years of experience in technical sales, starting with MVS, then eight years with AIX. For the past 12 years, he has worked in the storage and SAN product areas. He holds a degree as Graduate Engineer in Computer Science from the University of Dresden. He has co-authored eight IBM Redbooks publications. His current focus is working in the open storage area, storage virtualization, SAN, and PureSystems.

Peter Demharter is an IBM certified Senior Architect IT Infrastructure and Cisco Certified Internetwork Expert in Germany. He has over 20 years of experience in the data center and networking area and has worked for large companies, such as Daimler-Benz and Vodafone. He holds a degree in Administration and Information Science from the University of Constance. He has worked for IBM GTS for 10 years and has served as lead Architect in IBM projects such as ABB worldwide WAN migration from Equant to AT&T, and Deutschland Online Infrastructure, one of the first corporate IPv4/IPv6 dual stack wide area networks in Germany. He works for the IBM Research and Development Global Design Center in Boeblingen and focuses on IPv6, DC Networking, and Cloud Computing.

Stephan Fleck is a System Network Architect for IBM Systems & Technology Group, Europe. He has 19 years experience in the IT industry. His areas of expertise include network architecture assessments and network designs for data centers, and implementation proposals for network virtualization and network convergence solutions. Stephan also conducts training sessions for technical and sales personnel and he speaks regularly at technical conferences. He has worked as Network Security Lead Architect for the IBM Global Account and as support specialist for the European Network Support Back Office. Stephan is a Cisco Certified Internetwork Expert and holds a degree in electrical engineering from the Technical University Darmstadt, Germany.

Joachim Gross is an IT Architect and expert for network infrastructure in Germany. He has 20 years of experience in the networking area field as a Cisco Certified Internetwork Expert since 1995. He holds a degree in Information Technology from the FH in Esslingen, Germany. Working for IBM GTS for over 10 years, he has participated in worldwide networking and Voice over IP projects. His areas of expertise include data center networking and Voice over IP.

Ruediger Rissmann holds a Diploma Degree in Physics from the University of Heidelberg, Germany, and joined the IBM Zurich Research Laboratory in 1999. In his position as a network specialist, he has been involved in several pilot projects that explore new and emerging network technologies and has filed a number of patents. He leads the worldwide IPv6 deployment within the IBM Research Division. In March 2011, Ruediger became a research staff member and senior architect in the Services Innovation Lab. He holds the following certifications: IBM Certified IT Architect, Open Group Master Certified IT Architect, CCNP, CISSP, and GCFA.

Werner Sponer is a Senior IT Architect and expert for network infrastructure and security. He is responsible for network infrastructure and System Networking products in the System and Technology Group of IBM. He spent most of his 20-plus years at IBM growing the Global Services business through technical advancements. His assignments ranged from infrastructure to consulting and audit services, including projects and managed services. He brings over 18 years of IT experience in networking, data center, network architecture, local and wide area network, operation and support of IT infrastructure, in different customer industries and technologies. He evolves his leadership skills and customer orientation in different project scenarios in several countries, from consulting and planning, architecture, and design to operation and support. He is an engineer for electronic and biomedical technologies and IBM and Open Group Certified IT Architect.

Arwed Tschoeke is a Client Technical Architect in Hamburg, Germany. His focus areas are zEnterprise, virtualization solutions across IBM platforms, and Linux. He holds a degree in Physics from the University of Kaiserslautern, Germany.

Pietro Volante is a Certified IT Specialist for Networking Services. He has 20 years of experience in designing and implementing networks in many large client situations. He is certified as a Cisco Network and Design Professional (CCNP/CCDP) and has experience in designing data center networks and network performance analysis. In 2010, he worked on an assignment at STG to provide technical network support for the new BladeCenter network switches across north east Europe. He is responsible for projects in data center network integration and end-to-end network application performance analysis at key accounts.

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

- ▶ Erich Amrehn
- ▶ Bernhard Dierberger
- ▶ Oliver Raff
- ▶ Thomas Schwaller
- ▶ David Watts

Portions of this paper were based on the IBM Redbooks® publication, *Implementation of IBM j-type Ethernet Switches and Routers*, SG24-7882. Thanks to the authors of that paper.

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 this website:

<http://www.ibm.com/redbooks/residencies.html>

Comments welcome

Your comments are important to us!

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

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

<http://www.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

With the release of PureSystems™, IBM launched a second hybrid computing platform to the market. zEnterprise® with zBX is focused on mainframe affine applications with a simplified workload-oriented management approach. PureSystems offers various implementation possibilities that are focused on a cloud-oriented customer strategy.

To deliver value, PureSystems consists of the following building blocks:

- **Management**

The Flex System Manager simplifies and automates all management tasks. It also manages all physical and virtual resources within the solution. Hence, it offers a full integration and infrastructure-as-a-service-like management of PureSystems.

- **Compute Nodes**

To select the system that fits best to your requirements, it is possible to mix Power Systems and System x® compute nodes within the PureSystems Chassis.

- **Storage**

The Storwize V7000 storage controller delivers automatic EasyTierung of storage controller internal storage and the possibility to take advantage of external storage at the same time. With its built-in storage virtualization, simple and comprehensive management is possible via the integrated management of PureSystems.

- **Networking**

PureSystems provide a choice of adapters and switches. All components are standard-based and integrated into the management of PureSystems. This variety provides a combination of features that fits into the existing infrastructure. The modular concept offers the possibility to adapt to future requirements.

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

1.1 Networking

Many customers are currently migrating their networking infrastructure from 1 Gb Ethernet to 10 Gb Ethernet. This transformation exceeds the simple change of technology and requirements increased significantly. The complexity of modern application infrastructures requires networks of low latency at high bandwidth. Additionally, growing security awareness affects the design of a network and increases the complexity (for example, router, firewalls, filters). Because of virtualization and the adoption of cloud concepts, the physical network infrastructure merges with a logical and virtual networking environment that is represented by software components that are running on server systems.

As a result, there is no average network or general blueprint. Each network is unique because it depends on the customer's demands. Often, customers choose individual components from vendors that meet their requirements. From this decision, the following challenges arise:

- ▶ The administration of such mixed infrastructures is rather complex and often requires more management concepts.
- ▶ Testing and maintaining interoperability is elaborate and time-consuming.

To overcome these challenges, customers' adopt a single-vendor strategy. This strategy offers a simplification in the daily routine but can restrict the adoption of new solutions if they are not supported by the infrastructure vendor.

To support their customers, the industry defines standards. Based on those standards, interoperability between vendors can be achieved. This interoperability offers the opportunity to adapt the latest technology and limit the risk to the administration.

However, new standards are adopted by vendors at different times and not all choose to follow standards rigorously. Instead, they might provide their own extensions. One example of this issue is the integration of virtualized environments into the networking infrastructure. The networking branch of IBM is investing with other vendors a significant amount of energy to define global standards that support the mobility of virtual systems and infrastructures, such as vswitches. This effort delivers the availability of functions that allow a guest relocation between different systems that are independent of the hypervisor or the networking components within the physical infrastructure.

1.2 IBM PureSystems

The PureSystems platform is a new approach to deliver scalable hybrid systems for the adoption of modern cloud concepts. Its design delivers value to the customer by fulfilling the following requirements:

- ▶ Simplification to ease the implementation of complex solutions and operation
- ▶ Built in expertise to ease deployment and capacity planning
- ▶ Integration within the existing architectures and infrastructure

These advantages are achieved by a new hardware and system management concept. To reflect this concept, the systems are labeled *Expert Integrated Systems*. The following PureSystems offerings are available:

- ▶ PureFlex™ System: An infrastructure system that monitors capacity and performance to optimize the infrastructure (Infrastructure-as-a-Service within the cloud terminology).

- **PureApplication System:** A platform system that is based on a flexible infrastructure that provides the means of deploying and maintaining an application infrastructure that is based on patterns (Platform-as-a-Service within the cloud terminology).
- **PureData System:** Based on the PureApplication concept, this solution is focused on delivering data services by providing a fully managed, flexible, and highly available database platform that meets all demands.

The foundation of these Expert Integrated Systems is the PureSystems hardware, which consists out-of-server hardware (Power and x86), storage, and network, such as storage area network (SAN) and local area network (LAN). The design principle inherits the BladeCenter philosophy of IBM to open standards, manageability, serviceability, and an existing roadmap for investment protection.

To provide full flexibility, many active infrastructure components are available. The LAN components are derived from the networking technology of IBM, which ensures that an in-depth integration into virtual environments is possible. Because of the broad support of networking standards, this ability applies to physical networks as well.

For more information about IBM PureSystems, see *Overview of IBM PureSystems*, TIPS0892, which is available at this website:

<http://www.redbooks.ibm.com/abstracts/tips0892.html>

1.3 Switch configuration

IBM System Networking switches can be configured through multiple configuration interfaces. For this paper, the iSCLI method was chosen. Its syntax should be familiar to network administrators with experience in switches from other vendors.

Important: This Redpaper uses the show running-config configuration dumps to demonstrate how the switches were configured. These dumps include all of the command sequences that are required to configure the switch manually.

For more information, see the Configuration Dump section of the Configuration Commands chapter in *ISCLI—Industry Standard CLI Command Reference for the IBM Flex System Fabric EN4093 10Gb Scalable Switch*, which is available at this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.acc.networkdevices.doc%2Fio_module_compass.html

1.4 How to use this paper

We recommend that you read Chapter 2, “Layer 2 Network protocols and technologies” on page 5 first to clarify the use of technical terms. Then, based on the networking hardware you have, select the following appropriate chapter to read next:

- ▶ Chapter 3, “IBM RackSwitch G8264 connectivity” on page 15
- ▶ Chapter 4, “Cisco Nexus 5000 connectivity” on page 95
- ▶ Chapter 5, “Cisco Catalyst 6500 switch connectivity” on page 171

Within each of these chapters, you can review subsections that relate to the choice of Spanning Tree Protocol that you use.

Finally, Appendix A, “Troubleshooting” on page 177, describes different aspects of problem analysis and identifies information that is required for efficient troubleshooting.



Layer 2 Network protocols and technologies

Open systems interconnection (OSI) Layer 2 (or, the DataLink Layer) provides the functional means for data transfer between adjacent nodes in the network. Layer 2 also provides the lowest level of addressability in an Ethernet network that uses MAC addresses.

The MAC address contains 48 bits that are split into two, 24-bit sections. The first 24-bit section is assigned by IEEE to reflect the organizationally unique identifier (OUI)). Each Ethernet hardware manufacturer has one or more of these OUIs. The second 24-bit section is created by the manufacturer. The combination of these two 24-bit sections should guarantee that the MAC address is always unique in a LAN.

This chapter includes the following topics:

- ▶ Basic frame forwarding concepts
- ▶ Virtual local area network
- ▶ Spanning tree
- ▶ Link aggregation

2.1 Basic frame forwarding concepts

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 that is based on the destination MAC address.

Figure 2-1 shows the simplified principle of frame forwarding.

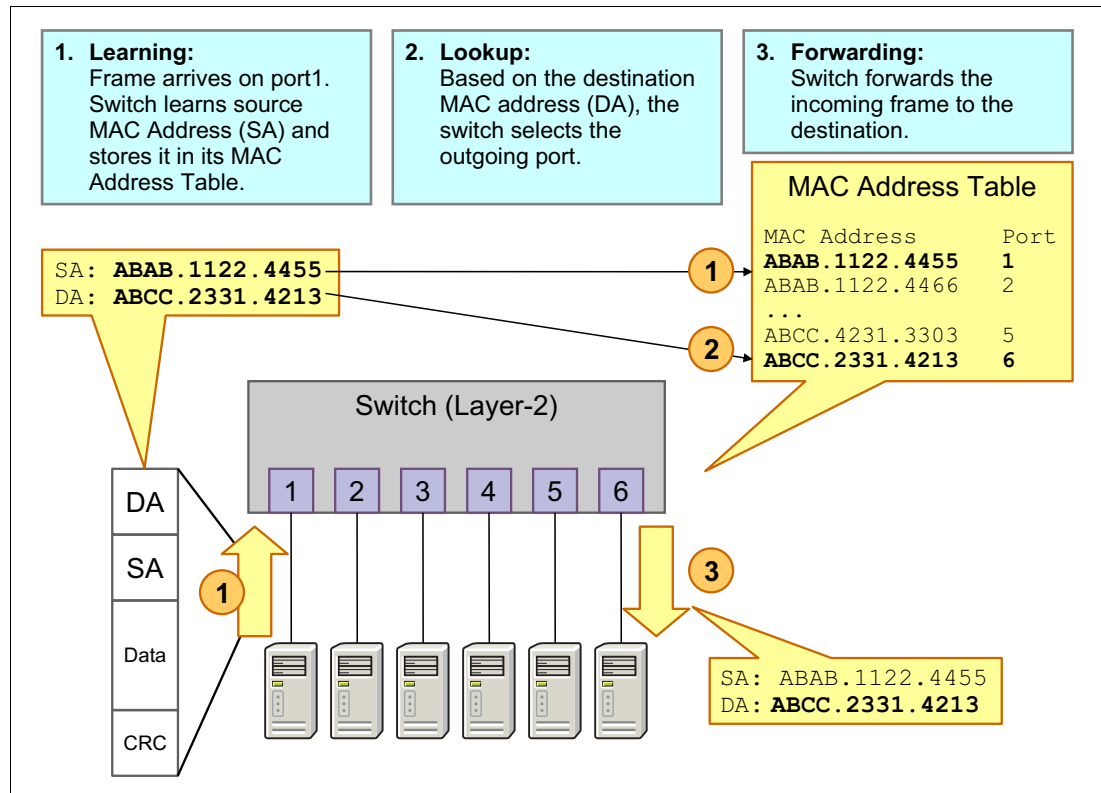


Figure 2-1 Simplified principle of frame forwarding

The forwarding of an 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 this fact in its MAC Address Table.

- Lookup

Based on the destination MAC address (DA), the switch looks up the correct routing in its MAC address table and selects the outgoing port (port 6).

- Forwarding

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

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

During this forwarding process, the frame header persists unmodified.

2.2 Virtual local area network

A virtual local area network (VLAN) is a networking concept in which a network is logically divided into smaller virtual LANs. The Layer 2 traffic in one VLAN is logically isolated from other VLANs, as shown in Figure 2-2.




Figure 2-2 Isolation at Layer 2

The simplest way to keep the isolated VLANs separately on an inter-switch link is to use one physical link for each VLAN, as shown in Figure 2-3.

However, 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.




Figure 2-3 Inter-switch link: one link for each VLAN

The second method is VLAN tagging over a single link in which each frame is tagged with its VLAN ID (see Figure 2-4 on page 8). This method is highly scalable because only a single link is required to provide connectivity to many VLANs. This configuration provides for better utilization of the link capacity when VLAN traffic is not uniform.

The protocol for VLAN tagging of frames in a LAN environment is defined by the IEEE 802.1 P/Q standard.




Figure 2-4 Inter-switch link that uses VLAN tagging

2.3 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 (TTL), which would allow for a looping packet to be detected and discarded. Packets that are 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 in which frames with a certain target can take only one path on each individual switch element. For Ethernet, the tree topology was chosen, which is the simplest topology that guarantees 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. This configuration has been further enhanced and is now used in the following forms:

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

STP was the initial implementation of Spanning-Tree Protocol, which was invented 1985 and published 1990 in the IEEE as 802.1D.

Rapid Spanning Tree (RSTP) became standard in IEEE in 2001 as 802.1w. It provides faster convergence times than STP.

Multiple Spanning Tree (MSTP) was first defined in IEEE as 802.1s and later merged into 802.1Q-2005 as an extension to RSTP. It uses more than one Spanning Tree process to distribute the VLANs into different STP topologies.

Cisco provides a proprietary version of VLAN-based STP. For each VLAN, it uses a separate Spanning Tree. Even if it is not an IEEE standard, many network vendors allow compatible setup to interoperate with Cisco's STP.

2.3.1 Spanning Tree Protocol: IEEE 802.1D

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

The following types of BPDUs are available:

- Configuration BPDU

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) BPDU

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 TCN for its own root port. The process continues until the TCN reaches the root bridge.

- Topology Change Notification Acknowledgement (TCA) BPDU

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 that are 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.

Figure 2-5 shows a switch topology with five interconnected switches. To avoid Layer 2-looped frames, Spanning Tree blocks all ports that include an indirect, redundant path to the root bridge. As shown in Figure 2-5, the resulting logical switch topology is based on the STP calculation.

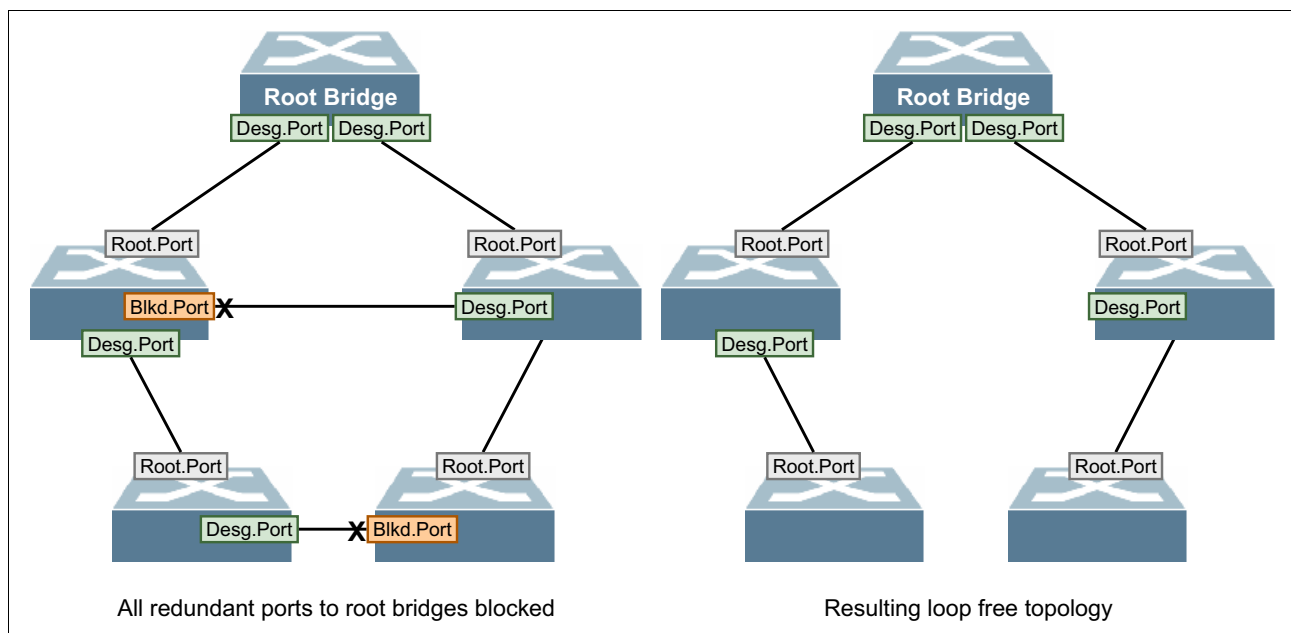


Figure 2-5 Switch topology with five interconnected switches

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.

2.3.2 Rapid Spanning Tree Protocol: IEEE 802.1w

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

An RSTP domain includes the following components:

- ▶ Root port: The “best path” to the root device.
- ▶ Designated port: Indicates that the switch is the designated bridge for the other switch that connects to this port.
- ▶ Alternative port: Provides an alternative root port.
- ▶ Backup port: Provides a designated alternative port. This configuration is used if there is more than one link that is connected to the same switch without link aggregation.

RSTP uses the following port states by using the show spanning tree command:

- ▶ Discarding: Like the blocking-state in STP, this port does not forward traffic to avoid loops.
- ▶ Learning: The port builds its MAC address table but does not forward traffic.
- ▶ Forwarding: The port forwards traffic.

The RSTP reconvergence time often is less than 1 second. The standard STP (802.1d) requires 30 seconds or more.

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

2.3.3 Multi-instance Spanning Tree Protocol: IEEE 802.1s

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

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 originally defined in the IEEE 802.1s draft specification and later incorporated into the IEEE 802.1Q-2003 specification.

2.3.4 Per VLAN Rapid Spanning Tree

Per VLAN Rapid Spanning Tree (PVRST) is a nonstandard spanning tree extension that is based on RSTP that was introduced by Cisco Systems. In PVRST mode, each VLAN is assigned to its own spanning-tree group. A maximum of 127 spanning tree groups are allowed in IBM System Networking switches.

PVRST use 802.1Q tagged frames to differentiate STP BPDUs for each VLAN. The IIBM System Networking implementation of PVRST is fully compatible with Cisco RSTP/PVRST+ protocol.

2.4 Link aggregation

A link aggregation group (LAG) combines physical links to operate as a single, larger logical link. The member links no longer function as independent physical connections, but as members of the larger logical link, as shown in Figure 2-6.




Figure 2-6 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 the devices that are 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 decide. This algorithm might include a source MAC, destination MAC, source IP, destination IP, and more. It might also include a combination of these values.

Link aggregation can be defined as static or by using a dynamic negotiation protocol, such as Link Aggregation Control Protocol (LACP). Aggregated links often are referred to as *Ether-Channels* or *Trunk-Links*.

Aggregated links appear to the STP as single logical links. Therefore, STP does not enable or disable individual physical links of an aggregated link.

2.4.1 Link Aggregation Control Protocol

LACP (also known as 802.3ad and, more recently, 802.1AX-2008) is a vendor-independent standard for dynamically building aggregated links between switches. On an LACP-defined link, the switches are sending LACP Data Units (LACPDU) to share information about the current state of the aggregated link. Compared to static LAG, LACP provides better failure detection and, therefore, a higher redundancy.

2.4.2 Virtual Link Aggregation Groups

Virtual Link Aggregation Groups (VLAGs) is an extension to link aggregation to allow more redundancy. For a standard LAG (static or dynamic), all ports that build an aggregated link must be on the same switch. VLAG allows two switches to pair 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 trunk that is connected to a single logical device.

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.

As shown in Figure 2-7, VLAG helps to avoid blocked ports by STP and allows higher performance and full redundancy.




Figure 2-7 Comparing STP with blocked ports versus VLAG loop-free topology

Important: The protocol for VLAG peer links is not standardized, so the switches in a pair of switches must belong to the same product family.

2.4.3 Cisco Virtual Port Channel

On the Nexus platform, Cisco implemented the VLAG concept as a version of a Multichassis EtherChannel (MEC), called the *Virtual Port Channel* (vPC), as shown in Figure 2-8. The vPC combines the advantages of hardware redundancy and the loop management of an aggregated link. The pair of switches that is building the vPC appear to any Portchannel-attached device as a single switch from Layer 2 perspective, while they are still operating as two independent devices with independent switch control and management.

If a vPC is used, the STP is not needed to manage the loops. Therefore, it could be disabled on these links and all disadvantages of the STP could be eliminated. The biggest advantage of this configuration is the usability of all bandwidth of the installed links and the fast handling of link failures within the vPC.




Figure 2-8 Schematic drawing of vPC

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

2.4.4 Link Layer Discovery Protocol: 802.1AB

The Link Layer Discovery Protocol (LLDP) is a vendor-neutral link-layer protocol that is used by network devices to enable standardized discovery of network nodes.

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



IBM RackSwitch G8264 connectivity

In this chapter, various network configuration scenarios for a PureSystem that is connected to an IBM Rack Switch infrastructure are described.

Configuration tests have been done for commonly used network technologies, VLAN trunking (IEEE 802.1Q), static and dynamic link aggregation (LACP), Spanning Tree (PVRST, MSTP), and network virtualization with VLAG (virtual Link Aggregation).

Link Layer Discovery Protocol (LLDP) as vendor independent protocol is used to verify Layer 2 topology.

In this chapter, we show the configuration dumps of the network devices and the commands that are used to verify the proper operation of the switches. We explain the configurations with use cases that show examples of how to configure the devices for this setup.

This chapter includes the following topics:

- ▶ Prerequisites
- ▶ Use Case 1: PVRST
- ▶ Use Case 2: Link aggregation and PVRST
- ▶ Use Case 3: Link aggregation and MST
- ▶ Use Case 4: Link aggregation, MSTP and VLAG
- ▶ Use Case 5: Link aggregation and VLAG without STP

3.1 Prerequisites

We started by physically connecting a triangle with two IBM RackSwitch™ G8264 switches and one IBM Flex System™ EN2092 1 Gb switch. We configured four VLANs and set up Per VLAN Rapid Spanning Tree (PVSTP). To test connectivity, we used a test PC.

We used the following switches and one PC to test connectivity:

- ▶ Two IBM RackSwitch G8264 switches
- ▶ One IBM Flex System EN2092 1 Gb Ethernet Scalable Switch
- ▶ One test PC

The links between the switches always are 10 Gigabit Ethernet.

3.2 Use Case 1: PVRST

In Use Case 1, we have a pair of IBM RackSwitch G8264 switches connected to Flex System EN2092 1-Gb Ethernet Scalable Switch with PVRST.

In this use case, we used three 10 GE links to connect the switches. We also configured 802.1q trunks and PVRST. For load balancing, odd VLANs 10 and 30 and even VLANs 20 and 40 are used, as shown in Figure 3-1 (odd VLANs) and Figure 3-2 on page 17 (even VLANs)

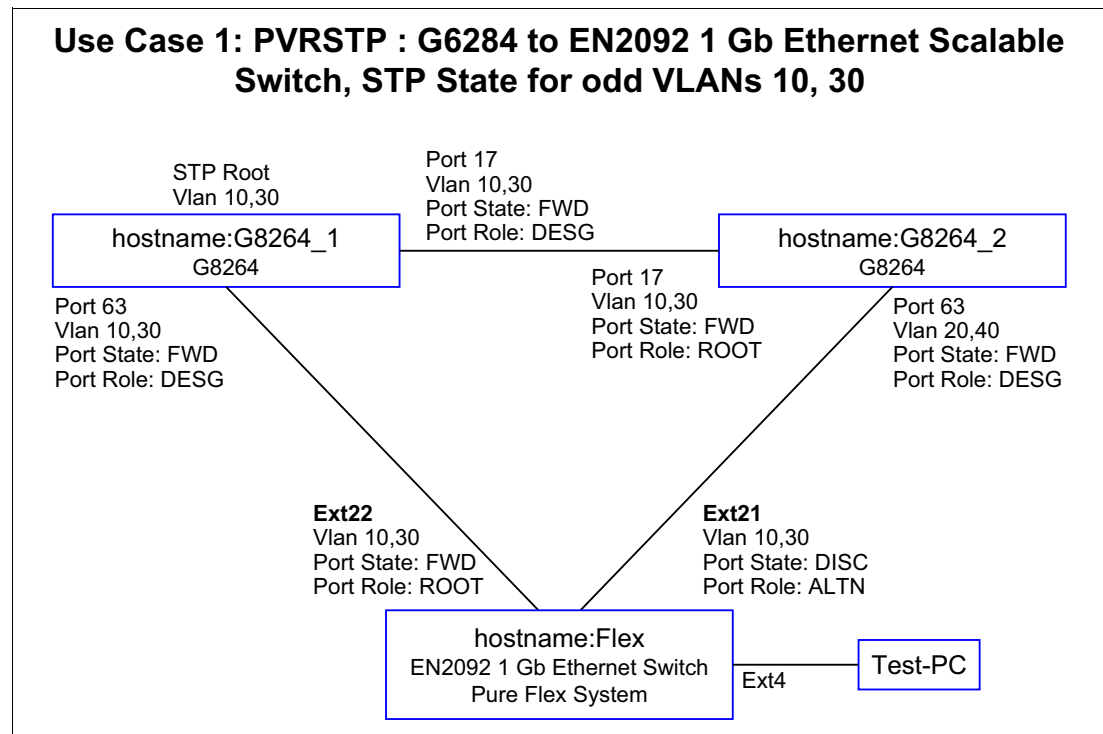


Figure 3-1 Use Case 1: PVRST: Odd-numbered VLANs

Use Case 1: PVRSTP : G6284 to EN2092 1 Gb Ethernet Scalable Switch, STP State for even VLANs 20, 40

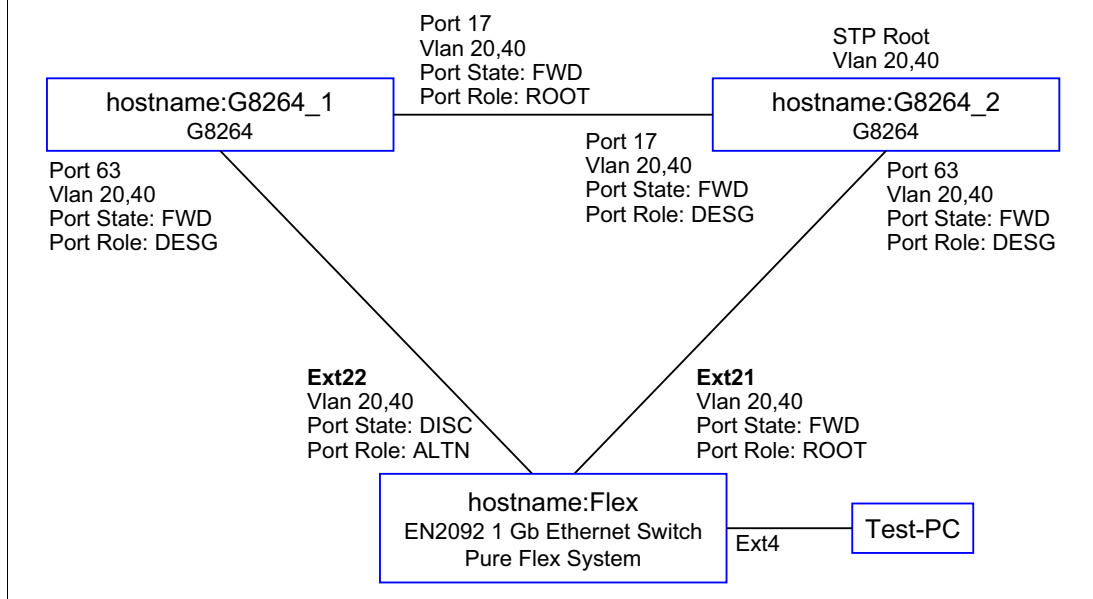


Figure 3-2 Use Case 1: PVRST: Even-numbered VLANs

3.2.1 Verifying the topology by using lldp

To verify the topology, we used the **lldp remote-device** command on the three switches, as shown in Example 3-1. Important parameters and details are highlighted in red.

Example 3-1 Checking the topology use show lldp remote-device

```
Flex#sh lldp remote-device
```

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
EXT22	1	08 17 f4 32 bb 00	63	G8264_1
EXT21	2	fc cf 62 9d 67 00	63	G8264_2
INTA1	3	5c f3 fc 5f 43 9d	5c-f3-fc-5f-43-9d	
EXT5	6	00 0d ec a3 8f 81	mgmt0	vie
EXT7	7	00 05 9b 7b 84 01	mgmt0	str

!--- Display the LLDP remote devices.

!--- The local Port Numbers of the Pure Flex System Ethernet Switch

!--- distinguish between internal and external Ethernet ports.

!--- The EXT4 port connecting to the Test PC is not shown as this device does not support LLDP .

```
G8264_1#show lldp remote-device
```

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
17	1	fc cf 62 9d 67 00	17	G8264_2
63	2	08 17 f4 76 78 00	50	Flex

!--- The port EXT22 of the Flex switch is mapped to remote port number 50.

```
G8264_2#show lldp remote-device
```

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
17	1	08 17 f4 32 bb 00	17	G8264_1
63	2	08 17 f4 76 78 00	49	Flex

!--- The port EXT21 of the Flex switch is mapped to remote port number 49.

3.2.2 Verifying trunks

To verify which VLANs are active on which trunk, we used the **show interface trunk** command on the switches, as shown in Example 3-2. Important parameters and details are highlighted in red.

Example 3-2 Output from the show interface trunk command

```
Flex#show interface trunk
Alias  Port Tag RMON Lrn Fld PVID      NAME                               VLAN(s)
-----
...
EXT4   32   y   d   e   e   1   TEST_PC                               1 10 20 30 40
...
EXT21  49   y   d   e   e   10  T0_G8264_2_Port63 10 20 30 40
EXT22  50   y   d   e   e   10  T0_G8264_1_Port63 10 20 30 40
EXT23  51   y   d   e   e   10  T0_G8264_1_Port64 10 20 30 40
EXT24  52   y   d   e   e   10  T0_G8264_2_Port64 10 20 30 40
```

* = PVID is tagged.

```
G8264_2#sh int trunk
Alias  Port Tag RMON Lrn Fld PVID      NAME                               VLAN(s)
-----
17     17   y   d   e   e   10  CrossLink           10 20 30 40
18     18   y   d   e   e   10  CrossLink           10 20 30 40
63     63   y   d   e   e   10  UPLINK_TO_FLEX     10 20 30 40
64     64   y   d   e   e   10  UPLINK_TO_FLEX     10 20 30 40
```

* = PVID is tagged.

3.2.3 Verifying PVRST spanning tree configurations

We verified the PVRST spanning tree configuration of the switches by executing the **show spanning-tree** command, which produced the following outputs. Important parameters and details are highlighted in red:

- ▶ EN2029: Example 3-3 on page 20
- ▶ G8264 switch 1: Example 3-4 on page 22
- ▶ G8264 switch 2: Example 3-5 on page 24

As shown in Figure 3-1 on page 16, we have two spanning trees, one for even-numbered VLANs and one for odd-numbered VLANs. By using the **show spanning-tree** command, you can verify the status of the respective Ethernet interface's VLAN, port state, and port role.

Example 3-3 Verifying the PVRST spanning tree configuration: EN2092 switch

```
Flex#sh spanning-tree
```

```
-----  
Pvst+ compatibility mode enabled  
-----
```

```
Spanning Tree Group 1: On (PVRST)
```

```
VLANs: 1
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel  
8000 00:16:ca:a1:c1:00    20000    EXT3    2    20    15
```

```
Parameters:  Priority  Hello  MaxAge  FwdDel  Aging  Topology Change Counts  
             61441    2      20      15     300           3
```

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
INTA1	0	0	FWD *				
INTA2	0	0	FWD *				
INTA4	0	0	FWD *				
EXT1	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00	801d	P2P
EXT2	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00	801e	P2P
EXT3	128	20000!	FWD	ROOT	8000-00:16:ca:a1:c1:00	8011	P2P
EXT4	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00	8020	P2P
EXT5	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00	8021	P2P
EXT7	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00	8023	P2P

* = STP turned off for this port.

! = Automatic path cost.

```
-----  
Spanning Tree Group 10: On (PVRST)
```

```
VLANs: 10
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel  
600a 08:17:f4:32:bb:00    2000    EXT22    2    20    15
```

!--- Compare the ID of the Root with the LLDP output to identify the root switch.

```
Parameters:  Priority  Hello  MaxAge  FwdDel  Aging  Topology Change Counts  
             61450    2      20      15     300           4
```

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	FWD	DESG	f00a-08:17:f4:76:78:00	8020	P2P
EXT21	128	2000!	DISC	ALTN	700a-fc:cf:62:9d:67:00	803f	Shared
EXT22	128	2000!	FWD	ROOT	600a-08:17:f4:32:bb:00	803f	Shared

! = Automatic path cost.

```
-----  
Spanning Tree Group 20: On (PVRST)
```

```
VLANs: 20
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel  
6014 fc:cf:62:9d:67:00    2000    EXT21    2    20    15
```

```
Parameters:  Priority  Hello  MaxAge  FwdDel  Aging  Topology Change Counts  
             61460    2      20      15     300           3
```

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
------	------	------	-------	------	-------------------	----------	------

```

-----
EXT4          128      20000! FWD  DESG f014-08:17:f4:76:78:00      8020      P2P
EXT21         128      2000! FWD  ROOT 6014-fc:cf:62:9d:67:00      803f      Shared
EXT22         128      2000! DISC ALTN 7014-08:17:f4:32:bb:00      803f      Shared
! = Automatic path cost.

```

```

-----
Spanning Tree Group 30: On (PVRST)
VLANs: 30

```

```

Current Root:          Path-Cost  Port Hello MaxAge FwdDel
601e 08:17:f4:32:bb:00      2000  EXT22   2    20    15

```

```

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
              61470     2     20     15    300              4

```

```

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
-----
EXT4          128      20000! FWD  DESG f01e-08:17:f4:76:78:00      8020      P2P
EXT21         128      2000! DISC ALTN 701e-fc:cf:62:9d:67:00      803f      Shared
EXT22         128      2000! FWD  ROOT 601e-08:17:f4:32:bb:00      803f      Shared
! = Automatic path cost.

```

```

-----
Spanning Tree Group 40: On (PVRST)
VLANs: 40

```

```

Current Root:          Path-Cost  Port Hello MaxAge FwdDel
6028 fc:cf:62:9d:67:00      2000  EXT21   2    20    15

```

```

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
              61480     2     20     15    300              3

```

```

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
-----
EXT4          128      20000! FWD  DESG f028-08:17:f4:76:78:00      8020      P2P
EXT21         128      2000! FWD  ROOT 6028-fc:cf:62:9d:67:00      803f      Shared
EXT22         128      2000! DISC ALTN 7028-08:17:f4:32:bb:00      803f      Shared
! = Automatic path cost.

```

```

-----
Spanning Tree Group 128: Off (PVRST), FDB aging timer 300
VLANs: 4095

```

```

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
-----
MGT1           0         0    FWD  *
* = STP turned off for this port.

```

Example 3-4 Verifying the PVRST spanning tree configuration: G8264 switch 1

```
G8264_1#sh spanning-tree
```

```
-----  
Pvst+ compatibility mode enabled  
  
-----
```

```
Spanning Tree Group 1: On (PVRST)
```

```
VLANs: 1
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel  
8001 08:17:f4:32:bb:00      0      0   2    20    15
```

```
Parameters:  Priority  Hello  MaxAge  FwdDel  Aging  Topology Change Counts  
              32769    2      20      15     300           7
```

```
      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type  
-----
```

Note: There is no active STP port in Spanning Tree Group 1.

```
-----  
Spanning Tree Group 10: On (PVRST)
```

```
VLANs: 10
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel  
600a 08:17:f4:32:bb:00      0      0   2    20    15
```

```
Parameters:  Priority  Hello  MaxAge  FwdDel  Aging  Topology Change Counts  
              24586    2      20      15     300           3
```

```
      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type  
-----  
17              128     2000!  FWD    DESG 600a-08:17:f4:32:bb:00  8011        P2P  
63              128     2000!  FWD    DESG 600a-08:17:f4:32:bb:00  803f        P2P
```

! = Automatic path cost.

```
-----  
Spanning Tree Group 20: On (PVRST)
```

```
VLANs: 20
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel  
6014 fc:cf:62:9d:67:00    2000    17   2    20    15
```

```
Parameters:  Priority  Hello  MaxAge  FwdDel  Aging  Topology Change Counts  
              28692    2      20      15     300           2
```

```
      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type  
-----  
17              128     2000!  FWD    ROOT 6014-fc:cf:62:9d:67:00  8011        P2P  
63              128     2000!  FWD    DESG 7014-08:17:f4:32:bb:00  803f        P2P
```

! = Automatic path cost.

```
-----  
Spanning Tree Group 30: On (PVRST)
```

```
VLANs: 30
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel  
601e 08:17:f4:32:bb:00      0      0   2    20    15
```

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
 24606 2 20 15 300 3

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17	128	2000!	FWD	DESG	601e-08:17:f4:32:bb:00	8011	P2P
63	128	2000!	FWD	DESG	601e-08:17:f4:32:bb:00	803f	P2P

! = Automatic path cost.

 Spanning Tree Group 40: On (PVRST)
 VLANs: 40

Current Root: Path-Cost Port Hello MaxAge FwdDel
 6028 fc:cf:62:9d:67:00 2000 17 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
 28712 2 20 15 300 2

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17	128	2000!	FWD	ROOT	6028-fc:cf:62:9d:67:00	8011	P2P
63	128	2000!	FWD	DESG	7028-08:17:f4:32:bb:00	803f	P2P

! = Automatic path cost.

 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.

Example 3-5 Verifying the PVRST spanning tree configuration: G8264 switch 2

```
G8264_2#sh spanning-tree
```

```
-----  
Pvst+ compatibility mode enabled  
  
-----
```

```
Spanning Tree Group 1: On (PVRST)
```

```
VLANs: 1
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel  
8001 fc:cf:62:9d:67:00      0      0   2    20    15
```

```
Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts  
32769      2    20    15    300          0
```

```
Port      Prio    Cost    State  Role Designated Bridge      Des Port  Type  
-----
```

Note: There is no active STP port in Spanning Tree Group 1.

```
-----  
Spanning Tree Group 10: On (PVRST)
```

```
VLANs: 10
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel  
600a 08:17:f4:32:bb:00    2000    17   2    20    15
```

```
Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts  
28682      2    20    15    300          3
```

```
Port      Prio    Cost    State  Role Designated Bridge      Des Port  Type  
-----  
17          128    2000!  FWD    ROOT 600a-08:17:f4:32:bb:00    8011      P2P  
63          128    2000!  FWD    DESG 700a-fc:cf:62:9d:67:00    803f      P2P  
! = Automatic path cost.
```

```
-----  
Spanning Tree Group 20: On (PVRST)
```

```
VLANs: 20
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel  
6014 fc:cf:62:9d:67:00      0      0   2    20    15
```

```
Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts  
24596      2    20    15    300          2
```

```
Port      Prio    Cost    State  Role Designated Bridge      Des Port  Type  
-----  
17          128    2000!  FWD    DESG 6014-fc:cf:62:9d:67:00    8011      P2P  
63          128    2000!  FWD    DESG 6014-fc:cf:62:9d:67:00    803f      P2P  
! = Automatic path cost.
```

```
-----  
Spanning Tree Group 30: On (PVRST)
```

```
VLANs: 30
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel  
601e 08:17:f4:32:bb:00    2000    17   2    20    15
```



```

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
            28702      2      20      15      300                      3

      Port      Prio      Cost      State Role Designated Bridge      Des Port      Type
-----
17              128      2000! FWD  ROOT 601e-08:17:f4:32:bb:00      8011          P2P
63              128      2000! FWD  DESG 701e-fc:cf:62:9d:67:00      803f          P2P
! = Automatic path cost.

-----

Spanning Tree Group 40: On (PVRST)
VLANs: 40

Current Root:          Path-Cost Port Hello MaxAge FwdDel
6028 fc:cf:62:9d:67:00      0      0      2      20      15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
            24616      2      20      15      300                      2

      Port      Prio      Cost      State Role Designated Bridge      Des Port      Type
-----
17              128      2000! FWD  DESG 6028-fc:cf:62:9d:67:00      8011          P2P
63              128      2000! FWD  DESG 6028-fc:cf:62:9d:67:00      803f          P2P
! = Automatic path cost.

-----

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.

```

3.2.4 Show running-config of all switches in Use Case 1

In the configuration output of the IBM Flex Switch and the IBM rack switches that are shown in the following examples, you can see the necessary configuration steps we did during our test. Important parameters and details are highlighted in red:

- EN2029: Example 3-6 on page 26
- G8264 switch 1: Example 3-7 on page 27
- G8264 switch 2: Example 3-8 on page 29

Example 3-6 Output from show running: EN2092 switch

```
Flex#sh run
Current configuration:
!
version "7.2.2.2"
switch-type "IBM Flex System EN2092 1Gb Ethernet Scalable Switch"
!
...
!
hostname "Flex"
...
!
interface port INTA2
    tagging
    exit
!
interface port INTA7
    shutdown
    exit
!
interface port EXT4
    name "TEST_PC"
    tagging
    exit
!
interface port EXT21
    name "T0_G8264_2_Port63"
    tagging
    pvid 10
    exit
!
interface port EXT22
    name "T0_G8264_1_Port63"
    tagging
    pvid 10
    exit
!
interface port EXT23
    name "T0_G8264_1_Port64"
    shutdown
    tagging
    pvid 10
    exit
!
interface port EXT24
    name "T0_G8264_2_Port64"
    shutdown
    tagging
    pvid 10
    exit
!
vlan 1
    member INTA1-EXT20
    no member EXT21-EXT24
!
!
vlan 10
    enable
    name "Server"
    member EXT4,EXT21-EXT24
```

```

!
!
vlan 20
    enable
    name "Data20"
    member EXT4,EXT21-EXT24
!
!
vlan 30
    enable
    name "Data30"
    member EXT4,EXT21-EXT24
!
!
vlan 40
    enable
    name "Data40"
    member EXT4,EXT21-EXT24
!
!
!
spanning-tree stp 10 vlan 10

spanning-tree stp 20 vlan 20

spanning-tree stp 30 vlan 30

spanning-tree stp 40 vlan 40

!
lldp enable
!
...
!
end

```

Example 3-7 Output from show running command: 8264 switch 1

```

G8264_1#sh run
Current configuration:
!
version "7.2.2"
switch-type "IBM Networking Operating System RackSwitch G8264"
!
!
!
!
no system dhcp
hostname "G8264_1"
system idle 60
!
!
interface port 17
    name "CrossLink"
    tagging
    pvid 10
    exit
!

```

```

interface port 18
    shutdown
    tagging
    pvid 10
    exit
!
interface port 63
    tagging
    pvid 10
    exit
!
interface port 64
    shutdown
    tagging
    pvid 10
    exit
!
vlan 1
    member 1-16,19-62
    no member 17-18,63-64
!
!
vlan 10
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 20
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 30
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 40
    enable
    name "VLAN 40"
    member 17-18,63-64
!
!
!
spanning-tree stp 10 bridge priority 24576
spanning-tree stp 10 vlan 10

spanning-tree stp 20 bridge priority 28672
spanning-tree stp 20 vlan 20

spanning-tree stp 30 bridge priority 24576
spanning-tree stp 30 vlan 30

spanning-tree stp 40 bridge priority 28672
spanning-tree stp 40 vlan 40
!

```

```
!  
lldp enable  
!  
...  
!  
end
```

Example 3-8 Output from show running command: G8264 switch 2

```
G8264_2#sh run  
Current configuration:  
!  
version "7.2.2"  
switch-type "IBM Networking Operating System RackSwitch G8264"  
!  
!  
  
!  
!  
no system dhcp  
hostname "G8264_2"  
system idle 60  
!  
!  
interface port 17  
    name "CrossLink"  
    tagging  
    pvid 10  
    exit  
  
!  
interface port 18  
    shutdown  
    tagging  
    pvid 10  
    exit  
  
!  
interface port 63  
    tagging  
    pvid 10  
    exit  
  
!  
interface port 64  
    shutdown  
    tagging  
    pvid 10  
    exit  
  
!  
vlan 1  
    member 1-16,19-62  
    no member 17-18,63-64  
  
!  
!  
vlan 10  
    enable  
    name "none"  
    member 17-18,63-64  
  
!  
!
```

```

vlan 20
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 30
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 40
    enable
    name "VLAN 40"
    member 17-18,63-64
!
!
!
spanning-tree stp 10 bridge priority 28672
spanning-tree stp 10 vlan 10

spanning-tree stp 20 bridge priority 24576
spanning-tree stp 20 vlan 20

spanning-tree stp 30 bridge priority 28672
spanning-tree stp 30 vlan 30

spanning-tree stp 40 bridge priority 24576
spanning-tree stp 40 vlan 40

!
!
lldp enable
!
...
!
!
end

```

3.3 Use Case 2: Link aggregation and PVRST

In our second use case, we added aggregation links and used three pairs of 10 GE links to connect the switches. We also configured 802.1q trunks with LACP and PVRST. For load balancing, odd VLANS 10 and 30 and even VLANS 20 and 40 were used (see Figure 3-3 on page 31 and Figure 3-4 on page 31).

Use Case 2: PVRSTP : G6284 to EN2092 1 Gb Ethernet Scalable Switch with LACP, STP State for even VLANs 20, 40

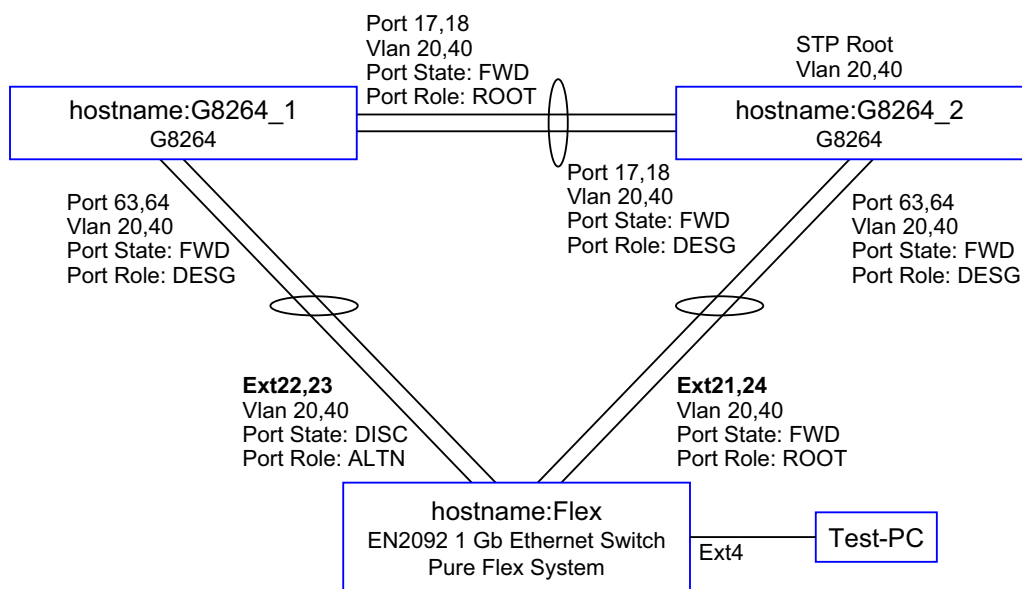


Figure 3-3 Use Case 2: Even-numbered VLANs

Use Case 2: PVRSTP : G6284 to EN2092 1 Gb Ethernet Scalable Switch with LACP, STP State for odd VLANs 10, 30

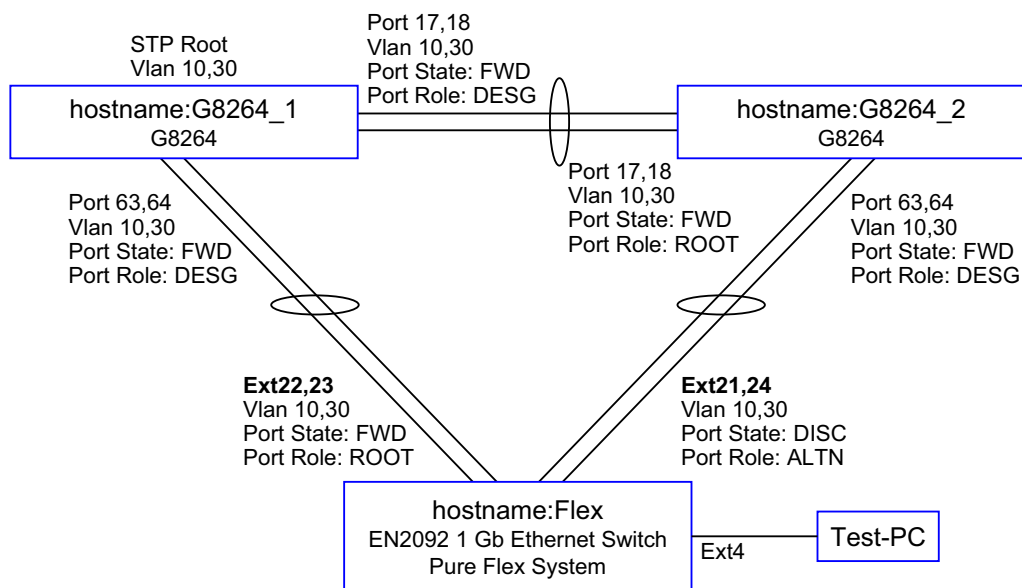


Figure 3-4 Use Case 2: Odd-numbered VLANs

3.3.1 Verifying the topology that is used by using lldp

To verify the topology, we used the **show lldp remote-device** command on the three switches, as shown in Example 3-9.

Example 3-9 Checking the topology use show lldp remote-device command

Flex#**show lldp remote-device**

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
-----	-----	-----	-----	-----
EXT22	1	08 17 f4 32 bb 00	63	G8264_1
EXT21	2	fc cf 62 9d 67 00	63	G8264_2
EXT23	5	08 17 f4 32 bb 00	64	G8264_1
EXT24	8	fc cf 62 9d 67 00	64	G8264_2

G8264_1#**sh lldp remote-device**

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
-----	-----	-----	-----	-----
17	1	fc cf 62 9d 67 00	17	G8264_2
63	2	08 17 f4 76 78 00	50	Flex
18	3	fc cf 62 9d 67 00	18	G8264_2
64	4	08 17 f4 76 78 00	51	Flex

G8264_2#**show lldp remote-device**

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
-----	-----	-----	-----	-----
17	1	08 17 f4 32 bb 00	17	G8264_1
63	2	08 17 f4 76 78 00	49	Flex
18	3	08 17 f4 32 bb 00	18	G8264_1
64	4	08 17 f4 76 78 00	52	Flex

3.3.2 Verifying trunks

To verify which VLANs are active on which trunk, we used the **show interface trunk** command on the three switches, as shown in Example 3-10.

Example 3-10 Show interface trunk command

```
Flex#show interface trunk
Alias  Port Tag RMON Lrn Flt PVID      NAME                               VLAN(s)
-----
EXT21  49   y   d   e   e   10   TO_G8264_2_Port63 10 20 30 40
EXT22  50   y   d   e   e   10   TO_G8264_1_Port63 10 20 30 40
EXT23  51   y   d   e   e   10   TO_G8264_1_Port64 10 20 30 40
EXT24  52   y   d   e   e   10   TO_G8264_2_Port64 10 20 30 40
* = PVID is tagged.
```



```
G8264_1#sh int trunk
Alias  Port Tag RMON Lrn Flt PVID      NAME                               VLAN(s)
-----
17     17   y   d   e   e   10   CrossLink           10 20 30 40
18     18   y   d   e   e   10   CrossLink           10 20 30 40
63     63   y   d   e   e   10   UPLINK_TO_FLEX     10 20 30 40
64     64   y   d   e   e   10   UPLINK_TO_FLEX     10 20 30 40
...
* = PVID is tagged.
```



```
G8264_2#sh int trunk
Alias  Port Tag RMON Lrn Flt PVID      NAME                               VLAN(s)
-----
17     17   y   d   e   e   10   CrossLink           10 20 30 40
18     18   y   d   e   e   10   CrossLink           10 20 30 40
63     63   y   d   e   e   10   UPLINK_TO_FLEX     10 20 30 40
64     64   y   d   e   e   10   UPLINK_TO_FLEX     10 20 30 40
* = PVID is tagged.
```

3.3.3 Verifying link aggregation by using lacp

We verified the link aggregation configuration of the three switches by executing the **show lacp information** command, as shown in Example 3-11.

Example 3-11 Show lacp information command

```
Flex#sh lacp information
```

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
EXT21	active	121	121	yes	32768	49	53	up	1
EXT22	active	122	122	yes	32768	50	54	up	1
EXT23	active	122	122	yes	32768	50	54	up	1
EXT24	active	121	121	yes	32768	49	53	up	1

!--- The “aggr” and “trunk” column identifies the ports which are configured together as link aggregation, i.e.trunk 53 is made of EXT21 and EXT24 .

```
G8264_1(config)#sh lacp information
```

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
17	active	117	117	yes	32768	17	65	up	1
18	active	117	117	yes	32768	17	65	up	1
63	active	163	163	yes	32768	63	66	up	1
64	active	163	163	yes	32768	63	66	up	1

```
G8264_2#sh lacp information
```

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
17	active	117	117	yes	32768	17	65	up	1
18	active	117	117	yes	32768	17	65	up	1
63	active	163	163	yes	32768	63	66	up	1
64	active	163	163	yes	32768	63	66	up	1

3.3.4 Verifying PVRST spanning tree configuration

In the next step, we verified the PVRST spanning tree configuration of the switches by executing the **show spanning-tree** command. As shown in Figure 3-3 on page 31 and Figure 3-4 on page 31, we have two spanning trees, one for even VLANs and one for odd VLANs. By using the show spanning tree command, you can verify the status of the respective Ethernet interface's VLAN, port state, and port role.

The commands that were run on the three switches produced the following outputs:

- ▶ EN2029: Example 3-12 on page 35
- ▶ G8264 switch 1: Example 3-13 on page 37
- ▶ G8264 switch 2: Example 3-14 on page 39

Example 3-12 Output from show spanning tree command: Flex System switch

Flex#show spanning-tree

Pvst+ compatibility mode enabled

Spanning Tree Group 1: On (PVRST)

VLANs: 1

Current Root: Path-Cost Port Hello MaxAge FwdDel
8000 00:16:ca:a1:c1:00 20000 EXT3 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
 61441 2 20 15 300 3

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
INTA1	0	0	FWD *				
INTA2	0	0	FWD *				
INTA4	0	0	FWD *				
EXT1	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00	801d	P2P
EXT2	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00	801e	P2P
EXT3	128	20000!	FWD	ROOT	8000-00:16:ca:a1:c1:00	8011	P2P
EXT4	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00	8020	P2P
EXT5	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00	8021	P2P
EXT7	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00	8023	P2P

* = STP turned off for this port.

! = Automatic path cost.

Spanning Tree Group 10: On (PVRST)

VLANs: 10

Current Root: Path-Cost Port Hello MaxAge FwdDel
600a 08:17:f4:32:bb:00 990 EXT22 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
 61450 2 20 15 300 8

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	FWD	DESG	f00a-08:17:f4:76:78:00	8020	P2P
EXT21 (pc53)	128	990!+	DISC	ALTN	700a-fc:cf:62:9d:67:00	8083	P2P
EXT22 (pc54)	128	990!+	FWD	ROOT	600a-08:17:f4:32:bb:00	8083	P2P
EXT23 (pc54)	128	990!+	FWD	ROOT	600a-08:17:f4:32:bb:00	8083	P2P
EXT24 (pc53)	128	990!+	DISC	ALTN	700a-fc:cf:62:9d:67:00	8083	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

!--- Please note the portchannel identifier after the port number, i.e. pc53, pc54

Spanning Tree Group 20: On (PVRST)

VLANs: 20

Current Root: Path-Cost Port Hello MaxAge FwdDel
6014 fc:cf:62:9d:67:00 990 EXT21 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
 61460 2 20 15 300 10

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	FWD	DESG	f014-08:17:f4:76:78:00	8020	P2P
EXT21 (pc53)	128	990!+	FWD	ROOT	6014-fc:cf:62:9d:67:00	8083	P2P
EXT22 (pc54)	128	990!+	DISC	ALTN	7014-08:17:f4:32:bb:00	8083	P2P
EXT23 (pc54)	128	990!+	DISC	ALTN	7014-08:17:f4:32:bb:00	8083	P2P
EXT24 (pc53)	128	990!+	FWD	ROOT	6014-fc:cf:62:9d:67:00	8083	P2P

! = Automatic path cost.
+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 30: On (PVRST)
VLANs: 30

Current Root: Path-Cost Port Hello MaxAge FwdDel
601e 08:17:f4:32:bb:00 990 EXT22 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
61470 2 20 15 300 8

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	FWD	DESG	f01e-08:17:f4:76:78:00	8020	P2P
EXT21 (pc53)	128	990!+	DISC	ALTN	701e-fc:cf:62:9d:67:00	8083	P2P
EXT22 (pc54)	128	990!+	FWD	ROOT	601e-08:17:f4:32:bb:00	8083	P2P
EXT23 (pc54)	128	990!+	FWD	ROOT	601e-08:17:f4:32:bb:00	8083	P2P
EXT24 (pc53)	128	990!+	DISC	ALTN	701e-fc:cf:62:9d:67:00	8083	P2P

! = Automatic path cost.
+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 40: On (PVRST)
VLANs: 40

Current Root: Path-Cost Port Hello MaxAge FwdDel
6028 fc:cf:62:9d:67:00 990 EXT21 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
61480 2 20 15 300 10

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	FWD	DESG	f028-08:17:f4:76:78:00	8020	P2P
EXT21 (pc53)	128	990!+	FWD	ROOT	6028-fc:cf:62:9d:67:00	8083	P2P
EXT22 (pc54)	128	990!+	DISC	ALTN	7028-08:17:f4:32:bb:00	8083	P2P
EXT23 (pc54)	128	990!+	DISC	ALTN	7028-08:17:f4:32:bb:00	8083	P2P
EXT24 (pc53)	128	990!+	FWD	ROOT	6028-fc:cf:62:9d:67:00	8083	P2P

! = Automatic path cost.
+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 128: Off (PVRST), FDB aging timer 300
VLANs: 4095

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
MGT1	0	0	FWD *				

* = STP turned off for this port.

Example 3-13 Output from show spanning tree command: G8264 switch 1

G8264_1(config)#sh spanning-tree

Pvst+ compatibility mode enabled

Spanning Tree Group 1: On (PVRST)
VLANs: 1

Current Root: Path-Cost Port Hello MaxAge FwdDel
8001 08:17:f4:32:bb:00 0 0 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
32769 2 20 15 300 7

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
------	------	------	-------	------	-------------------	----------	------

!--- Note: There is no active STP port in Spanning Tree Group 1.

Spanning Tree Group 10: On (PVRST)
VLANs: 10

Current Root: Path-Cost Port Hello MaxAge FwdDel
600a 08:17:f4:32:bb:00 0 0 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
24586 2 20 15 300 7

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17 (pc65)	128	990!	FWD	DESG	600a-08:17:f4:32:bb:00	8082	P2P
18 (pc65)	128	990!	FWD	DESG	600a-08:17:f4:32:bb:00	8082	P2P
63 (pc66)	128	990!	FWD	DESG	600a-08:17:f4:32:bb:00	8083	P2P
64 (pc66)	128	990!	FWD	DESG	600a-08:17:f4:32:bb:00	8083	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 20: On (PVRST)
VLANs: 20

Current Root: Path-Cost Port Hello MaxAge FwdDel
6014 fc:cf:62:9d:67:00 990 17 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
28692 2 20 15 300 9

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17 (pc65)	128	990!	FWD	ROOT	6014-fc:cf:62:9d:67:00	8082	P2P
18 (pc65)	128	990!	FWD	ROOT	6014-fc:cf:62:9d:67:00	8082	P2P
63 (pc66)	128	990!	FWD	DESG	7014-08:17:f4:32:bb:00	8083	P2P
64 (pc66)	128	990!	FWD	DESG	7014-08:17:f4:32:bb:00	8083	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

```

-----
Spanning Tree Group 30: On (PVRST)
VLANs: 30

Current Root:          Path-Cost  Port Hello MaxAge FwdDel
601e 08:17:f4:32:bb:00      0      0    2    20    15

Parameters:  Priority  Hello  MaxAge  FwdDel  Aging  Topology Change Counts
              24606      2      20      15      300          7

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
-----
17  (pc65)    128      990!+ FWD  DESG 601e-08:17:f4:32:bb:00  8082      P2P
18  (pc65)    128      990!+ FWD  DESG 601e-08:17:f4:32:bb:00  8082      P2P
63  (pc66)    128      990!+ FWD  DESG 601e-08:17:f4:32:bb:00  8083      P2P
64  (pc66)    128      990!+ FWD  DESG 601e-08:17:f4:32:bb:00  8083      P2P
! = Automatic path cost.
+ = Portchannel cost, not the individual port cost.

-----
Spanning Tree Group 40: On (PVRST)
VLANs: 40

Current Root:          Path-Cost  Port Hello MaxAge FwdDel
6028 fc:cf:62:9d:67:00      990      17    2    20    15

Parameters:  Priority  Hello  MaxAge  FwdDel  Aging  Topology Change Counts
              28712      2      20      15      300          9

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
-----
17  (pc65)    128      990!+ FWD  ROOT 6028-fc:cf:62:9d:67:00  8082      P2P
18  (pc65)    128      990!+ FWD  ROOT 6028-fc:cf:62:9d:67:00  8082      P2P
63  (pc66)    128      990!+ FWD  DESG 7028-08:17:f4:32:bb:00  8083      P2P
64  (pc66)    128      990!+ FWD  DESG 7028-08:17:f4:32:bb:00  8083      P2P
! = Automatic path cost.
+ = Portchannel cost, not the individual port cost.

-----
Spanning Tree Group 128: Off (PVRST), FDB aging timer 300

```

Example 3-14 Output from show spanning tree command: G8264 switch 2

G8264_2#sh spanning-tree

Pvst+ compatibility mode enabled

Spanning Tree Group 1: On (PVRST)

VLANs: 1

Current Root: Path-Cost Port Hello MaxAge FwdDel
8001 fc:cf:62:9d:67:00 0 0 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
32769 2 20 15 300 0

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
------	------	------	-------	------	-------------------	----------	------

Note: There is no active STP port in Spanning Tree Group 1.

Spanning Tree Group 10: On (PVRST)

VLANs: 10

Current Root: Path-Cost Port Hello MaxAge FwdDel
600a 08:17:f4:32:bb:00 990 17 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
28682 2 20 15 300 6

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17 (pc65)	128	990!+	FWD	ROOT	600a-08:17:f4:32:bb:00	8082	P2P
18 (pc65)	128	990!+	FWD	ROOT	600a-08:17:f4:32:bb:00	8082	P2P
63 (pc66)	128	990!+	FWD	DESG	700a-fc:cf:62:9d:67:00	8083	P2P
64 (pc66)	128	990!+	FWD	DESG	700a-fc:cf:62:9d:67:00	8083	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 20: On (PVRST)

VLANs: 20

Current Root: Path-Cost Port Hello MaxAge FwdDel
6014 fc:cf:62:9d:67:00 0 0 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
24596 2 20 15 300 9

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17 (pc65)	128	990!+	FWD	DESG	6014-fc:cf:62:9d:67:00	8082	P2P
18 (pc65)	128	990!+	FWD	DESG	6014-fc:cf:62:9d:67:00	8082	P2P
63 (pc66)	128	990!+	FWD	DESG	6014-fc:cf:62:9d:67:00	8083	P2P
64 (pc66)	128	990!+	FWD	DESG	6014-fc:cf:62:9d:67:00	8083	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 30: On (PVRST)

VLANs: 30

Current Root: Path-Cost Port Hello MaxAge FwdDel
601e 08:17:f4:32:bb:00 990 17 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
28702 2 20 15 300 6

	Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17	(pc65)	128	990!	FWD	ROOT	601e-08:17:f4:32:bb:00	8082	P2P
18	(pc65)	128	990!	FWD	ROOT	601e-08:17:f4:32:bb:00	8082	P2P
63	(pc66)	128	990!	FWD	DESG	701e-fc:cf:62:9d:67:00	8083	P2P
64	(pc66)	128	990!	FWD	DESG	701e-fc:cf:62:9d:67:00	8083	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 40: On (PVRST)

VLANs: 40

Current Root: Path-Cost Port Hello MaxAge FwdDel
6028 fc:cf:62:9d:67:00 0 0 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
24616 2 20 15 300 9

	Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17	(pc65)	128	990!	FWD	DESG	6028-fc:cf:62:9d:67:00	8082	P2P
18	(pc65)	128	990!	FWD	DESG	6028-fc:cf:62:9d:67:00	8082	P2P
63	(pc66)	128	990!	FWD	DESG	6028-fc:cf:62:9d:67:00	8083	P2P
64	(pc66)	128	990!	FWD	DESG	6028-fc:cf:62:9d:67:00	8083	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

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.

3.3.5 Show running-config of all switches in Use Case 2

In the configuration output of the switches that is shown in Example 3-15, Example 3-16 on page 43, and Example 3-17 on page 45, you can see the configuration steps that we performed during our test. Important parameters and detail are highlighted in red.

Example 3-15 Output of the show running command: EN2092 switch

```
Flex#sh run
Current configuration:
!
version "7.2.2.2"
switch-type "IBM Flex System EN2092 1Gb Ethernet Scalable Switch"
!
hostname "Flex"
!
interface port INTA2
    tagging
    exit
!
interface port INTA7
    shutdown
    exit
!
interface port EXT4
    name "TEST_PC"
    tagging
    exit
!
interface port EXT21
    name "TO_G8264_2_Port63"
    tagging
    pvid 10
    exit
!
interface port EXT22
    name "TO_G8264_1_Port63"
    tagging
    pvid 10
    exit
!
interface port EXT23
    name "TO_G8264_1_Port64"
    tagging
    pvid 10
    exit
!
interface port EXT24
    name "TO_G8264_2_Port64"
    tagging
    pvid 10
    exit
!
vlan 1
    member INTA1-EXT20
    no member EXT21-EXT24
!
vlan 10
    enable
    name "Server"
```

```
        member EXT4,EXT21-EXT24
!
vlan 20
    enable
    name "Data20"
    member EXT4,EXT21-EXT24
!
vlan 30
    enable
    name "Data30"
    member EXT4,EXT21-EXT24
!
vlan 40
    enable
    name "Data40"
    member EXT4,EXT21-EXT24
!
spanning-tree stp 10 vlan 10

spanning-tree stp 20 vlan 20

spanning-tree stp 30 vlan 30

spanning-tree stp 40 vlan 40

!
interface port EXT21
    lacp mode active
    lacp key 121
!
interface port EXT22
    lacp mode active
    lacp key 122
!
interface port EXT23
    lacp mode active
    lacp key 122
!
interface port EXT24
    lacp mode active
    lacp key 121
!
lldp enable
!
end
```

Example 3-16 Output of the show running command: G8264 switch 1

```
G8264_1#sh run
Current configuration:
!
version "7.2.2"
switch-type "IBM Networking Operating System RackSwitch G8264"
!
hostname "G8264_1"
!
interface port 17
    name "CrossLink"
    tagging
    pvid 10
    exit
!
interface port 18
    name "CrossLink"
    tagging
    pvid 10
    exit
!
interface port 63
    name "UPLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
interface port 64
    name "UPLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
vlan 1
    member 1-16,19-62
    no member 17-18,63-64
!
vlan 10
    enable
    name "none"
    member 17-18,63-64
!
vlan 20
    enable
    name "none"
    member 17-18,63-64
!
vlan 30
    enable
    name "none"
    member 17-18,63-64
!
vlan 40
    enable
    name "VLAN 40"
    member 17-18,63-64
!
spanning-tree stp 10 bridge priority 24576
spanning-tree stp 10 vlan 10
```

```
spanning-tree stp 20 bridge priority 28672
spanning-tree stp 20 vlan 20

spanning-tree stp 30 bridge priority 24576
spanning-tree stp 30 vlan 30

spanning-tree stp 40 bridge priority 28672
spanning-tree stp 40 vlan 40

!
interface port 17
    lacp mode active
    lacp key 117
!
interface port 18
    lacp mode active
    lacp key 117
!
interface port 63
    lacp mode active
    lacp key 163
!
interface port 64
    lacp mode active
    lacp key 163
!
lldp enable
!
end
```

Example 3-17 Output of the show running command: G8264 switch 2

```
G8264_2#sh run
Current configuration:
!
version "7.2.2"
switch-type "IBM Networking Operating System RackSwitch G8264"
!
hostname "G8264_2"
!
interface port 17
    name "CrossLink"
    tagging
    pvid 10
    exit
!
interface port 18
    name "CrossLink"
    tagging
    pvid 10
    exit
!
interface port 63
    name "UPLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
interface port 64
    name "UPLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
vlan 1
    member 1-16,19-62
    no member 17-18,63-64
!
vlan 10
    enable
    name "none"
    member 17-18,63-64
!
vlan 20
    enable
    name "none"
    member 17-18,63-64
!
vlan 30
    enable
    name "none"
    member 17-18,63-64
!
vlan 40
    enable
    name "VLAN 40"
    member 17-18,63-64
!
spanning-tree stp 10 bridge priority 28672
spanning-tree stp 10 vlan 10
```

```
spanning-tree stp 20 bridge priority 24576
spanning-tree stp 20 vlan 20

spanning-tree stp 30 bridge priority 28672
spanning-tree stp 30 vlan 30

spanning-tree stp 40 bridge priority 24576
spanning-tree stp 40 vlan 40

!
interface port 17
    lacp mode active
    lacp key 117
!
interface port 18
    lacp mode active
    lacp key 117
!
interface port 63
    lacp mode active
    lacp key 163
!
interface port 64
    lacp mode active
    lacp key 163
!
lldp enable
!
end
```

3.4 Use Case 3: Link aggregation and MST

For this use case, we replaced the PVRST with MST. Again, we have three pairs of 10 GE links between the three switches, which were running 802.1q trunking and LACP. The VLANs 10 and 30, and 20 and 40 are manually distributed over the uplinks from the Flex switch, as shown in Figure 3-5 and Figure 3-6 on page 48.

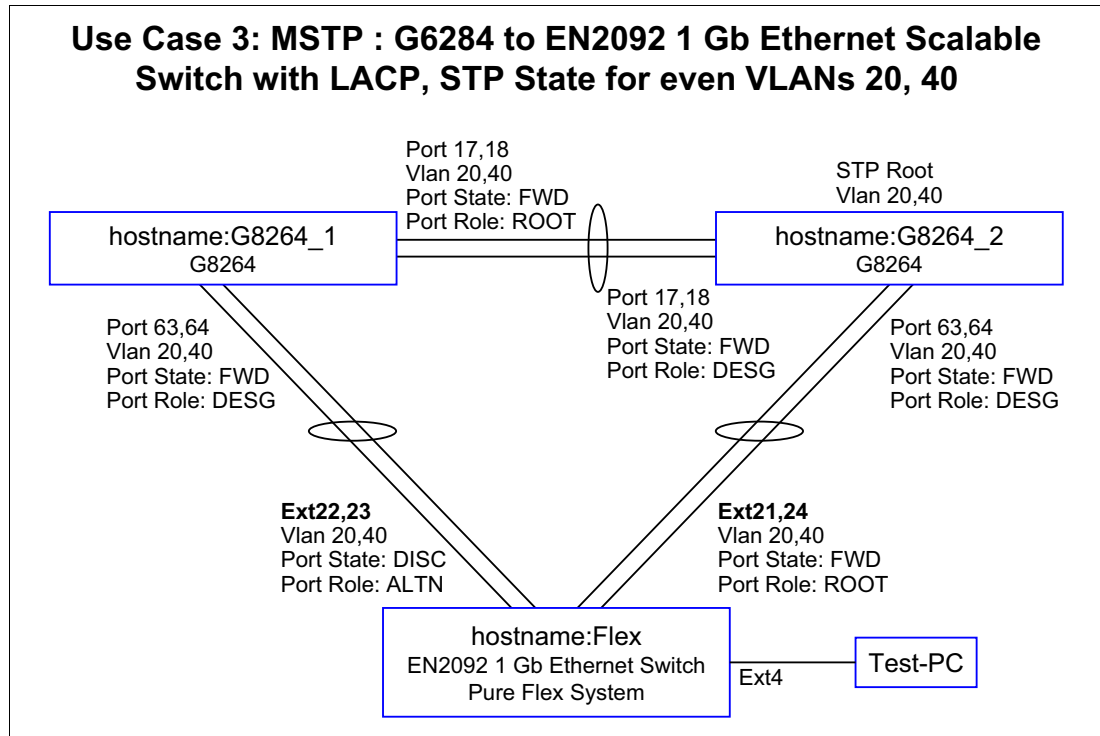


Figure 3-5 Use Case 3: Even-numbered VLANs

Use Case 3: MSTP : G6284 to EN2092 1 Gb Ethernet Scalable Switch with LACP, STP State for odd VLANs 10, 30

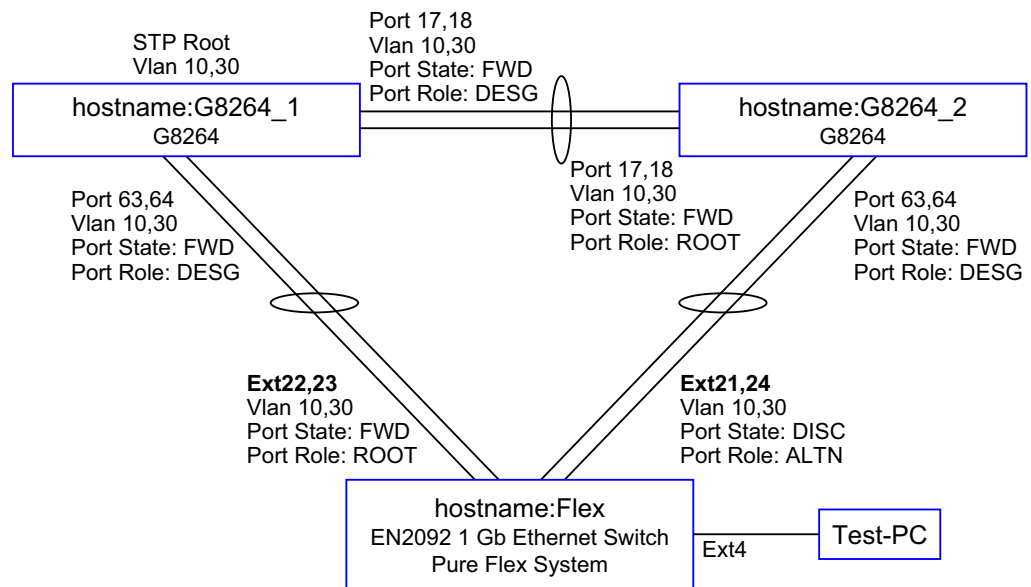


Figure 3-6 Use Case 3: Odd-numbered VLANs

3.4.1 Verifying the topology that was used by using lldp

To verify the topology, we used the **show lldp remote-device** command on the switches, as shown in Example 3-18.

Example 3-18 Checking the topology use show lldp remote-device command

Flex#**sh lldp remote-device**

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
-----	-----	-----	-----	-----
EXT22	1	08 17 f4 32 bb 00	63	G8264_1
EXT21	2	fc cf 62 9d 67 00	63	G8264_2
EXT23	5	08 17 f4 32 bb 00	64	G8264_1
EXT24	8	fc cf 62 9d 67 00	64	G8264_2

G8264_1#**sh lldp remote-device**

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
-----	-----	-----	-----	-----
17	1	fc cf 62 9d 67 00	17	G8264_2
63	2	08 17 f4 76 78 00	50	Flex
18	3	fc cf 62 9d 67 00	18	G8264_2
64	4	08 17 f4 76 78 00	51	Flex

G8264_2#**sh lldp remote-device**

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
-----	-----	-----	-----	-----
17	1	08 17 f4 32 bb 00	17	G8264_1
63	2	08 17 f4 76 78 00	49	Flex
18	3	08 17 f4 32 bb 00	18	G8264_1
64	4	08 17 f4 76 78 00	52	Flex

3.4.2 Verifying trunks

To verify which VLANs are active on which trunk, we used the **show interface trunk** command on the switches, as shown in Example 3-19.

Example 3-19 Show interface trunk command

```
Flex#sh interface trunk
Alias  Port Tag RMON Lrn Fld PVID      NAME                               VLAN(s)
-----
EXT21  49   y   d   e   e   10  TO_G8264_2_Port63 10 20 30 40
EXT22  50   y   d   e   e   10  TO_G8264_1_Port63 10 20 30 40
EXT23  51   y   d   e   e   10  TO_G8264_1_Port64 10 20 30 40
EXT24  52   y   d   e   e   10  TO_G8264_2_Port64 10 20 30 40
```

* = PVID is tagged.

```
G8264_1#sh interface trunk
Alias  Port Tag RMON Lrn Fld PVID      NAME                               VLAN(s)
-----
16     16   n   d   e   e    1   1
17     17   y   d   e   e   10  CrossLink          10 20 30 40
18     18   y   d   e   e   10  CrossLink          10 20 30 40
63     63   y   d   e   e   10  UPLINK_TO_FLEX    10 20 30 40
64     64   y   d   e   e   10  UPLINK_TO_FLEX    10 20 30 40
```

* = PVID is tagged.

```
G8264_2#sh interface trunk
Alias  Port Tag RMON Lrn Fld PVID      NAME                               VLAN(s)
-----
17     17   y   d   e   e   10  CrossLink          10 20 30 40
18     18   y   d   e   e   10  CrossLink          10 20 30 40
63     63   y   d   e   e   10  UPLINK_TO_FLEX    10 20 30 40
64     64   y   d   e   e   10  UPLINK_TO_FLEX    10 20 30 40
```

* = PVID is tagged.

3.4.3 Verifying link aggregation by using lacp

We verified the link aggregation configuration of the switches by executing the **show lacp information** command, as shown in Example 3-20.

Example 3-20 Show lacp information command

Flex#sh lacp info

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
EXT21	active	121	121	yes	32768	49	53	up	1
EXT22	active	122	122	yes	32768	50	54	up	1
EXT23	active	122	122	yes	32768	50	54	up	1
EXT24	active	121	121	yes	32768	49	53	up	1

G8264_1#sh lacp information

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
17	active	117	117	yes	32768	17	65	up	1
18	active	117	117	yes	32768	17	65	up	1
63	active	163	163	yes	32768	63	66	up	1
64	active	163	163	yes	32768	63	66	up	1

G8264_2#sh lacp information

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
17	active	117	117	yes	32768	17	65	up	1
18	active	117	117	yes	32768	17	65	up	1
63	active	163	163	yes	32768	63	66	up	1
64	active	163	163	yes	32768	63	66	up	1

3.4.4 Verifying MST spanning tree configuration

In the next step, we verified the MST spanning tree configuration of the switches by executing the **show spanning-tree** command. As shown in Figure 3-5 on page 47 and Figure 3-6 on page 48, we have two spanning trees, one for even VLANs and one for odd VLANs. By using the show spanning tree command, you can verify the status of the respective Ethernet interface's VLAN, port state, and port role.

The commands that were run on the three switches produced the following outputs:

- ▶ EN2029: Example 3-21 on page 52
- ▶ G8264 switch 1: Example 3-22 on page 53
- ▶ G8264 switch 2: Example 3-23 on page 54

Example 3-21 Verifying the MST spanning tree configuration: Flex System switch

```
Flex#sh spanning-tree
```

```
-----  
Pvst+ compatibility mode enabled
```

```
Mstp Digest: 0xe821ccee7501115289b37c79a72e07c9
```

```
-----  
Spanning Tree Group 1: On (MSTP)
```

```
VLANs MAPPED: 10 30
```

```
VLANs: 10 30
```

```
Current Root:          Path-Cost  Port  
6000 08:17:f4:32:bb:00    990  EXT22
```

```
Parameters:  Priority  Aging  Topology Change Counts  
              61440    300           4
```

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	FWD	DESG	f000-08:17:f4:76:78:00	8020	P2P
EXT21 (pc53)	128	990!+	DISC	ALTN	7000-fc:cf:62:9d:67:00	8083	P2P
EXT22 (pc54)	128	990!+	FWD	ROOT	6000-08:17:f4:32:bb:00	8083	P2P
EXT23 (pc54)	128	990!+	FWD	ROOT	6000-08:17:f4:32:bb:00	8083	P2P
EXT24 (pc53)	128	990!+	DISC	ALTN	7000-fc:cf:62:9d:67:00	8083	P2P

```
! = Automatic path cost.
```

```
+ = Portchannel cost, not the individual port cost.
```

```
-----  
Spanning Tree Group 2: On (MSTP)
```

```
VLANs MAPPED: 20 40
```

```
VLANs: 20 40
```

```
Current Root:          Path-Cost  Port  
6000 fc:cf:62:9d:67:00    990  EXT21
```

```
Parameters:  Priority  Aging  Topology Change Counts  
              61440    300           6
```

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	FWD	DESG	f000-08:17:f4:76:78:00	8020	P2P
EXT21 (pc53)	128	990!+	FWD	ROOT	6000-fc:cf:62:9d:67:00	8083	P2P
EXT22 (pc54)	128	990!+	DISC	ALTN	7000-08:17:f4:32:bb:00	8083	P2P
EXT23 (pc54)	128	990!+	DISC	ALTN	7000-08:17:f4:32:bb:00	8083	P2P
EXT24 (pc53)	128	990!+	FWD	ROOT	6000-fc:cf:62:9d:67:00	8083	P2P

```
! = Automatic path cost.
```

```
+ = Portchannel cost, not the individual port cost.
```

Example 3-22 Verifying the MST spanning tree configuration: G8264 switch 1

```
G8264_1(config)#sh spanning-tree
```

```
-----  
Pvst+ compatibility mode enabled
```

```
Mstp Digest: 0xe821ccee7501115289b37c79a72e07c9
```

```
-----  
Spanning Tree Group 1: On (MSTP)
```

```
VLANs MAPPED: 10 30
```

```
VLANs: 10 30
```

```
Current Root:          Path-Cost  Port  
6000 08:17:f4:32:bb:00      0      0
```

```
Parameters:  Priority  Aging  Topology Change Counts  
              24576    300           8
```

	Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17	(pc65)	128	990!	FWD	DESG	6000-08:17:f4:32:bb:00	8082	P2P
18	(pc65)	128	990!	FWD	DESG	6000-08:17:f4:32:bb:00	8082	P2P
63	(pc66)	128	990!	FWD	DESG	6000-08:17:f4:32:bb:00	8083	P2P
64	(pc66)	128	990!	FWD	DESG	6000-08:17:f4:32:bb:00	8083	P2P

```
! = Automatic path cost.
```

```
+ = Portchannel cost, not the individual port cost.
```

```
-----  
Spanning Tree Group 2: On (MSTP)
```

```
VLANs MAPPED: 20 40
```

```
VLANs: 20 40
```

```
Current Root:          Path-Cost  Port  
6000 fc:cf:62:9d:67:00    990    17
```

```
Parameters:  Priority  Aging  Topology Change Counts  
              28672    300           8Press q to quit, any other key to
```

	Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17	(pc65)	128	990!	FWD	ROOT	6000-fc:cf:62:9d:67:00	8082	P2P
18	(pc65)	128	990!	FWD	ROOT	6000-fc:cf:62:9d:67:00	8082	P2P
63	(pc66)	128	990!	FWD	DESG	7000-08:17:f4:32:bb:00	8083	P2P
64	(pc66)	128	990!	FWD	DESG	7000-08:17:f4:32:bb:00	8083	P2P

```
! = Automatic path cost.
```

```
+ = Portchannel cost, not the individual port cost.
```

Example 3-23 Verifying the MST spanning tree configuration: G8264 switch 2

G8264_2(config)#sh spanning-tree

Pvst+ compatibility mode enabled

Mstp Digest: 0xe821ccee7501115289b37c79a72e07c9

Spanning Tree Group 1: On (MSTP)

VLANs MAPPED: 10 30

VLANs: 10 30

Current Root: Path-Cost Port
6000 08:17:f4:32:bb:00 990 17

Parameters: Priority Aging Topology Change Counts
 28672 300 2

	Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17	(pc65)	128	990!+	FWD	ROOT	6000-08:17:f4:32:bb:00	8082	P2P
18	(pc65)	128	990!+	FWD	ROOT	6000-08:17:f4:32:bb:00	8082	P2P
63	(pc66)	128	990!+	FWD	DESG	7000-fc:cf:62:9d:67:00	8083	P2P
64	(pc66)	128	990!+	FWD	DESG	7000-fc:cf:62:9d:67:00	8083	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 2: On (MSTP)

VLANs MAPPED: 20 40

VLANs: 20 40

Current Root: Path-Cost Port
6000 fc:cf:62:9d:67:00 0 0

Parameters: Priority Aging Topology Change Counts
 24576 300 6Press q to quit, any other key to

	Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17	(pc65)	128	990!+	FWD	DESG	6000-fc:cf:62:9d:67:00	8082	P2P
18	(pc65)	128	990!+	FWD	DESG	6000-fc:cf:62:9d:67:00	8082	P2P
63	(pc66)	128	990!+	FWD	DESG	6000-fc:cf:62:9d:67:00	8083	P2P
64	(pc66)	128	990!+	FWD	DESG	6000-fc:cf:62:9d:67:00	8083	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

3.4.5 Show running-config of all switches in Use Case 3

In the configuration output of the switches that is shown in Example 3-24, you can see the necessary configuration steps that we performed during our test. Important parameters and detail are highlighted in red.

The commands that were run on the three switches produced the following outputs:

- ▶ EN2029: Example 3-24
- ▶ G8264 switch 1: Example 3-25 on page 57
- ▶ G8264 switch 2: Example 3-26 on page 58

Example 3-24 Output of the show running command: EN2092 switch

```
Flex#sh run
Current configuration:
!
version "7.2.2.2"
switch-type "IBM Flex System EN2092 1Gb Ethernet Scalable Switch"
!
hostname "Flex"
!
interface port INTA2
    tagging
    exit
!
interface port INTA7
    shutdown
    exit
!
interface port EXT4
    name "TEST_PC"
    tagging
    exit
!
interface port EXT21
    name "T0_G8264_2_Port63"
    tagging
    pvid 10
    exit
!
interface port EXT22
    name "T0_G8264_1_Port63"
    tagging
    pvid 10
    exit
!
interface port EXT23
    name "T0_G8264_1_Port64"
    tagging
    pvid 10
    exit
!
interface port EXT24
    name "T0_G8264_2_Port64"
    tagging
    pvid 10
    exit
!
vlan 1
```

```

        member INTA1-EXT20
        no member EXT21-EXT24
    !
vlan 10
    enable
    name "Server"
    member EXT4,EXT21-EXT24
!
vlan 20
    enable
    name "Data20"
    member EXT4,EXT21-EXT24
!
vlan 30
    enable
    name "Data30"
    member EXT4,EXT21-EXT24
!
vlan 40
    enable
    name "Data40"
    member EXT4,EXT21-EXT24
!
spanning-tree mstp version 10
spanning-tree mstp name "PureFlex"
spanning-tree mode mst
spanning-tree mstp cist-add-vlan 1
spanning-tree mstp cist-add-vlan 4095
!
spanning-tree stp 1 vlan 10
spanning-tree stp 1 vlan 30

spanning-tree stp 2 vlan 20
spanning-tree stp 2 vlan 40

!
interface port EXT21
    lacp mode active
    lacp key 121
!
interface port EXT22
    lacp mode active
    lacp key 122
!
interface port EXT23
    lacp mode active
    lacp key 122
!
interface port EXT24
    lacp mode active
    lacp key 121
!
lldp enable
!
End

```

Example 3-25 Output of the show running command: G8264 switch 1

```
G8264_1#sh run
Current configuration:
!
version "7.2.2"
switch-type "IBM Networking Operating System RackSwitch G8264"
!
hostname "G8264_1"
!
!
interface port 17
    name "CrossLink"
    tagging
    pvid 10
    exit
!
interface port 18
    name "CrossLink"
    tagging
    pvid 10
    exit
!
interface port 63
    name "UPLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
interface port 64
    name "UPLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
vlan 1
    member 1-16,19-62
    no member 17-18,63-64
!
vlan 10
    enable
    name "none"
    member 17-18,63-64
!
vlan 20
    enable
    name "none"
    member 17-18,63-64
!
vlan 30
    enable
    name "none"
    member 17-18,63-64
!
vlan 40
    enable
    name "VLAN 40"
    member 17-18,63-64
!
spanning-tree mstp version 10
spanning-tree mstp name "PureFlex"
```

```

spanning-tree mode mst
spanning-tree mstp cist-add-vlan 1
spanning-tree mstp cist-add-vlan 4095
!
spanning-tree stp 1 bridge priority 24576
spanning-tree stp 1 vlan 10
spanning-tree stp 1 vlan 30

spanning-tree stp 2 bridge priority 28672
spanning-tree stp 2 vlan 20
spanning-tree stp 2 vlan 40

spanning-tree stp 40 bridge priority 28672

!
interface port 17
    lacp mode active
    lacp key 117
!
interface port 18
    lacp mode active
    lacp key 117
!
interface port 63
    lacp mode active
    lacp key 163
!
interface port 64
    lacp mode active
    lacp key 163
!
lldp enable
!
end

```

Example 3-26 Output of the show running command: G8264 switch 2

```

G8264_2#sh run

Current configuration:
!
version "7.2.2"
switch-type "IBM Networking Operating System RackSwitch G8264"
!
hostname "G8264_2"
!
interface port 17
    name "CrossLink"
    tagging
    pvid 10
    exit
!
interface port 18
    name "CrossLink"
    tagging
    pvid 10
    exit
!

```

```

interface port 63
    name "UPLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
interface port 64
    name "UPLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
vlan 1
    member 1-16,19-62
    no member 17-18,63-64
!
vlan 10
    enable
    name "none"
    member 17-18,63-64
!
vlan 20
    enable
    name "none"
    member 17-18,63-64
!
vlan 30
    enable
    name "none"
    member 17-18,63-64
!
vlan 40
    enable
    name "VLAN 40"
    member 17-18,63-64
!
spanning-tree mstp version 10
spanning-tree mstp name "PureFlex"
spanning-tree mode mst
spanning-tree mstp cist-add-vlan 1
spanning-tree mstp cist-add-vlan 4095
!
spanning-tree stp 1 bridge priority 28672
spanning-tree stp 1 vlan 10
spanning-tree stp 1 vlan 30

spanning-tree stp 2 bridge priority 24576
spanning-tree stp 2 vlan 20
spanning-tree stp 2 vlan 40

!
interface port 17
    lacp mode active
    lacp key 117
!
interface port 18
    lacp mode active
    lacp key 117
!

```

```

interface port 63
    lacp mode active
    lacp key 163
!
interface port 64
    lacp mode active
    lacp key 163
!
lldp enable
!
end

```

3.5 Use Case 4: Link aggregation, MSTP and VLAG

The concept of virtual link aggregation (VLAG) shows the pair of G8264 switches logically as one switch entity. Together with LACP, this configuration allows the typical *triangle design* to be run, as shown in Figure 3-7.

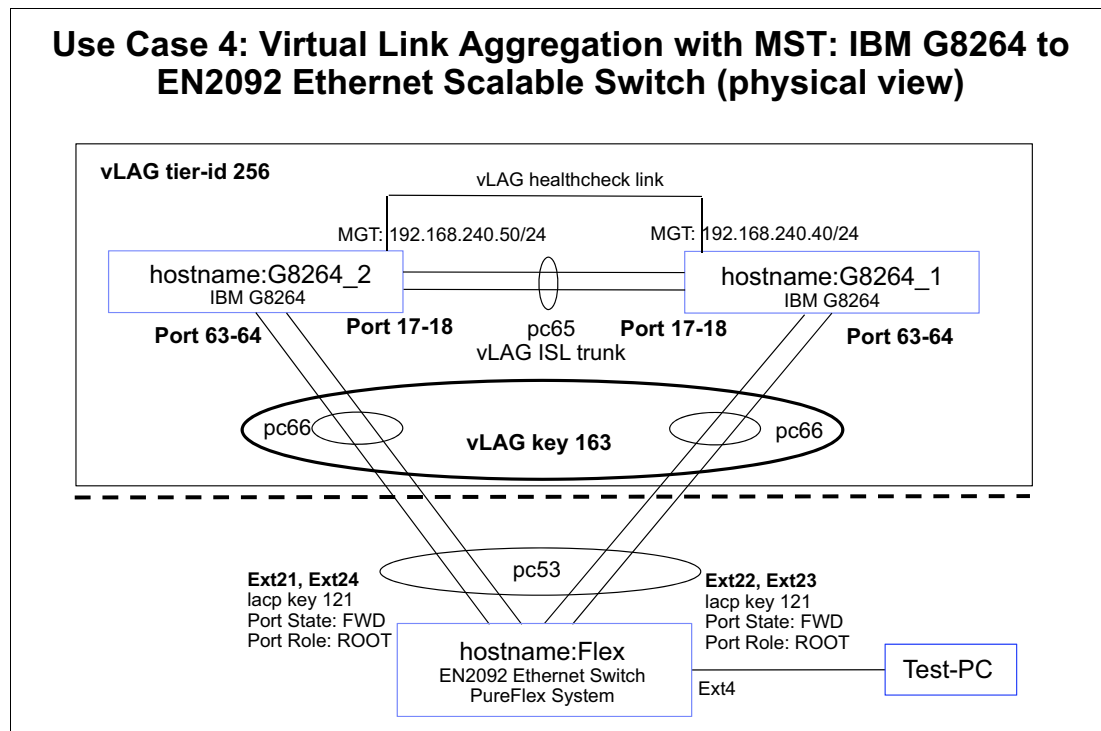


Figure 3-7 VLAG with MST

Figure 3-8 on page 61 shows the logical view of the setup. To the IBM Flex Switch, the pair of IBM RackSwitch G8264 switches looks like one switch.

Use Case 4: Virtual Link Aggregation with MST: IBM G8264 to EN2092 Ethernet Scalable Switch (logical view)

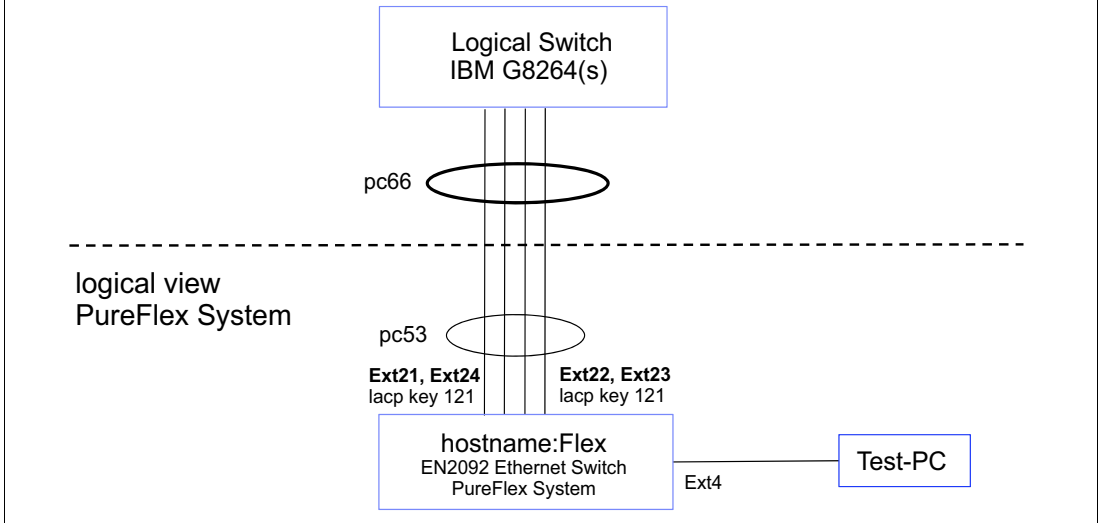


Figure 3-8 VLAG with MST (logical view)

3.5.1 Verifying the topology by using lldp

To verify the topology, we used the **show lldp remote-device** command on the switches, as shown in Example 3-27.

Example 3-27 Verifying the topology by using lldp

```
G8264_1#show lldp remote-device
LLDP Remote Devices Information
```

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
17	1	fc cf 62 9d 67 00	17	G8264_2
63	2	08 17 f4 76 78 00	50	Flex
18	3	fc cf 62 9d 67 00	18	G8264_2
64	4	08 17 f4 76 78 00	51	Flex

```
G8264_2#sh lldp remote-device
LLDP Remote Devices Information
```

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
17	1	08 17 f4 32 bb 00	17	G8264_1
63	2	08 17 f4 76 78 00	49	Flex
18	3	08 17 f4 32 bb 00	18	G8264_1
64	4	08 17 f4 76 78 00	52	Flex

```
Flex#sh lldp remote-device
LLDP Remote Devices Information
```

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
EXT22	1	08 17 f4 32 bb 00	63	G8264_1
EXT21	2	fc cf 62 9d 67 00	63	G8264_2
INTA1	3	5c f3 fc 5f 43 9d	5c-f3-fc-5f-43-9d	
INTA4	4	5c f3 fc 6e 23 41	5c-f3-fc-6e-23-41	
EXT23	5	08 17 f4 32 bb 00	64	G8264_1
EXT5	6	00 0d ec a3 8f 81	mgmt0	vie
EXT7	7	00 05 9b 7b 84 01	mgmt0	str
EXT24	8	fc cf 62 9d 67 00	64	G8264_2

3.5.2 Verify interface status

To verify the interface, we used the **show interface status** command on the switches, as shown in Example 3-28.

Example 3-28 Verify interface status

```
G8264_1#sh interface st
```

Alias	Port	Speed	Duplex	Flow Ctrl	Link	Name
				--TX--RX--		
1	1	10000	full	no no	down	1
--						
16	16	10000	full	no no	down	16
17	17	10000	full	no no	up	CrossLink
18	18	10000	full	no no	up	CrossLink
19	19	1G/10G	full	no no	down	19
--						

62	62	1G/10G	full	no	no	down	62
63	63	10000	full	no	no	up	UPLINK_TO_FLEX
64	64	10000	full	no	no	up	UPLINK_TO_FLEX
MGT	65	1000	full	yes	yes	up	MGT

G8264_2#sh interface status

Alias	Port	Speed	Duplex	Flow Ctrl		Link	Name
				--TX--	--RX--		
1	1	40000	full	no	no	down	1
--							
16	16	10000	full	no	no	down	16
17	17	10000	full	no	no	up	CrossLink
18	18	10000	full	no	no	up	CrossLink
19	19	1G/10G	full	no	no	down	19
--							
62	62	1G/10G	full	no	no	down	62
63	63	10000	full	no	no	up	UPLINK_TO_FLEX
64	64	10000	full	no	no	up	UPLINK_TO_FLEX
MGT	65	1000	full	yes	yes	up	MGT

Flex#show interface status

Alias	Port	Speed	Duplex	Flow Ctrl		Link	Name
				--TX--	--RX--		
INTA1	1	1000	full	yes	yes	up	INTA1
INTA2	2	1000	full	yes	yes	up	INTA2
INTA3	3	1000	full	yes	yes	down	INTA3
INTA4	4	1000	full	yes	yes	up	INTA4
INTA5	5	1000	full	yes	yes	down	INTA5
INTA6	6	1000	full	yes	yes	down	INTA6
INTA7	7	1000	full	yes	yes	disabled	INTA7
INTA8	8	1000	full	yes	yes	down	INTA8
INTA9	9	1000	full	yes	yes	down	INTA9
INTA10	10	1000	full	yes	yes	down	INTA10
INTA11	11	1000	full	yes	yes	down	INTA11
INTA12	12	1000	full	yes	yes	down	INTA12
INTA13	13	1000	full	yes	yes	down	INTA13
INTA14	14	1000	full	yes	yes	down	INTA14
INTB1	15	1000	full	yes	yes	down	INTB1
--							
INTB14	28	1000	full	yes	yes	down	INTB14
EXT1	29	1000	full	no	no	up	EXT1
EXT2	30	1000	full	no	no	up	EXT2
EXT3	31	1000	full	no	no	up	EXT3
EXT4	32	1000	full	no	no	up	TEST_PC
EXT5	33	1000	full	no	no	up	EXT5
--							
EXT20	48	any	any	no	no	down	EXT20
EXT21	49	10000	full	no	no	up	T0_G8264_2_Port63
EXT22	50	10000	full	no	no	up	T0_G8264_1_Port63
EXT23	51	10000	full	no	no	up	T0_G8264_1_Port64
EXT24	52	10000	full	no	no	up	T0_G8264_2_Port64
MGT1	53	1000	full	no	no	up	MGT1

3.5.3 Verifying trunks

To verify which VLANs are active on which trunk, we used the **show interface trunk** command on the switches, as shown in Example 3-29.

Example 3-29 Verifying trunks

G8264_1#sh interface trunk									
Alias	Port	Tag	RMON	Lrn	Fld	PVID	NAME	VLAN(s)	
1	1	n	d	e	e	1		1	
--									
16	16	n	d	e	e	1		1	
17	17	y	d	d	e	4094	CrossLink	10 20 30 40	4094
18	18	y	d	d	e	4094	CrossLink	10 20 30 40	4094
19	19	n	d	e	e	1		1	
--									
62	62	n	d	e	e	1		1	
63	63	y	d	e	e	10	UPLINK_TO_FLEX	10 20 30 40	
64	64	y	d	e	e	10	UPLINK_TO_FLEX	10 20 30 40	
MGT	65	n	d	e	e	4095		4095	
G8264_2#sh interface trunk									
Alias	Port	Tag	RMON	Lrn	Fld	PVID	NAME	VLAN(s)	
1	1	n	d	e	e	1		1	
--									
16	16	n	d	e	e	1		1	
17	17	y	d	d	e	4094	CrossLink	10 20 30 40	4094
18	18	y	d	d	e	4094	CrossLink	10 20 30 40	4094
19	19	n	d	e	e	1		1	
--									
62	62	n	d	e	e	1		1	
63	63	y	d	e	e	10	UPLINK_TO_FLEX	10 20 30 40	
64	64	y	d	e	e	10	UPLINK_TO_FLEX	10 20 30 40	
MGT	65	n	d	e	e	4095		4095	
Flex#sh interface trunk									
Alias	Port	Tag	RMON	Lrn	Fld	PVID	NAME	VLAN(s)	
INTA1	1	n	d	e	e	1	INTA1	1	
INTA2	2	y	d	e	e	1	INTA2	1	
INTA3	3	n	d	e	e	1	INTA3	1	
INTA4	4	n	d	e	e	1	INTA4	1	
INTA5	5	n	d	e	e	1	INTA5	1	
INTA6	6	n	d	e	e	1	INTA6	1	
INTA7	7	n	d	e	e	1	INTA7	1	
INTA8	8	n	d	e	e	1	INTA8	1	
INTA9	9	n	d	e	e	1	INTA9	1	
INTA10	10	n	d	e	e	1	INTA10	1	
INTA11	11	n	d	e	e	1	INTA11	1	
INTA12	12	n	d	e	e	1	INTA12	1	
INTA13	13	n	d	e	e	1	INTA13	1	
INTA14	14	n	d	e	e	1	INTA14	1	
INTB1	15	n	d	e	e	1	INTB1	1	
--									
INTB14	28	n	d	e	e	1	INTB14	1	
EXT1	29	n	d	e	e	1	EXT1	1	
EXT2	30	n	d	e	e	1	EXT2	1	
EXT3	31	n	d	e	e	1	EXT3	1	

EXT4	32	y	d	e	e	1	TEST_PC	1	10	20	30	40
EXT5	33	n	d	e	e	1	EXT5	1				
--												
EXT20	48	n	d	e	e	1	EXT20	1				
EXT21	49	y	d	e	e	10	TO_G8264_2_Port63	10	20	30	40	
EXT22	50	y	d	e	e	10	TO_G8264_1_Port63	10	20	30	40	
EXT23	51	y	d	e	e	10	TO_G8264_1_Port64	10	20	30	40	
EXT24	52	y	d	e	e	10	TO_G8264_2_Port64	10	20	30	40	
MGT1	53	y	d	e	e	4095	MGT1	4095				

3.5.4 Verify spanning tree

We verified the spanning tree configuration of the switches by executing the **show spanning-tree** command, as shown in Example 3-30.

Example 3-30 Verify spanning tree

```
G8264_1#sh spanning-tree
-----
Pvst+ compatibility mode enabled

Mstp Digest: 0xe821ccee7501115289b37c79a72e07c9

-----
Spanning Tree Group 1: On (MSTP)
VLANs MAPPED: 10 30
VLANs: 10 30

Current Root:          Path-Cost  Port
6000 08:17:f4:32:bb:00      0      0

Parameters:  Priority  Aging  Topology Change Counts
              24576    300      20

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
-----
17  (pc65)    128      990!+ FWD  DESG 6000-08:17:f4:32:bb:00    8082        P2P
18  (pc65)    128      990!+ FWD  DESG 6000-08:17:f4:32:bb:00    8082        P2P
63  (pc66)    128      200!+ FWD  DESG 6000-08:17:f4:32:bb:00    8102        P2P
64  (pc66)    128      200!+ FWD  DESG 6000-08:17:f4:32:bb:00    8102        P2P
! = Automatic path cost.
+ = Portchannel cost, not the individual port cost.

-----
Spanning Tree Group 2: On (MSTP)
VLANs MAPPED: 20 40
VLANs: 20 40

Current Root:          Path-Cost  Port
6000 fc:cf:62:9d:67:00    990    17

Parameters:  Priority  Aging  Topology Change Counts
              28672    300      19

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
-----
17  (pc65)    128      990!+ FWD  ROOT 6000-fc:cf:62:9d:67:00    8082        P2P
18  (pc65)    128      990!+ FWD  ROOT 6000-fc:cf:62:9d:67:00    8082        P2P
63  (pc66)    128      200!+ FWD  DESG 7000-08:17:f4:32:bb:00    8102        P2P
```

```

64      (pc66)  128      200!+ FWD   DESG 7000-08:17:f4:32:bb:00      8102      P2P
! = Automatic path cost.
+ = Portchannel cost, not the individual port cost.

```

```

-----
Spanning Tree Group 32: Off (MSTP), FDB aging timer 300
VLANs MAPPED: 4094
VLANs: 4094

```

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
17	(pc65)	0	0	FWD *			
18	(pc65)	0	0	FWD *			

* = STP turned off for this port.

3.5.5 Verify virtual link aggregation

We verified the link aggregation configuration of the switches by executing various **show** commands, as shown in Example 3-31.

Example 3-31 Verify virtual link aggregation

```

G8264_1#show lacp information
port    mode    adminkey  operkey  selected  prio  aggr  trunk  status  minlinks
-----
1        off      1         1        no        32768 --    --    --      1
--
16       off      16        16       no        32768 --    --    --      1
17       active   117       117      yes       32768 17    65    up      1
18       active   117       117      yes       32768 17    65    up      1
19       off      19        19       no        32768 --    --    --      1
--
62       off      62        62       no        32768 --    --    --      1
63       active   163       163      yes       32768 63    66    up      1
64       active   163       163      yes       32768 63    66    up      1

```

```

G8264_1#sh lacp aggregator 63

```

```

Aggregator Id 63

```

```

-----
Aggregator MAC address - 08:17:f4:32:bb:a0
Actor System Priority   - 32768
Actor System ID        - 08:17:f4:c3:dd:ff
Individual              - FALSE
Actor Oper Key          - 163
Partner System Priority - 32768
Partner System ID      - 08:17:f4:76:78:00
Partner Oper Key        - 121
ready                   - TRUE
Min-Links               - 1
Number of Ports in aggr - 2
index 0   port 63
index 1   port 64

```

```

G8264_1#show vlag adminkey 163

```

```

vLAG is enabled on admin key 163

```

```

Current LACP params for 63: active, Priority 32768, Admin Key 163, Min-Links 1

```

```

Current LACP params for 64: active, Priority 32768, Admin Key 163, Min-Links 1

```

```
G8264_1#show vlag information
vLAG Tier ID: 256
vLAG system MAC: 08:17:f4:c3:dd:ff
Local MAC 08:17:f4:32:bb:00 Priority 0 Admin Role PRIMARY (Operational Role PRIMARY)
Peer MAC fc:cf:62:9d:67:00 Priority 0
Health local 192.168.240.40 peer 192.168.240.50 State UP
ISL trunk id 65
ISL state Up
Startup Delay Interval: 120s (Finished)
```

vLAG 65: config with admin key 163, associated trunk 66, state formed

```
G8264_1#show vlag isl
ISL_ID   ISL_Vlan   ISL_Trunk   ISL_Members   Link_State   Trunk_State
65       4094      Adminkey 117      17            UP           UP
          18            18            UP           UP
```

```
G8264_1#show vlag statistics
```

```
vLAG PDU sent:
Role Election:      2      System Info:      1
Peer Instance Enable: 2      Peer Instance Disable: 0
FDB Dynamic Add:    4      FDB Dynamic Del:    4
FDB Inactive Add:   0      FDB Inactive Del:   0
Health Check:       384    ISL Hello:         31
Other:              0      Unknown:           0
```

```
vLAG PDU received:
Role Election:      2      System Info:      1
Peer Instance Enable: 2      Peer Instance Disable: 0
FDB Dynamic Add:    4      FDB Dynamic Del:    4
FDB Inactive Add:   0      FDB Inactive Del:   0
Health Check:       382    ISL Hello:         1
Other:              0      Unknown:           0
```

```
vLAG IGMP packets forwarded:
IGMP Reports:      0
IGMP Leaves:       0
```

```
G8264_2#sh lacp information
port   mode   adminkey   operkey   selected   prio   aggr   trunk   status   minlinks
-----
1      off    1          1         no        32768  --    --      --      1
--
16     off    16         16        no        32768  --    --      --      1
17     active 117        117       yes       32768  17    65     up      1
18     active 117        117       yes       32768  17    65     up      1
19     off    19         19        no        32768  --    --      --      1
--
62     off    62         62        no        32768  --    --      --      1
63     active 163        163       yes       32768  64    66     up      1
64     active 163        163       yes       32768  64    66     up      1
```

```
G8264_2#show lacp aggregator 64
Aggregator Id 64
-----
Aggregator MAC address - fc:cf:62:9d:67:a0
```

```

Actor System Priority - 32768
Actor System ID      - 08:17:f4:c3:dd:ff
Individual           - FALSE
Actor Oper Key       - 163
Partner System Priority - 32768
Partner System ID    - 08:17:f4:76:78:00
Partner Oper Key     - 121
ready                - TRUE
Min-Links            - 1
Number of Ports in aggr - 2
index 0    port 63
index 1    port 64

```

G8264_2#sh vlag information

```

vLAG Tier ID: 256
vLAG system MAC: 08:17:f4:c3:dd:ff
Local MAC fc:cf:62:9d:67:00 Priority 0 Admin Role SECONDARY (Operational Role SECONDARY)
Peer MAC 08:17:f4:32:bb:00 Priority 0
Health local 192.168.240.50 peer 192.168.240.40 State UP
ISL trunk id 65
ISL state Up
Startup Delay Interval: 120s (Finished)

```

vLAG 65: config with admin key 163, associated trunk 66, state formed

G8264_2#sh vlag adminkey 163

```

vLAG is enabled on admin key 163
Current LACP params for 63: active, Priority 32768, Admin Key 163, Min-Links 1

Current LACP params for 64: active, Priority 32768, Admin Key 163, Min-Links 1

```

G8264_2#sh vlag isl

ISL_ID	ISL_Vlan	ISL_Trunk	ISL_Members	Link_State	Trunk_State
65	4094	Adminkey 117	17	UP	UP
			18	UP	UP

G8264_2#sh vlag statistics

vLAG PDU sent:

Role Election:	2	System Info:	1
Peer Instance Enable:	2	Peer Instance Disable:	0
FDB Dynamic Add:	4	FDB Dynamic Del:	4
FDB Inactive Add:	0	FDB Inactive Del:	0
Health Check:	530	ISL Hello:	2
Other:	0	Unknown:	0

vLAG PDU received:

Role Election:	2	System Info:	1
Peer Instance Enable:	2	Peer Instance Disable:	0
FDB Dynamic Add:	4	FDB Dynamic Del:	4
FDB Inactive Add:	0	FDB Inactive Del:	0
Health Check:	529	ISL Hello:	31
Other:	0	Unknown:	0

vLAG IGMP packets forwarded:

IGMP Reports:	0
IGMP Leaves:	0

```
Flex#sh lacp information
port    mode    adminkey  operkey  selected  prio  aggr  trunk  status  minlinks
-----
INTA1   off      1         1        no        32768 --    --    --      1
--
INTB14  off      28        28       no        32768 --    --    --      1
EXT1    off      29        29       no        32768 --    --    --      1
--
EXT20   off      48        48       no        32768 --    --    --      1
EXT21   active   121       121      yes        32768 52    53    up      1
EXT22   active   121       121      yes        32768 52    53    up      1
EXT23   active   121       121      yes        32768 52    53    up      1
EXT24   active   121       121      yes        32768 52    53    up      1
```

```
Flex#sh lacp
Current LACP system ID: 08:17:f4:76:78:00
Current LACP system Priority: 32768
Current LACP timeout scale: long
```

Current LACP params for EXT21: active, Priority 32768, Admin Key 121, Min-Links 1

Current LACP params for EXT22: active, Priority 32768, Admin Key 121, Min-Links 1

Current LACP params for EXT23: active, Priority 32768, Admin Key 121, Min-Links 1

Current LACP params for EXT24: active, Priority 32768, Admin Key 121, Min-Links 1

```
Flex#sh lacp aggregator 52
Aggregator Id 52
-----
Aggregator MAC address - 08:17:f4:76:78:86
Actor System Priority - 32768
Actor System ID - 08:17:f4:76:78:00
Individual - FALSE
Actor Oper Key - 121
Partner System Priority - 32768
Partner System ID - 08:17:f4:c3:dd:ff
Partner Oper Key - 163
ready - TRUE
Min-Links - 1
Number of Ports in aggr - 4
index 0 port EXT21
index 1 port EXT22
index 2 port EXT23
index 3 port EXT24
```

The Flex switch now has one aggregated link (port channel) consisting of four connections to the logically unified pair of IBM G8264 switches. Previously, the Flex switch featured two aggregated links that consisted of two connections each to two separate IBM G8264.

The MST spanning tree is still configured. In contrast to the configurations without VLAG, all four ports now are in spanning tree status forwarding because they all belong to the same LCAP channel.

3.5.6 Show running-config of all switches in Use Case 4

The following configuration memory dumps of the three switches show the successfully tested setup. The essential parameters for this use case are highlighted in red.

The commands that were run on the three switches produced the following outputs:

- ▶ EN2029: Example 3-32
- ▶ G8264 switch 1: Example 3-33 on page 73
- ▶ G8264 switch 2: Example 3-34 on page 75

Example 3-32 Output of the show running command: EN2092 switch

```
Flex#sh run
Current configuration:
!
version "7.2.2.2"
switch-type "IBM Flex System EN2092 1Gb Ethernet Scalable Switch"
!
!

snmp-server user 4 name "DirectorServerSNMPv3User"
snmp-server user 4 authentication-protocol sha authentication-password
"602e911d40088008ac26f2f683b823fa38bbdaca61af87e7367acc3d627979a016507d179fd43edc664137aa7e
2b40f63d"
snmp-server user 4 privacy-protocol des privacy-password
"7f068e355a008a20b62ee7f699b029d28afa8626040f6b48106531c7dcf753ad33117273b4a73403720bee4701
1b065f9c"
!
snmp-server group 4 user-name DirectorServerSNMPv3User
snmp-server group 4 group-name "ibmd_grp_4"
!
snmp-server access 4 name "ibmd_grp_4"
snmp-server access 4 level authPriv
snmp-server access 4 notify-view "iso"
!
snmp-server target-address 1 name "ibmd_taddr_1" address 192.168.10.103
snmp-server target-address 1 parameters-name "ibmd_tparam_1"
!
snmp-server target-parameters 1 name "ibmd_tparam_1"
snmp-server target-parameters 1 user-name "DirectorServerSNMPv3User"
snmp-server target-parameters 1 level authPriv
!
snmp-server version v1v2v3
!
snmp-server name "Flex"
!
hostname "Flex"
system idle 60
!
!
access http enable
access telnet enable
!
interface port INTA2
    tagging
    exit
!
interface port INTA7
    shutdown
    exit
```

```

!
interface port EXT4
    name "TEST_PC"
    tagging
    exit
!
interface port EXT21
    name "T0_G8264_2_Port63"
    tagging
    pvid 10
    exit
!
interface port EXT22
    name "T0_G8264_1_Port63"
    tagging
    pvid 10
    exit
!
interface port EXT23
    name "T0_G8264_1_Port64"
    tagging
    pvid 10
    exit
!
interface port EXT24
    name "T0_G8264_2_Port64"
    tagging
    pvid 10
    exit
!
vlan 1
    member INTA1-EXT20
    no member EXT21-EXT24
!
!
vlan 10
    enable
    name "Server"
    member EXT4,EXT21-EXT24
!
!
vlan 20
    enable
    name "Data20"
    member EXT4,EXT21-EXT24
!
!
vlan 30
    enable
    name "Data30"
    member EXT4,EXT21-EXT24
!
!
vlan 40
    enable
    name "Data40"
    member EXT4,EXT21-EXT24
!
!
!

```

```

spanning-tree mstp version 10
spanning-tree mstp name "PureFlex"
spanning-tree mode mst
spanning-tree mstp cist-add-vlan 1
spanning-tree mstp cist-add-vlan 4095
!
spanning-tree stp 1 vlan 10
spanning-tree stp 1 vlan 30

spanning-tree stp 2 vlan 20
spanning-tree stp 2 vlan 40

!
interface port EXT21
    lacp mode active
    lacp key 121
!
interface port EXT22
    lacp mode active
    lacp key 121
!
interface port EXT23
    lacp mode active
    lacp key 121
!
interface port EXT24
    lacp mode active
    lacp key 121
!
!
!
!
!
!
lldp enable
!
!
!
!
!
ntp enable
ntp ipv6 primary-server fe80::211:25ff:fec3:1420 MGT
ntp interval 15
ntp authenticate
ntp primary-key 49909
!
ntp message-digest-key 103 md5-ekey
4264b3504204a200ae2df2b381b401f2d384e6827376b623d79c78c89f3b4288a2619aa3f05c0d5dc8a369a956a
81063a4203a5a34993a54288393f9264b42da
!
! SNIP
! ...more lines of "ntp message-digest-key"
! SNIP
!
ntp message-digest-key 64248 md5-ekey
f42d0519500d0008bc24e6f293bda3fadbbbc2899f01c55d586637020e1f9dd332028f2e1b627438abbd5bbe8350
5dc965b43752daacb2751446c122610608374
!
ntp trusted-key
103,1821,2416,3343,4617,6903,7255,9094,10386,10939,12266,12389,13261,13280,13640,14424,1641

```



```
7,17555,17944,18537,19291,19742,19776,20027,21166,21710,22141,22512,23917,25162,25988,27418
,27687,27964,28200,29005,29180,29297,29395,31615,31972,32287,32782,34183,35544,35571,37155,
37414,37968,38424,38865,38947,39752,40976,41343,41997,42080,42261,42816,42898,43020,48745,4
9909,50872,51266,54111,54278,55616,57966,61370,62043,62789,63696,63785,64175,64248
!
end
```

Example 3-33 Output of the show running command: G8264 switch 1

```
G8264_1#sh run
version "7.2.2"
switch-type "IBM Networking Operating System RackSwitch G8264"
!
!

!
!
no system dhcp
hostname "G8264_1"
system idle 60
!
!
interface port 17
    name "CrossLink"
    tagging
    pvid 4094
    exit
!
interface port 18
    name "CrossLink"
    tagging
    pvid 4094
    exit
!
interface port 63
    name "DOWNLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
interface port 64
    name "DOWNLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
vlan 1
    member 1-16,19-62
    no member 17-18,63-64
!
!
vlan 10
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 20
```

```

        enable
        name "none"
        member 17-18,63-64
    !
    !
vlan 30
    enable
    name "none"
    member 17-18,63-64
    !
    !
vlan 40
    enable
    name "VLAN 40"
    member 17-18,63-64
    !
    !
vlan 4094
    enable
    name "VLAG_ISL"
    member 17-18
    !
    !
    !
spanning-tree mstp version 10
spanning-tree mstp name "PureFlex"
spanning-tree mode mst
spanning-tree mstp cist-add-vlan 1
spanning-tree mstp cist-add-vlan 4095
!
spanning-tree stp 1 bridge priority 24576
spanning-tree stp 1 vlan 10
spanning-tree stp 1 vlan 30

spanning-tree stp 2 bridge priority 28672
spanning-tree stp 2 vlan 20
spanning-tree stp 2 vlan 40

no spanning-tree stp 32 enable
spanning-tree stp 32 vlan 4094

!
interface port 17
    lacp mode active
    lacp key 117
!
interface port 18
    lacp mode active
    lacp key 117
!
interface port 63
    lacp mode active
    lacp key 163
!
interface port 64
    lacp mode active
    lacp key 163
!
!

```

Example 3-34 Output of the show running command: G8264 switch 2

Chapter 3. IBM RackSwitch G8264 connectivity **75**

```

!
interface port 63
    name "DOWNLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
interface port 64
    name "DOWNLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
vlan 1
    member 1-16,19-62
    no member 17-18,63-64
!
!
vlan 10
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 20
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 30
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 40
    enable
    name "VLAN 40"
    member 17-18,63-64
!
!
vlan 4094
    enable
    name "VLAG_ISL"
    member 17-18
!
!
!
spanning-tree mstp version 10
spanning-tree mstp name "PureFlex"
spanning-tree mode mst
spanning-tree mstp cist-add-vlan 1
spanning-tree mstp cist-add-vlan 4095
!
spanning-tree stp 1 bridge priority 28672
spanning-tree stp 1 vlan 10
spanning-tree stp 1 vlan 30

spanning-tree stp 2 bridge priority 24576

```

```

spanning-tree stp 2 vlan 20
spanning-tree stp 2 vlan 40

no spanning-tree stp 32 enable
spanning-tree stp 32 vlan 4094

!
interface port 17
    lacp mode active
    lacp key 117
!
interface port 18
    lacp mode active
    lacp key 117
!
interface port 63
    lacp mode active
    lacp key 163
!
interface port 64
    lacp mode active
    lacp key 163
!
!
!
vlag enable
vlag tier-id 256
vlag isl vlan 4094
vlag hlthchk peer-ip 192.168.240.40
vlag isl adminkey 117
vlag adminkey 163 enable
!
!
!
!
!
!
!
!
!
!
lldp enable
!
interface ip 128
    ip address 192.168.240.50
    enable
    exit
!
ip gateway 4 address 192.168.240.1
ip gateway 4 enable
!
!
!
!
!
!
end

```

3.6 Use Case 5: Link aggregation and VLAG without STP

The concept of virtual link aggregation (VLAG) shows the pair of G8264 switch logically as one switch entity. Together with LACP, this configuration allows the typical triangle design to be run, as shown in Figure 3-9, without spanning tree.

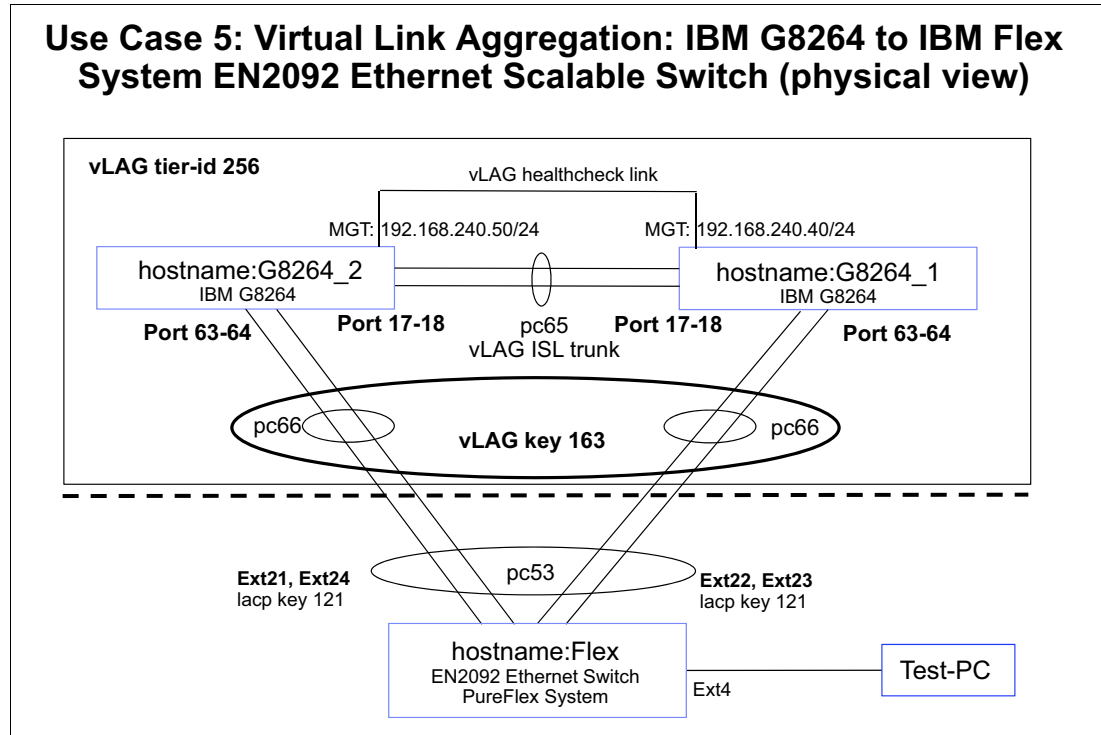


Figure 3-9 Use Case 5

3.6.1 Verifying the topology by using lldp

To verify the topology, we used the **show lldp remote-device** command on the switches, as shown in Example 3-35.

Example 3-35 Verifying the topology by using lldp

```
G8264_1#show lldp remote-device
LLDP Remote Devices Information
```

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
17	1	fc cf 62 9d 67 00	17	G8264_2
63	2	08 17 f4 76 78 00	50	Flex
18	3	fc cf 62 9d 67 00	18	G8264_2
64	4	08 17 f4 76 78 00	51	Flex

```
G8264_2#sh lldp remote-device
LLDP Remote Devices Information
```

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
17	1	08 17 f4 32 bb 00	17	G8264_1
63	2	08 17 f4 76 78 00	49	Flex
18	3	08 17 f4 32 bb 00	18	G8264_1
64	4	08 17 f4 76 78 00	52	Flex

```
Flex#sh lldp remote-device
LLDP Remote Devices Information
```

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
EXT22	1	08 17 f4 32 bb 00	63	G8264_1
EXT21	2	fc cf 62 9d 67 00	63	G8264_2
INTA1	3	5c f3 fc 5f 43 9d	5c-f3-fc-5f-43-9d	
INTA4	4	5c f3 fc 6e 23 41	5c-f3-fc-6e-23-41	
EXT23	5	08 17 f4 32 bb 00	64	G8264_1
EXT5	6	00 0d ec a3 8f 81	mgmt0	vie
EXT7	7	00 05 9b 7b 84 01	mgmt0	str
EXT24	8	fc cf 62 9d 67 00	64	G8264_2

3.6.2 Verify interface status

To verify the interface, we used the **show interface status** command on the switches, as shown in Example 3-36.

Example 3-36 Verify interface status

```
G8264_1#sh int status
```

Alias	Port	Speed	Duplex	Flow Ctrl	Link	Name
				--TX--RX--		
1	1	10000	full	no no	down	1
--						
16	16	10000	full	no no	down	16
17	17	10000	full	no no	up	CrossLink
18	18	10000	full	no no	up	CrossLink
19	19	1G/10G	full	no no	down	19
--						

62	62	1G/10G	full	no	no	down	62
63	63	10000	full	no	no	up	UPLINK_TO_FLEX
64	64	10000	full	no	no	up	UPLINK_TO_FLEX
MGT	65	1000	full	yes	yes	up	MGT

G8264_2#show interface status

Alias	Port	Speed	Duplex	Flow Ctrl		Link	Name
-----	----	-----	-----	--TX--	--RX--	-----	-----
1	1	40000	full	no	no	down	1
--							
16	16	10000	full	no	no	down	16
17	17	10000	full	no	no	up	CrossLink
18	18	10000	full	no	no	up	CrossLink
19	19	1G/10G	full	no	no	down	19
--							
62	62	1G/10G	full	no	no	down	62
63	63	10000	full	no	no	up	UPLINK_TO_FLEX
64	64	10000	full	no	no	up	UPLINK_TO_FLEX
MGT	65	1000	full	yes	yes	up	MGT

Flex#sh interface status

Alias	Port	Speed	Duplex	Flow Ctrl		Link	Name
-----	----	-----	-----	--TX--	--RX--	-----	-----
INTA1	1	1000	full	yes	yes	up	INTA1
INTA2	2	1000	full	yes	yes	up	INTA2
INTA3	3	1000	full	yes	yes	down	INTA3
INTA4	4	1000	full	yes	yes	up	INTA4
INTA5	5	1000	full	yes	yes	down	INTA5
INTA6	6	1000	full	yes	yes	down	INTA6
INTA7	7	1000	full	yes	yes	disabled	INTA7
INTA8	8	1000	full	yes	yes	down	INTA8
INTA9	9	1000	full	yes	yes	down	INTA9
INTA10	10	1000	full	yes	yes	down	INTA10
INTA11	11	1000	full	yes	yes	down	INTA11
INTA12	12	1000	full	yes	yes	down	INTA12
INTA13	13	1000	full	yes	yes	down	INTA13
INTA14	14	1000	full	yes	yes	down	INTA14
INTB1	15	1000	full	yes	yes	down	INTB1
--							
INTB14	28	1000	full	yes	yes	down	INTB14
EXT1	29	1000	full	no	no	up	EXT1
EXT2	30	1000	full	no	no	up	EXT2
EXT3	31	1000	full	no	no	up	EXT3
EXT4	32	1000	full	no	no	up	TEST_PC
EXT5	33	1000	full	no	no	up	EXT5
--							
EXT20	48	any	any	no	no	down	EXT20
EXT21	49	10000	full	no	no	up	T0_G8264_2_Port63
EXT22	50	10000	full	no	no	up	T0_G8264_1_Port63
EXT23	51	10000	full	no	no	up	T0_G8264_1_Port64
EXT24	52	10000	full	no	no	up	T0_G8264_2_Port64
MGT1	53	1000	full	no	no	up	MGT1

3.6.3 Verifying trunks

To verify which VLANs are active on which trunk, we used the **show interface trunk** command on the switches, as shown in Example 3-37.

Example 3-37 Verifying trunks

G8264_1#sh interface trunk									
Alias	Port	Tag	RMON	Lrn	Fld	PVID	NAME	VLAN(s)	
1	1	n	d	e	e	1		1	
--									
16	16	n	d	e	e	1		1	
17	17	y	d	d	e	4094	CrossLink	10 20 30 40	4094
18	18	y	d	d	e	4094	CrossLink	10 20 30 40	4094
19	19	n	d	e	e	1		1	
--									
62	62	n	d	e	e	1		1	
63	63	y	d	e	e	10	UPLINK_TO_FLEX	10 20 30 40	
64	64	y	d	e	e	10	UPLINK_TO_FLEX	10 20 30 40	
MGT	65	n	d	e	e	4095		4095	

G8264_2#sh interface trunk									
Alias	Port	Tag	RMON	Lrn	Fld	PVID	NAME	VLAN(s)	
1	1	n	d	e	e	1		1	
--									
16	16	n	d	e	e	1		1	
17	17	y	d	d	e	4094	CrossLink	10 20 30 40	4094
18	18	y	d	d	e	4094	CrossLink	10 20 30 40	4094
19	19	n	d	e	e	1		1	
--									
62	62	n	d	e	e	1		1	
63	63	y	d	e	e	10	UPLINK_TO_FLEX	10 20 30 40	
64	64	y	d	e	e	10	UPLINK_TO_FLEX	10 20 30 40	
MGT	65	n	d	e	e	4095		4095	

Flex#sh interface trunk									
Alias	Port	Tag	RMON	Lrn	Fld	PVID	NAME	VLAN(s)	
INTA1	1	n	d	e	e	1	INTA1	1	
INTA2	2	y	d	e	e	1	INTA2	1	
INTA3	3	n	d	e	e	1	INTA3	1	
INTA4	4	n	d	e	e	1	INTA4	1	
INTA5	5	n	d	e	e	1	INTA5	1	
INTA6	6	n	d	e	e	1	INTA6	1	
INTA7	7	n	d	e	e	1	INTA7	1	
INTA8	8	n	d	e	e	1	INTA8	1	
INTA9	9	n	d	e	e	1	INTA9	1	
INTA10	10	n	d	e	e	1	INTA10	1	
INTA11	11	n	d	e	e	1	INTA11	1	
INTA12	12	n	d	e	e	1	INTA12	1	
INTA13	13	n	d	e	e	1	INTA13	1	
INTA14	14	n	d	e	e	1	INTA14	1	
INTB1	15	n	d	e	e	1	INTB1	1	
--									
INTB14	28	n	d	e	e	1	INTB14	1	
EXT1	29	n	d	e	e	1	EXT1	1	
EXT2	30	n	d	e	e	1	EXT2	1	
EXT3	31	n	d	e	e	1	EXT3	1	

EXT4	32	y	d	e	e	1	TEST_PC	1	10	20	30	40
EXT5	33	n	d	e	e	1	EXT5	1				
--												
EXT20	48	n	d	e	e	1	EXT20	1				
EXT21	49	y	d	e	e	10	TO_G8264_2_Port63	10	20	30	40	
EXT22	50	y	d	e	e	10	TO_G8264_1_Port63	10	20	30	40	
EXT23	51	y	d	e	e	10	TO_G8264_1_Port64	10	20	30	40	
EXT24	52	y	d	e	e	10	TO_G8264_2_Port64	10	20	30	40	
MGT1	53	y	d	e	e	4095	MGT1		4095			

3.6.4 Verify virtual link aggregation

We verified the link aggregation configuration of the switches by executing various **show** commands, as shown in Example 3-38.

Example 3-38 Verify virtual link aggregation

```
G8264_1#show lacp information
port    mode    adminkey  operkey  selected  prio  aggr  trunk  status  minlinks
-----
1        off      1         1        no        32768 --    --    --      1
--
16       off      16        16       no        32768 --    --    --      1
17       active   117       117      yes       32768 17    65    up      1
18       active   117       117      yes       32768 17    65    up      1
19       off      19        19       no        32768 --    --    --      1
--
62       off      62        62       no        32768 --    --    --      1
63       active   163       163      yes       32768 63    66    up      1
64       active   163       163      yes       32768 63    66    up      1

G8264_1#sh lacp aggregator 63
Aggregator Id 63
-----
Aggregator MAC address - 08:17:f4:32:bb:a0
Actor System Priority   - 32768
Actor System ID        - 08:17:f4:c3:dd:ff
Individual              - FALSE
Actor Oper Key          - 163
Partner System Priority - 32768
Partner System ID      - 08:17:f4:76:78:00
Partner Oper Key        - 121
ready                   - TRUE
Min-Links               - 1
Number of Ports in aggr - 2
index 0    port 63
index 1    port 64

G8264_1#show spanning-tree
Spanning Tree is shut down.

G8264_1#sh vlag
vLAG status: enabled
vLAG Tier ID: 256
vLAG system MAC: 08:17:f4:c3:dd:ff
Local Priority: 0
ISL Information: VLAN 4094, Trunk 0, LACP Key 117
Health check Peer IP Address: 192.168.240.50
Health check connection retry interval: 30 seconds
```

Health check number of keepalive attempts: 3
 Health check keepalive interval: 5 seconds
 vLAG startup delay interval: 120 seconds
 Current LACP system ID: 08:17:f4:32:bb:00
 Current LACP system Priority: 32768
 Current LACP timeout scale: long

vLAG 65 : active
 Current LACP params for 63: active, Priority 32768, Admin Key 163, Min-Links 1
 Current LACP params for 64: active, Priority 32768, Admin Key 163, Min-Links 1

G8264_1#sh vlag information
 vLAG Tier ID: 256
 vLAG system MAC: 08:17:f4:c3:dd:ff
 Local MAC 08:17:f4:32:bb:00 Priority 0 Admin Role PRIMARY (Operational Role SECONDARY)
 Peer MAC fc:cf:62:9d:67:00 Priority 0
 Health local 192.168.240.40 peer 192.168.240.50 State UP
 ISL trunk id 65
 ISL state Up
 Startup Delay Interval: 120s (Finished)

vLAG 65: config with admin key 163, associated trunk 66, state formed

G8264_1#sh vlag isl

ISL_ID	ISL_Vlan	ISL_Trunk	ISL_Members	Link_State	Trunk_State
65	4094	Adminkey 117	17	UP	UP
			18	UP	UP

G8264_1#sh vlag statistics

vLAG PDU sent:			
Role Election:	6	System Info:	50
Peer Instance Enable:	5	Peer Instance Disable:	0
FDB Dynamic Add:	12	FDB Dynamic Del:	15
FDB Inactive Add:	0	FDB Inactive Del:	0
Health Check:	3392	ISL Hello:	292
Other:	0	Unknown:	0

vLAG PDU received:			
Role Election:	5	System Info:	3
Peer Instance Enable:	6	Peer Instance Disable:	0
FDB Dynamic Add:	12	FDB Dynamic Del:	12
FDB Inactive Add:	0	FDB Inactive Del:	0
Health Check:	3387	ISL Hello:	231
Other:	0	Unknown:	0

vLAG IGMP packets forwarded:
 IGMP Reports: 0
 IGMP Leaves: 0

G8264_2#sh lacp information

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
1	off	1	1	no	32768	--	--	--	1
--									
16	off	16	16	no	32768	--	--	--	1
17	active	117	117	yes	32768	17	65	up	1
18	active	117	117	yes	32768	17	65	up	1
19	off	19	19	no	32768	--	--	--	1
--									
62	off	62	62	no	32768	--	--	--	1
63	active	163	163	yes	32768	64	66	up	1
64	active	163	163	yes	32768	64	66	up	1

G8264_2#sh lacp aggregator 64

Aggregator Id 64

```

-----
Aggregator MAC address - fc:cf:62:9d:67:a0
Actor System Priority   - 32768
Actor System ID        - 08:17:f4:c3:dd:ff
Individual              - FALSE
Actor Oper Key          - 163
Partner System Priority - 32768
Partner System ID      - 08:17:f4:76:78:00
Partner Oper Key        - 121
ready                   - TRUE
Min-Links               - 1
Number of Ports in aggr - 2
index 0    port 63
index 1    port 64

```

G8264_2#sh vlag

```

vLAG status: enabled
vLAG Tier ID: 256
vLAG system MAC: 08:17:f4:c3:dd:ff
Local Priority: 0
ISL Information: VLAN 4094, Trunk 0, LACP Key 117
Health check Peer IP Address: 192.168.240.40
Health check connection retry interval: 30 seconds
Health check number of keepalive attempts: 3
Health check keepalive interval: 5 seconds
vLAG startup delay interval: 120 seconds
Current LACP system ID: fc:cf:62:9d:67:00
Current LACP system Priority: 32768
Current LACP timeout scale: long

```

vLAG 65 : active

Current LACP params for 63: active, Priority 32768, Admin Key 163, Min-Links 1

Current LACP params for 64: active, Priority 32768, Admin Key 163, Min-Links 1

G8264_2#sh vlag information

```

vLAG Tier ID: 256
vLAG system MAC: 08:17:f4:c3:dd:ff
Local MAC fc:cf:62:9d:67:00 Priority 0 Admin Role SECONDARY (Operational Role PRIMARY)
Peer MAC 08:17:f4:32:bb:00 Priority 0
Health local 192.168.240.50 peer 192.168.240.40 State UP
ISL trunk id 65
ISL state Up
Startup Delay Interval: 120s (Finished)

```

vLAG 65: config with admin key 163, associated trunk 66, state formed

G8264_2#sh vlag adminkey 163

vLAG is enabled on admin key 163

Current LACP params for 63: active, Priority 32768, Admin Key 163, Min-Links 1

Current LACP params for 64: active, Priority 32768, Admin Key 163, Min-Links 1

G8264_2#sh vlag isl

ISL_ID	ISL_Vlan	ISL_Trunk	ISL_Members	Link_State	Trunk_State
65	4094	Adminkey 117	17 18	UP UP	UP UP

G8264_2#sh vlag statistics

vLAG PDU sent:

Role Election:	5	System Info:	14
Peer Instance Enable:	6	Peer Instance Disable:	0
FDB Dynamic Add:	12	FDB Dynamic Del:	12
FDB Inactive Add:	0	FDB Inactive Del:	0
Health Check:	3546	ISL Hello:	321
Other:	0	Unknown:	0

vLAG PDU received:

Role Election:	6	System Info:	3
Peer Instance Enable:	5	Peer Instance Disable:	0
FDB Dynamic Add:	12	FDB Dynamic Del:	15
FDB Inactive Add:	0	FDB Inactive Del:	0
Health Check:	3540	ISL Hello:	321
Other:	0	Unknown:	0

vLAG IGMP packets forwarded:

IGMP Reports:	0
IGMP Leaves:	0

Flex#show lacp information

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks
INTA1	off	1	1	no	32768	--	--	--	1
--									
INTB14	off	28	28	no	32768	--	--	--	1
EXT1	off	29	29	no	32768	--	--	--	1
--									
EXT20	off	48	48	no	32768	--	--	--	1
EXT21	active	121	121	yes	32768	52	53	up	1
EXT22	active	121	121	yes	32768	52	53	up	1
EXT23	active	121	121	yes	32768	52	53	up	1
EXT24	active	121	121	yes	32768	52	53	up	1

Flex#sh lacp aggregator 52

Aggregator Id 52

Aggregator MAC address - 08:17:f4:76:78:86
Actor System Priority - 32768
Actor System ID - 08:17:f4:76:78:00
Individual - FALSE

```
Actor Oper Key          - 121
Partner System Priority - 32768
Partner System ID       - 08:17:f4:c3:dd:ff
Partner Oper Key        - 163
ready                   - TRUE
Min-Links                - 1
Number of Ports in aggr - 4
index 0    port EXT21
index 1    port EXT22
index 2    port EXT23
index 3    port EXT24
```

The Flex System switch now has one aggregated link (port channel) consisting of four connections to the logically unified pair of IBM G8264 switches. Previously, the Flex System switch featured two aggregated links that consisted of two connections each to two separate IBM G8264.

3.6.5 Show running-config of all switches in Use Case 5

The following configuration memory dumps of the IBM Flex Switch and both IBM System Network switches show the successfully tested setup. The essential parameters for this use case are highlighted in red.

The commands that were run on the three switches produced the following outputs:

- ▶ EN2029: Example 3-39
- ▶ G8264 switch 1: Example 3-40 on page 89
- ▶ G8264 switch 2: Example 3-41 on page 92

Example 3-39 Output of the show running command: EN2092

```
Flex#sh run
Current configuration:
!
version "7.2.2.2"
switch-type "IBM Flex System EN2092 1Gb Ethernet Scalable Switch"
!
!

snmp-server user 4 name "DirectorServerSNMPv3User"
snmp-server user 4 authentication-protocol sha authentication-password
"448edc340000882085a7b7f7c3b02bd2f0520e931ea46bc5b7eded9972fe8261a0ef96428215042c04724d220
c902acd9"
snmp-server user 4 privacy-protocol des privacy-password
"453edd840110888084b7b6e7c2a02b7269f0ab694f0b3fefcd1dc2cefc9b2755a977e48dfffb7f2c02ae685e8fd
38cfc425"
!
snmp-server group 4 user-name DirectorServerSNMPv3User
snmp-server group 4 group-name "ibmd_grp_4"
!
snmp-server access 4 name "ibmd_grp_4"
snmp-server access 4 level authPriv
snmp-server access 4 notify-view "iso"
!
snmp-server target-address 1 name "ibmd_taddr_1" address 192.168.10.103
snmp-server target-address 1 parameters-name "ibmd_tparam_1"
!
snmp-server target-parameters 1 name "ibmd_tparam_1"
snmp-server target-parameters 1 user-name "DirectorServerSNMPv3User"
```

```

snmp-server target-parameters 1 level authPriv
!
snmp-server version v1v2v3
!
snmp-server name "Flex"
!
hostname "Flex"
system idle 60
!
!
access http enable
access telnet enable
!
interface port INTA2
    tagging
    exit
!
interface port INTA7
    shutdown
    exit
!
interface port EXT4
    name "TEST_PC"
    tagging
    exit
!
interface port EXT21
    name "T0_G8264_2_Port63"
    tagging
    pvid 10
    exit
!
interface port EXT22
    name "T0_G8264_1_Port63"
    tagging
    pvid 10
    exit
!
interface port EXT23
    name "T0_G8264_1_Port64"
    tagging
    pvid 10
    exit
!
interface port EXT24
    name "T0_G8264_2_Port64"
    tagging
    pvid 10
    exit
!
vlan 1
    member INTA1-EXT20
    no member EXT21-EXT24
!
!
vlan 10
    enable
    name "Server"
    member EXT4,EXT21-EXT24
!

```

```

!
vlan 20
    enable
    name "Data20"
    member EXT4,EXT21-EXT24
!
!
vlan 30
    enable
    name "Data30"
    member EXT4,EXT21-EXT24
!
!
vlan 40
    enable
    name "Data40"
    member EXT4,EXT21-EXT24
!
!
!
spanning-tree mstp version 10
spanning-tree mstp name "PureFlex"
spanning-tree mode disable
!
spanning-tree stp 1 vlan 1
spanning-tree stp 1 vlan 10
spanning-tree stp 1 vlan 20
spanning-tree stp 1 vlan 30
spanning-tree stp 1 vlan 40
!
interface port EXT21
    lACP mode active
    lACP key 121
!
interface port EXT22
    lACP mode active
    lACP key 121
!
interface port EXT23
    lACP mode active
    lACP key 121
!
interface port EXT24
    lACP mode active
    lACP key 121
!
!
!
!
!
!
lldp enable
!
!
!
!
!
ntp enable
ntp ipv6 primary-server fe80::211:25ff:fec3:1420 MGT

```



```

ntp interval 15
ntp authenticate
ntp primary-key 49909
!
ntp message-digest-key 103 md5-ekey
0b87933c0300822886a6f2f7c0b021da71fedfcb71dca85400f52051d4db341ddc66d383102dc917aa13d6f2967
b6179f6d9396a95503e6e0217d9f7248c1c3a
!
! SNIP
! ...more lines of "ntp message-digest-key"
! SNIP
!
ntp message-digest-key 64248 md5-ekey
898311380100002884a6f2f3c2b0a3dae66cc6e9326e294b602f8fc11ca24cca6780d1f7d5b707d49f028be5635
b0932ffcf8aa484922018dc0863fb346e37a
!
ntp trusted-key
103,1821,2416,3343,4617,6903,7255,9094,10386,10939,12266,12389,13261,13280,13640,14424,1641
7,17555,17944,18537,19291,19742,19776,20027,21166,21710,22141,22512,23917,25162,25988,27418
,27687,27964,28200,29005,29180,29297,29395,31615,31972,32287,32782,34183,35544,35571,37155,
37414,37968,38424,38865,38947,39752,40976,41343,41997,42080,42261,42816,42898,43020,48745,4
9909,50872,51266,54111,54278,55616,57966,61370,62043,62789,63696,63785,64175,64248
!
end

```

Example 3-40 Output of the show running command: G8264 switch 1

```

G8264_1#sh run
Current configuration:
!
version "7.2.2"
switch-type "IBM Networking Operating System RackSwitch G8264"
!
!

!
!
no system dhcp
hostname "G8264_1"
system idle 60
!
!
interface port 17
    name "CrossLink"
    tagging
    pvid 4094
    exit
!
interface port 18
    name "CrossLink"
    tagging
    pvid 4094
    exit
!
interface port 63
    name "DOWNLINK_TO_FLEX"
    tagging
    pvid 10

```

```

        exit
    !
interface port 64
    name "DOWNLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
vlan 1
    member 1-16,19-62
    no member 17-18,63-64
!
!
vlan 10
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 20
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 30
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 40
    enable
    name "VLAN 40"
    member 17-18,63-64
!
!
vlan 4094
    enable
    name "VLAG_ISL"
    member 17-18
!
!
!
spanning-tree mstp version 10
spanning-tree mstp name "PureFlex"
spanning-tree mode disable
!
spanning-tree stp 1 bridge priority 24576
spanning-tree stp 1 vlan 1
spanning-tree stp 1 vlan 10
spanning-tree stp 1 vlan 20
spanning-tree stp 1 vlan 30
spanning-tree stp 1 vlan 40

spanning-tree stp 2 bridge priority 28672
no spanning-tree stp 32 enable
spanning-tree stp 32 vlan 4094
!

```


Example 3-41 Output of the show running command: G8264 switch 2

```
G8264_2#sh run
Current configuration:
!
version "7.2.2"
switch-type "IBM Networking Operating System RackSwitch G8264"
!
!

!
!
no system dhcp
hostname "G8264_2"
system idle 60
!
!
interface port 17
    name "CrossLink"
    tagging
    pvid 4094
    exit
!
interface port 18
    name "CrossLink"
    tagging
    pvid 4094
    exit
!
interface port 63
    name "DOWNLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
interface port 64
    name "DOWNLINK_TO_FLEX"
    tagging
    pvid 10
    exit
!
vlan 1
    member 1-16,19-62
    no member 17-18,63-64
!
!
vlan 10
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 20
    enable
    name "none"
    member 17-18,63-64
!
!
vlan 30
    enable
    name "none"
```

```

        member 17-18,63-64
    !
    !
vlan 40
    enable
    name "VLAN 40"
    member 17-18,63-64
!
!
vlan 4094
    enable
    name "VLAG_ISL"
    member 17-18
!
!
!
spanning-tree mstp version 10
spanning-tree mstp name "PureFlex"
spanning-tree mode disable
!
spanning-tree stp 1 bridge priority 28672
spanning-tree stp 1 vlan 1
spanning-tree stp 1 vlan 10
spanning-tree stp 1 vlan 20
spanning-tree stp 1 vlan 30
spanning-tree stp 1 vlan 40

spanning-tree stp 2 bridge priority 24576
no spanning-tree stp 32 enable
spanning-tree stp 32 vlan 4094

!
interface port 17
    lacp mode active
    lacp key 117
!
interface port 18
    lacp mode active
    lacp key 117
!
interface port 63
    lacp mode active
    lacp key 163
!
interface port 64
    lacp mode active
    lacp key 163
!
!
!
vlag enable
vlag tier-id 256
vlag isl vlan 4094
vlag hltchk peer-ip 192.168.240.40
vlag isl adminkey 117
vlag adminkey 163 enable
!
!
!
!
```

```
!  
!  
!  
!  
!  
!  
lldp enable  
!  
interface ip 128  
    ip address 192.168.240.50  
    enable  
    exit  
!  
ip gateway 4 address 192.168.240.1  
ip gateway 4 enable  
!  
!  
!  
!  
!  
end
```



Cisco Nexus 5000 connectivity

In this chapter, we describe the process that was used to test the Layer 2 interoperability between Cisco Nexus 5000 Switches and the embedded IBM Flex System switch. The embedded IBM Flex Switch was connected to two Cisco Nexus 5000 switches.

We tested Layer 2 connectivity trunking, channeling (link aggregation), and spanning tree. For trunking, we used 802.1q. For link aggregation, we tested static and LACP. The tested spanning trees were PVRST and MSTP. To show load balancing (even if spanning tree is active), we configured even and odd VLANS. Finally, we tested vPC to activate all of the links.

To verify Layer 2 topology, we used Link Layer Discovery Protocol (LLDP) as the vendor independent protocol.

Important: IBM switches do not support the proprietary Cisco Discovery Protocol (CDP) protocol.

This chapter includes the following topics:

- ▶ Prerequisites
- ▶ Use Case 1: PVRST
- ▶ Use Case 2: PVRST with LACP Channeling
- ▶ Use Case 3: MST with LACP Channeling
- ▶ Use Case 4: MST with LACP Channeling and vPC
- ▶ Use Case 5: LACP Channeling and vPC without spanning tree

4.1 Prerequisites

We started by physically connecting a triangle with two Cisco Nexus 5000 switches and one IBM Systems Networking embedded Flex Switch. We configured four VLANs and set up Per VLAN Rapid Spanning Tree (PVRST). To test connectivity, we used a test PC.

We used the following switches and one PC to test connectivity:

- ▶ One Cisco Nexus 5010 Switch
- ▶ One Cisco Nexus 5020 Switch
- ▶ One IBM Flex System EN2092 1-Gb Ethernet Scalable Switch
- ▶ One test PC

All of the links between the switches are 10 Gigabit Ethernet.

4.2 Use Case 1: PVRST

In our first use case, we used three 10 GE links to connect the switches. We also configured 802.1q trunks and PVRST. For load balancing, odd VLANs 10 and 30, and even VLANs 20 and 40 are used, as shown in Figure 4-1.

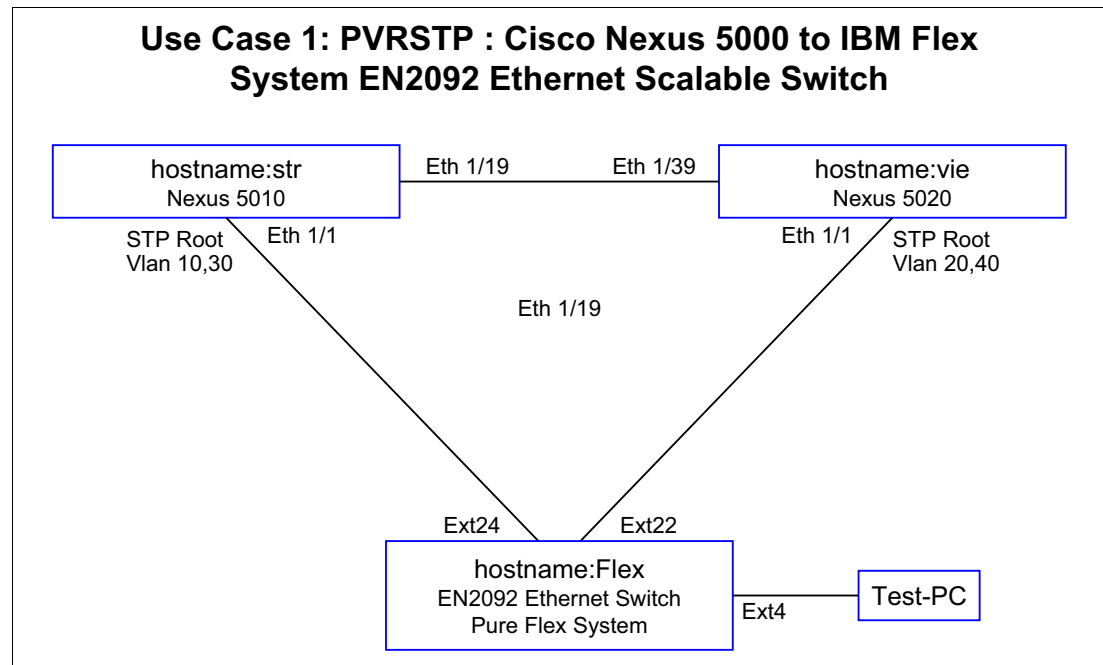


Figure 4-1 Use Case 1

4.2.1 Verifying the topology that is used by using lldp

To verify our configurations, we used several show commands on the IBM and Cisco switches, as shown in Example 4-1 on page 97. The essential parameters for this use case are highlighted in red.

To check the topology, we used the **show lldp remote-device** command on the IBM Flex System switch and the **show lldp neighbors** command on the Cisco Nexus switch. The important parameters and details are highlighted in red.

Example 4-1 Verifying configurations

```
Flex#show lldp remote-device
```

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
EXT22	1	00 0d ec a3 8f 88	Eth1/1	vie
EXT24	2	00 05 9b 7b 84 08	Eth1/1	str
INTA1	3	5c f3 fc 5f 43 9d	5c-f3-fc-5f-43-9d	

!--- Display the LLDP remote devices. Note that you must enable
!--- "feature lldp" on the N5000.
!--- The local Port Numbers of the Pure Flex System Ethernet Switch
!--- distinguish between internal and external Ethernet ports.

```
str# show lldp neighbors
```

Capability codes:

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

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

Device ID	Local Intf	Hold-time	Capability	Port ID
Flex	Eth1/1	120	BR	52
vie	Eth1/19	120	B	Eth1/39

Total entries displayed: 2

!--- The Port named EXT22 at the Pure Flex System Ethernet Switch has the
!--- port ID 52 which is shown in the show lldp neighbors here.

```
vie# show lldp neighbors
```

Capability codes:

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

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

Device ID	Local Intf	Hold-time	Capability	Port ID
Flex	Eth1/1	120	BR	50
str	Eth1/39	120	B	Eth1/19

Total entries displayed: 2

!--- The Port named EXT24 at the Pure Flex System Ethernet Switch has the
!--- port ID 50 which is shown in the show lldp neighbors here.
Port EXT22 = Port ID 50

4.2.2 Verifying trunks

To verify which VLANs are active on which trunk, we used the **show interface trunk** command on the IBM Flex System switch and the Cisco Nexus switch, as shown in Example 4-2. The important parameters and details are highlighted in red.

Example 4-2 Output of show interface trunk command

Flex#**show interface trunk**

Alias	Port	Tag	RMON	Lrn	Fld	PVID	NAME	VLAN(s)
...								
EXT4	32	y	d	e	e	1	TEST_PC	1 10 20 30 40
...								
EXT22	50	y	d	e	e	10	TO_VIE_ETH1/1	10 20 30 40
...								
EXT24	52	y	d	e	e	10	TO_STR_ETH1/1	10 20 30 40

str# **show interface trunk**

Port	Native Vlan	Status	Port Channel
Eth1/1	10	trunking	--
Eth1/2	1	trunking	--
Eth1/19	1	trunking	--
Eth1/20	1	trunking	--

Port	Vlans Allowed on Trunk
Eth1/1	10,20,30,40
Eth1/2	1-3967,4048-4093
Eth1/19	1-3967,4048-4093
Eth1/20	1-3967,4048-4093

Port	Vlans Err-disabled on Trunk
Eth1/1	none
Eth1/2	none
Eth1/19	none
Eth1/20	none

Port	STP Forwarding
Eth1/1	10,20,30,40
Eth1/2	none
Eth1/19	1,10,20,30,40
Eth1/20	none

Port	Vlans in spanning tree forwarding state and not pruned
Eth1/1	--
Eth1/2	--
Eth1/19	--

Eth1/20 --

vie# **show interface trunk**

Port	Native Vlan	Status	Port Channel
Eth1/1	10	trunking	--
Eth1/2	1	trunking	--
Eth1/39	1	trunking	--
Eth1/40	1	trunking	--

Port	Vlans Allowed on Trunk
Eth1/1	10,20,30,40
Eth1/2	1-3967,4048-4093
Eth1/39	1-3967,4048-4093
Eth1/40	1-3967,4048-4093

Port	Vlans Err-disabled on Trunk
Eth1/1	none
Eth1/2	none
Eth1/39	none
Eth1/40	none

Port	STP Forwarding
Eth1/1	10,20,30,40
Eth1/2	none
Eth1/39	1,10,20,30,40
Eth1/40	none

Port	Vlans in spanning tree forwarding state and not pruned
Eth1/1	--
Eth1/2	--
Eth1/39	--
Eth1/40	--

As shown in Figure 4-2 on page 100 and Figure 4-3 on page 100, we have two spanning trees, one for even-numbered VLANs and one for odd-numbered VLANs. By using the show spanning tree command, you can verify the status of the respective Ethernet interface's VLAN, port state, and port role.

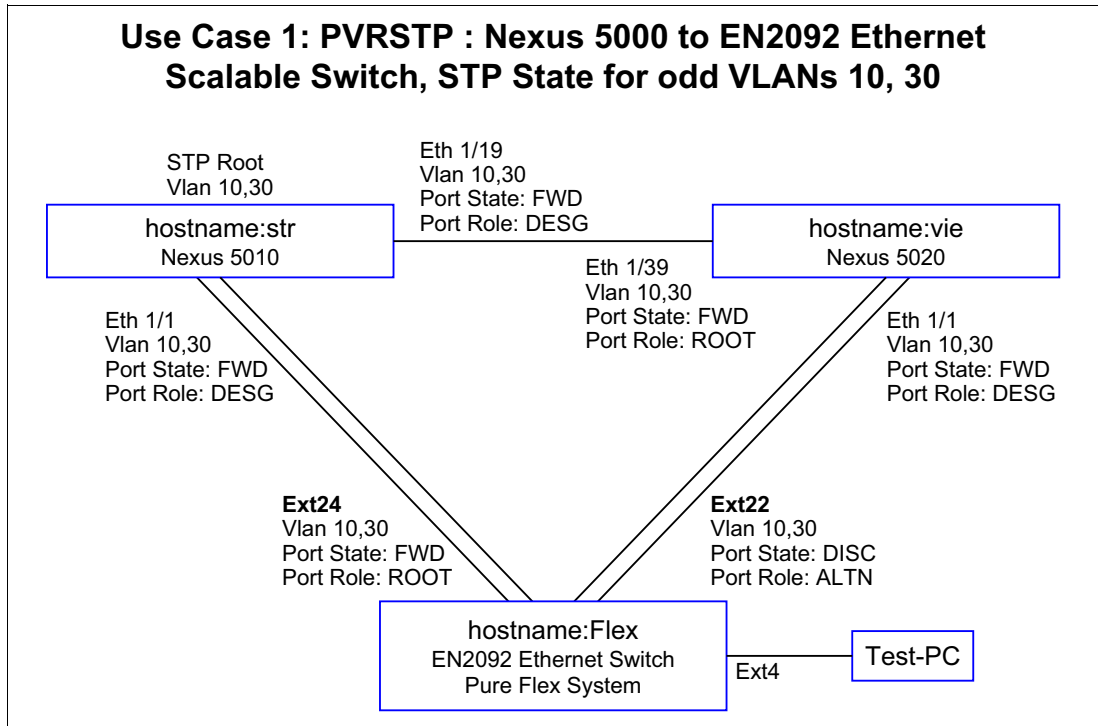


Figure 4-2 Use Case 1: Odd-numbered VLANs

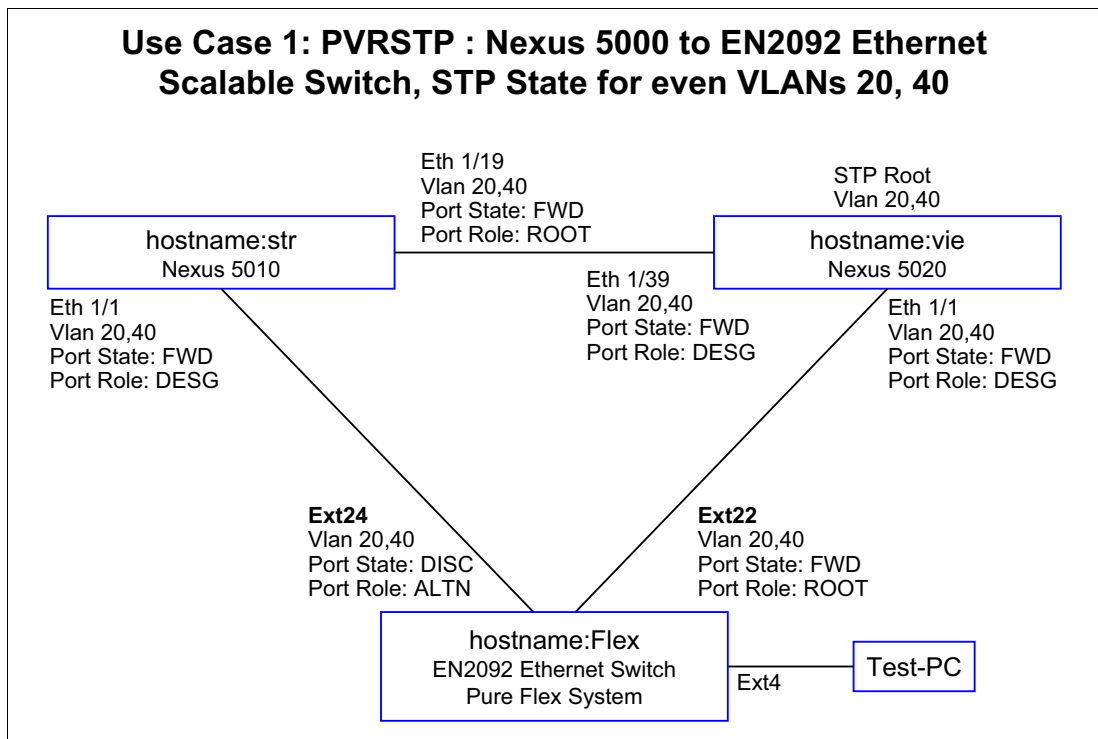


Figure 4-3 Use Case 1: Even-numbered VLANs

In the next step, we verified the PVRST spanning tree configuration of the switches by executing the **show spanning-tree** command.

The commands that were run on the three switches produced the following outputs:

- Flex System EN2029: Example 4-3
- G8264 STR switch: Example 4-4 on page 102
- G8264 VIE switch: Example 4-5 on page 104

Important parameters and details are highlighted in red.

Example 4-3 Outout of show spanning-tree command: Flex System switch

```
Flex#show spanning-tree
-----
Pvst+ compatibility mode enabled
-----
Spanning Tree Group 10: On (PVRST)
VLANs: 10

Current Root:          Path-Cost  Port Hello MaxAge FwdDel
600a 00:05:9b:7b:84:3c    2000   EXT24    2    20    15

!--- Compare the ID of the Root with the LLDP output to identify the root switch.

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
             61450    2    20    15    300             12Press q to

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
-----
EXT4            128    20000!  FWD    DESG f00a-08:17:f4:76:78:00    8020        P2P
EXT22           128    2000!  DISC   ALTN 700a-00:0d:ec:a3:8f:bc    8081        P2P
EXT24           128    2000!  FWD    ROOT 600a-00:05:9b:7b:84:3c    8081        P2P
! = Automatic path cost.

-----
Spanning Tree Group 20: On (PVRST)
VLANs: 20

Current Root:          Path-Cost  Port Hello MaxAge FwdDel
6014 00:0d:ec:a3:8f:bc    2000   EXT22    2    20    15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
             61460    2    20    15    300             1

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
-----
EXT4            128    20000!  DISC   DESG f014-08:17:f4:76:78:00    8020        P2P
EXT22           128    2000!  FWD    ROOT 6014-00:0d:ec:a3:8f:bc    8081        P2P
EXT24           128    2000!  DISC   ALTN 7014-00:05:9b:7b:84:3c    8081        P2P
! = Automatic path cost.

-----
Spanning Tree Group 30: On (PVRST)
VLANs: 30

Current Root:          Path-Cost  Port Hello MaxAge FwdDel
601e 00:05:9b:7b:84:3c    2000   EXT24    2    20    15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
             61470    2    20    15Press q to quit, any other key to cont    300
1

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
```

```

-----
EXT4          128      20000! DISC  DESG f01e-08:17:f4:76:78:00      8020      P2P
EXT22         128      2000!  DISC  ALTN 701e-00:0d:ec:a3:8f:bc      8081      P2P
EXT24         128      2000!  FWD  ROOT 601e-00:05:9b:7b:84:3c      8081      P2P
! = Automatic path cost.

```

```

-----
Spanning Tree Group 40: On (PVRST)

```

VLANs: 40

```

Current Root:          Path-Cost  Port Hello MaxAge FwdDel
6028 00:0d:ec:a3:8f:bc    2000  EXT22   2    20    15

```

```

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
             61480     2     20    15    300              1

```

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	DISC	DESG	f028-08:17:f4:76:78:00	8020	P2P
EXT22	128	2000!	FWD	ROOT	6028-00:0d:ec:a3:8f:bc	8081	P2P
EXT24	128	2000!	DISC	ALTN	7028-00:05:9b:7b:84:3c	8081	P2P

! = Automatic path cost.

```

-----
Spanning Tree Group 128: Off (PVRST), FDB aging timer 300

```

VLANs: 4095

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
MGT1	0	0	FWD *				

* = STP turned off for this port.

Example 4-4 Output of show spanning-tree command: STR switch

```
str# show spanning-tree
```

VLAN0001

```

Spanning tree enabled protocol rstp
Root ID    Priority    32769
Address    0005.9b7b.843c
This bridge is the root
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

```

Bridge ID  Priority    32769 (priority 32768 sys-id-ext 1)
Address    0005.9b7b.843c
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Eth1/19	Desg	FWD	2	128.147	P2p

VLAN0010

```

Spanning tree enabled protocol rstp
Root ID    Priority    24586
Address    0005.9b7b.843c
This bridge is the root
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

!--- Compare the address (ID) of the Root with the LLDP output to identify the root switch
.

```
Bridge ID  Priority    24586 (priority 24576 sys-id-ext 10)
Address    0005.9b7b.843c
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Eth1/1	Desg	FWD	2	128.129	P2p
Eth1/19	Desg	FWD	2	128.147	P2p

VLAN0020

Spanning tree enabled protocol rstp

```
Root ID  Priority    24596
Address  000d.eca3.8fbc
Cost     2
Port     147 (Ethernet1/19)
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID  Priority    28692 (priority 28672 sys-id-ext 20)
Address    0005.9b7b.843c
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Eth1/1	Desg	FWD	2	128.129	P2p
Eth1/19	Root	FWD	2	128.147	P2p

VLAN0030

Spanning tree enabled protocol rstp

```
Root ID  Priority    24606
Address  0005.9b7b.843c
This bridge is the root
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID  Priority    24606 (priority 24576 sys-id-ext 30)
Address    0005.9b7b.843c
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Eth1/1	Desg	FWD	2	128.129	P2p
Eth1/19	Desg	FWD	2	128.147	P2p

VLAN0040

Spanning tree enabled protocol rstp

```
Root ID  Priority    24616
Address  000d.eca3.8fbc
Cost     2
Port     147 (Ethernet1/19)
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID  Priority    28712 (priority 28672 sys-id-ext 40)
Address    0005.9b7b.843c
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Eth1/1	Desg	FWD	2	128.129	P2p
Eth1/19	Root	FWD	2	128.147	P2p

Example 4-5 Output of show spanning-tree command: VIE switch

vie# **show spanning-tree**

VLAN0001

Spanning tree enabled protocol rstp
 Root ID Priority 32769
 Address 0005.9b7b.843c
 Cost 2
 Port 167 (Ethernet1/39)
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 32769 (priority 32768 sys-id-ext 1)
 Address 000d.eca3.8fbc
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Eth1/39	Root	FWD	2	128.167	P2p

VLAN0010

Spanning tree enabled protocol rstp
 Root ID Priority 24586
 Address 0005.9b7b.843c
 Cost 2
 Port 167 (Ethernet1/39)
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 28682 (priority 28672 sys-id-ext 10)
 Address 000d.eca3.8fbc
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Eth1/1	Desg	FWD	2	128.129	P2p
Eth1/39	Root	FWD	2	128.167	P2p

VLAN0020

Spanning tree enabled protocol rstp
 Root ID Priority 24596
 Address 000d.eca3.8fbc
 This bridge is the root
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 24596 (priority 24576 sys-id-ext 20)
 Address 000d.eca3.8fbc
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
-----------	------	-----	------	----------	------


```

Eth1/1          Desg FWD 2      128.129 P2p
Eth1/39         Desg FWD 2      128.167 P2p

VLAN0030
Spanning tree enabled protocol rstp
  Root ID      Priority    24606
                Address    0005.9b7b.843c
                Cost        2
                Port        167 (Ethernet1/39)
                Hello Time  2 sec Max Age 20 sec Forward Delay 15 sec

  Bridge ID    Priority    28702 (priority 28672 sys-id-ext 30)
                Address    000d.eca3.8fbc
                Hello Time  2 sec Max Age 20 sec Forward Delay 15 sec

Interface      Role Sts Cost      Prio.Nbr Type
-----
Eth1/1         Desg FWD 2      128.129 P2p
Eth1/39        Root FWD 2      128.167 P2p

VLAN0040
Spanning tree enabled protocol rstp
  Root ID      Priority    24616
                Address    000d.eca3.8fbc
                This bridge is the root
                Hello Time  2 sec Max Age 20 sec Forward Delay 15 sec

  Bridge ID    Priority    24616 (priority 24576 sys-id-ext 40)
                Address    000d.eca3.8fbc
                Hello Time  2 sec Max Age 20 sec Forward Delay 15 sec

Interface      Role Sts Cost      Prio.Nbr Type
-----
Eth1/1         Desg FWD 2      128.129 P2p
Eth1/39        Desg FWD 2      128.167 P2p

vie#

```

4.2.3 Show running-config of all switches in Use Case 1

In the following configuration print outs of the IBM Flex System switch and the Cisco Nexus switches, you can comprehend the necessary configuration steps we did during our test. Important parameters and detail are highlighted in red.

Important: Sections of the configuration output in Example 4-6 on page 106, Example 4-7 on page 108, and Example 4-8 on page 109 were removed to highlight the important parts of the outputs. The omissions are indicated by “...”.

The commands that were run on the three switches produced the following outputs:

- ▶ Flex System EN2029: Example 4-6 on page 106
- ▶ G8264 STR switch: Example 4-7 on page 108
- ▶ G8264 VIE switch: Example 4-8 on page 109

Important parameters and details are highlighted in red.

Example 4-6 Output of show running-config command: Flex System switch

```
Flex# show running-config
!
version "7.2.2.2"
switch-type "IBM Flex System EN2092 1Gb Ethernet Scalable Switch"
!
!

...
hostname "Flex"
system idle 60
!
!
access http enable
access telnet enable
!
...
interface port EXT4
    name "TEST_PC"
    tagging
    exit
!
...
interface port EXT21
    tagging
    pvid 10
    exit
!
interface port EXT22
    name "TO_VIE_ETH1/1"
    tagging
    pvid 10
    exit
!
interface port EXT23
    tagging
    pvid 10
    exit
!
interface port EXT24
    name "TO_STR_ETH1/1"
    tagging
    pvid 10
    exit
!
vlan 1
    member INTA1-EXT20
    no member EXT21-EXT24
!
!
vlan 10
    enable
    name "Server"
    member EXT4,EXT21-EXT24
!
!
vlan 20
    enable
    name "Data20"
    member EXT4,EXT21-EXT24
```

```
!  
!  
vlan 30  
    enable  
    name "Data30"  
    member EXT4,EXT21-EXT24  
!  
!  
vlan 40  
    enable  
    name "Data40"  
    member EXT4,EXT21-EXT24  
!  
!  
!  
spanning-tree stp 10 vlan 10  
  
spanning-tree stp 20 vlan 20  
  
spanning-tree stp 30 vlan 30  
  
spanning-tree stp 40 vlan 40  
  
!  
!  
!  
!  
!  
!  
lldp enable  
!  
!  
!  
!  
...  
end
```

Example 4-7 Output of show running-config command: STR switch

```
str# show running-config
version 5.1(3)N2(1)
hostname str

feature telnet
no feature http-server
feature lldp

username admin password 5 $1$0c8ULbm7$bRaCJLmRCrkJRUIDcNaaJ0 role network-admin

...

vrf context management
  ip route 0.0.0.0/0 192.168.240.1
vlan 1
vlan 10
  name Server
vlan 20
  name Data20
vlan 30
  name Data30
vlan 40
  name Data40
spanning-tree vlan 10,30 priority 24576
spanning-tree vlan 20,40 priority 28672

interface Ethernet1/1
  description TO_FLEX_EXT24
  switchport mode trunk
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40

...

interface Ethernet1/19
  description TO_VIE_ETH1/39
  switchport mode trunk
  switchport access vlan 10

interface Ethernet1/20
  shutdown
  switchport mode trunk
  switchport access vlan 10

interface mgmt0
  ip address 192.168.240.30/24
  clock timezone MESZ 2 0
  line console
  line vty
  boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N2.1.bin

boot system bootflash:/n5000-uk9.5.1.3.N2.1.bin
```

Example 4-8 Output of show running-config command: VIE switch

```
vie# show running-config
version 5.1(3)N2(1)
hostname vie

feature telnet
feature lldp

username admin password 5 $1$3QkdUbKB$s1Ytem8Ty6FfYtQc9Zs0k1 role network-admin
...
vrf context management
    ip route 0.0.0.0/0 192.168.240.1
vlan 1
vlan 10
    name Server
vlan 20
    name Data20
vlan 30
    name Data30
vlan 40
    name Data40
spanning-tree vlan 10,30 priority 28672
spanning-tree vlan 20,40 priority 24576

...
interface Ethernet1/1
    description TO_FLEX_EXT22
    switchport mode trunk
    switchport trunk native vlan 10
    switchport trunk allowed vlan 10,20,30,40

interface Ethernet1/39
    switchport mode trunk
    switchport access vlan 10
...

interface mgmt0
    no snmp trap link-status
    vrf member management
    ip address 192.168.240.20/24
clock timezone MESZ 2 0
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N2.1.bin
boot system bootflash:/n5000-uk9.5.1.3.N2.1.bin
```

4.3 Use Case 2: PVRST with LACP Channeling

In this use case, we added a second link between each switch pair to test PVRST with LACP channeling (see Figure 4-4).

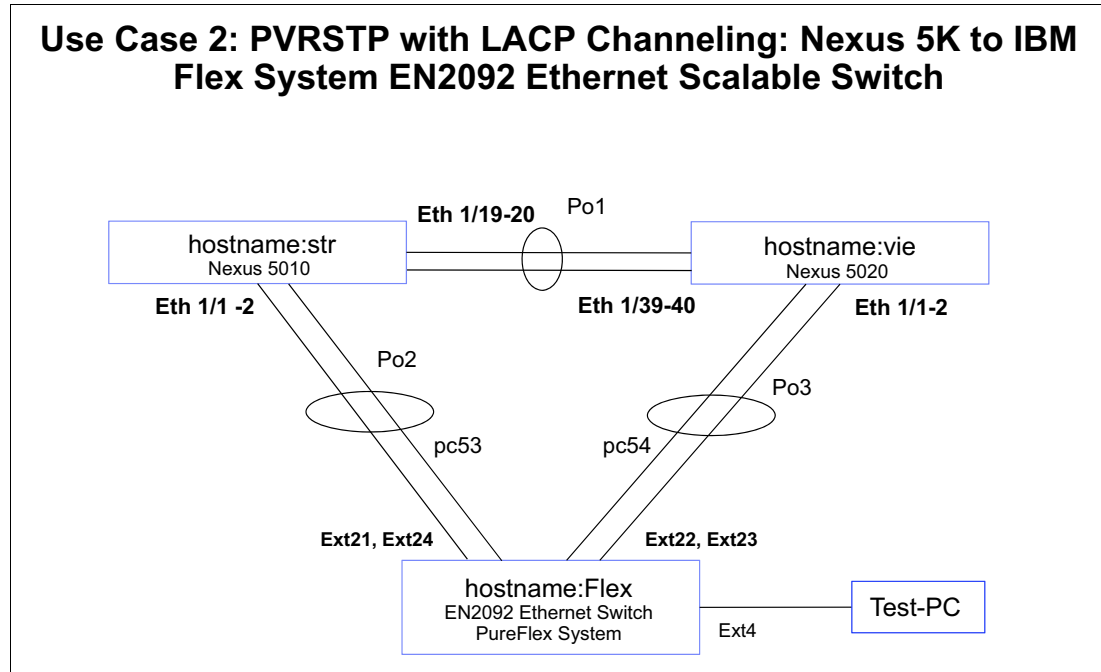


Figure 4-4 Use Case 2

4.3.1 Verifying the topology used by using lldp

As in Use Case 1, we verified the configurations with several show commands on the IBM and on the Cisco switches.

A best practice to check the topology is using **show lldp remote-device** on the IBM Flex System switch and **show lldp neighbors** on the Cisco Nexus switch. Important parameters and detail are highlighted in red.

The commands that were run on the three switches produced the following outputs:

- ▶ Flex System EN2029: Example 4-9 on page 111
- ▶ G8264 STR switch: Example 4-10 on page 111
- ▶ G8264 VIE switch: Example 4-11 on page 111

Example 4-9 Output of show lldp remote-device on the Flex System switch

Flex#**show lldp remote-device**

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
EXT22	1	00 0d ec a3 8f 88	Eth1/1	vie
EXT24	2	00 05 9b 7b 84 08	Eth1/1	str
INTA1	3	5c f3 fc 5f 43 9d	5c-f3-fc-5f-43-9d	
EXT21	4	00 05 9b 7b 84 09	Eth1/2	str
EXT23	5	00 0d ec a3 8f 89	Eth1/2	vie

Example 4-10 Output of show lldp neighbor on the STR switch

str# **show lldp neighbour**

Capability codes:

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

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

Device ID	Local Intf	Hold-time	Capability	Port ID
Flex	Eth1/1	120	BR	52
Flex	Eth1/2	120	BR	49
vie	Eth1/19	120	B	Eth1/39
vie	Eth1/20	120	B	Eth1/40
Total entries displayed: 4				

Example 4-11 Output of show lldp neighbors on the VIE switch

vie# **show lldp neighbors**

Capability codes:

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

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

Device ID	Local Intf	Hold-time	Capability	Port ID
Flex	Eth1/1	120	BR	50
Flex	Eth1/2	120	BR	51
str	Eth1/39	120	B	Eth1/19
str	Eth1/40	120	B	Eth1/20

4.3.2 Verifying trunks

To review which vlans are active on which trunk, we used the show interface trunk on IBM Flex switch and on the Cisco Nexus switch.

The commands that were run on the three switches produced the following outputs:

- Flex System EN2029: Example 4-12
- G8264 STR switch: Example 4-13 on page 113
- G8264 VIE switch: Example 4-14 on page 114

Important parameters and details are highlighted in red.

Example 4-12 Output of show interface trunk on the Flex System switch

Flex#**show interface trunk**

Alias	Port	Tag	RMON	Lrn	Fld	PVID	NAME	VLAN(s)

...								
EXT4	32	y	d	e	e	1	TEST_PC	1 10 20 30 40
...								
EXT21	49	y	d	e	e	10	TO_STR_ETH1/2	10 20 30 40
EXT22	50	y	d	e	e	10	TO_VIE_ETH1/1	10 20 30 40
EXT23	51	y	d	e	e	10	TO_VIE_ETH1/2	10 20 30 40
EXT24	52	y	d	e	e	10	TO_STR_ETH1/1	10 20 30 40
MGT1	53	y	d	e	e	4095	MGT1	4095

Example 4-13 Output of show interface trunk on the STR switch

str# **show interface trunk**

Port	Native Vlan	Status	Port Channel
Eth1/1	10	trnk-bndl	Po2
Eth1/2	10	trnk-bndl	Po2
Eth1/19	1	trnk-bndl	Po1
Eth1/20	1	trnk-bndl	Po1
Po1	1	trunking	--
Po2	10	trunking	--

Port	Vlans Allowed on Trunk
Eth1/1	10,20,30,40
Eth1/2	10,20,30,40
Eth1/19	1-3967,4048-4093
Eth1/20	1-3967,4048-4093
Po1	1-3967,4048-4093
Po2	10,20,30,40

Port	Vlans Err-disabled on Trunk
Eth1/1	none
Eth1/2	none
Eth1/19	none
Eth1/20	none
Po1	none
Po2	none

Port	STP Forwarding
Eth1/1	none
Eth1/2	none
Eth1/19	none
Eth1/20	none
Po1	1,10,20,30,40
Po2	10,20,30,40

Port	Vlans in spanning tree forwarding state and not pruned
Eth1/1	--
Eth1/2	--
Eth1/19	--
Eth1/20	--
Po1	--
Po2	--

Example 4-14 Output of show interface trunk on the VIE switch

```
vie# show interface trunk
```

Port	Native Vlan	Status	Port Channel
Eth1/1	10	trnk-bndl	Po3
Eth1/2	10	trnk-bndl	Po3
Eth1/39	1	trnk-bndl	Po1
Eth1/40	1	trnk-bndl	Po1
Po1	1	trunking	--
Po3	10	trunking	--

Port	Vlans Allowed on Trunk
Eth1/1	10,20,30,40
Eth1/2	10,20,30,40
Eth1/39	1-3967,4048-4093
Eth1/40	1-3967,4048-4093
Po1	1-3967,4048-4093
Po3	10,20,30,40

Port	Vlans Err-disabled on Trunk
Eth1/1	none
Eth1/2	none
Eth1/39	none
Eth1/40	none
Po1	none
Po3	none

Port	STP Forwarding
Eth1/1	none
Eth1/2	none
Eth1/39	none
Eth1/40	none
Po1	1,10,20,30,40
Po3	10,20,30,40

Port	Vlans in spanning tree forwarding state and not pruned
Eth1/1	--
Eth1/2	--
Eth1/39	--
Eth1/40	--
Po1	--
Po3	--

4.3.3 Verifying PVRST spanning tree configuration

In the next step, we verified the PVRST spanning tree configuration of the switches by executing the **show spanning-tree** command. In Figure 4-5 and Figure 4-6 on page 116, showing even and odd VLANs, you can verify the status on the respective Ethernet interface-referring VLAN, port state, and port role.

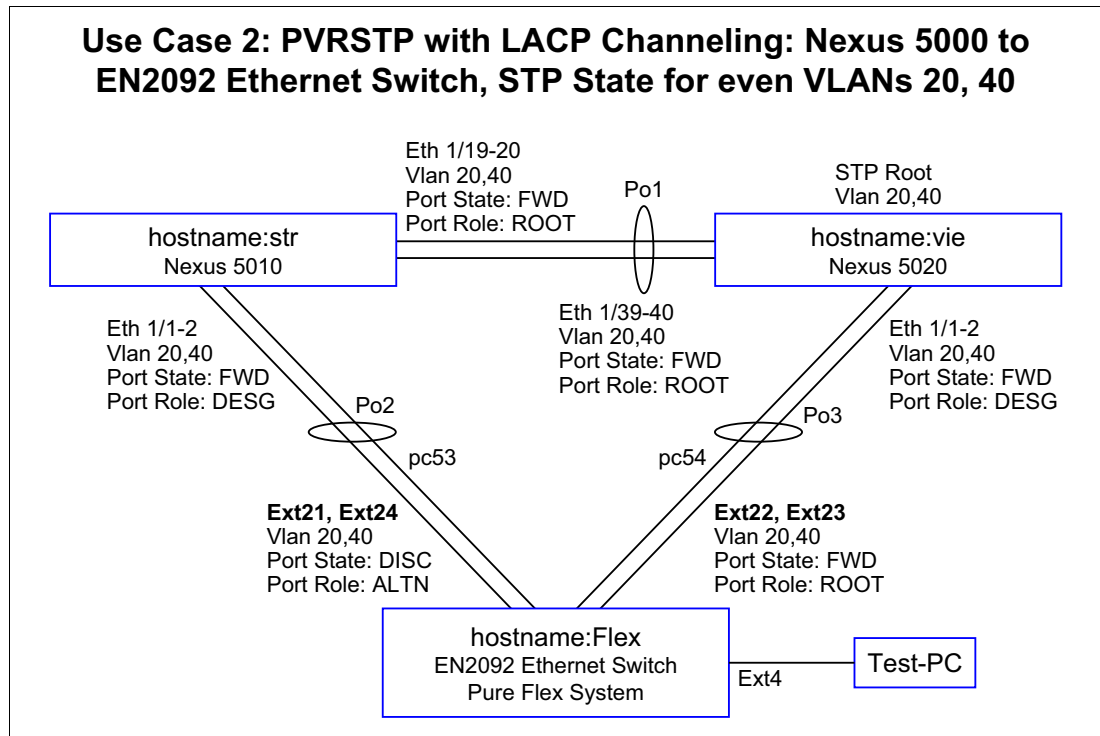


Figure 4-5 Use Case 2: Even-numbered VLANs

Use Case 3: PVRSTP with LACP Channeling: Nexus 5000 to En2092 Ethernet Switch, STP State for odd VLANs 10, 30

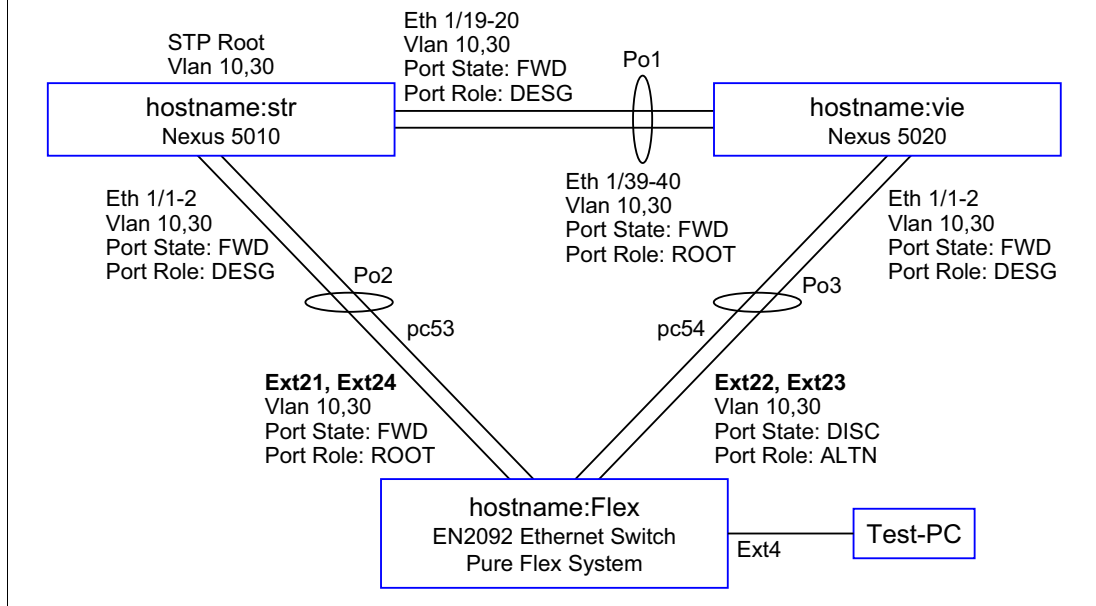


Figure 4-6 Use Case 2: Odd-numbered VLANs

In Example 4-15, the outputs of the show commands of the Flex System and Nexus switches show all of the link pairs are successfully channelled with LACP. The important parameters and details are highlighted in red.

Example 4-15 Configuration output

Flex#show spanning-tree

Pvst+ compatibility mode enabled

Spanning Tree Group 1: On (PVRST)
VLANs: 1

Current Root: Path-Cost Port Hello MaxAge FwdDel
8000 00:16:ca:a1:c1:00 20000 EXT3 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
61441 2 20 15 300 13

Port	Prio	Cost	State	Role	Designated	Bridge	Des Port	Type
INTA1	0	0	FWD *					
INTA2	0	0	FWD *					
INTA4	0	0	FWD *					
EXT1	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00		801d	P2P
EXT2	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00		801e	P2P
EXT3	128	20000!	FWD	ROOT	8000-00:16:ca:a1:c1:00		8011	P2P
EXT4	128	20000!	FWD	DESG	f001-08:17:f4:76:78:00		8020	P2P

* = STP turned off for this port.

! = Automatic path cost.

Spanning Tree Group 10: On (PVRST)
VLANs: 10

Current Root: Path-Cost Port Hello MaxAge FwdDel
600a 00:05:9b:7b:84:3c 990 EXT21 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
61450 2 20 15 300 28

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	FWD	DESG	f00a-08:17:f4:76:78:00	8020	P2P
EXT21 (pc53)	128	990!+	FWD	ROOT	600a-00:05:9b:7b:84:3c	9001	P2P
EXT22 (pc54)	128	990!+	DISC	ALTN	700a-00:0d:ec:a3:8f:bc	9002	P2P
EXT23 (pc54)	128	990!+	DISC	ALTN	700a-00:0d:ec:a3:8f:bc	9002	P2P
EXT24 (pc53)	128	990!+	FWD	ROOT	600a-00:05:9b:7b:84:3c	9001	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 20: On (PVRST)
VLANs: 20

Current Root: Path-Cost Port Hello MaxAge FwdDel
6014 00:0d:ec:a3:8f:bc 990 EXT22 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
61460 2 20 15 300 20

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	FWD	DESG	f014-08:17:f4:76:78:00	8020	P2P
EXT21 (pc53)	128	990!+	DISC	ALTN	7014-00:05:9b:7b:84:3c	9001	P2P
EXT22 (pc54)	128	990!+	FWD	ROOT	6014-00:0d:ec:a3:8f:bc	9002	P2P
EXT23 (pc54)	128	990!+	FWD	ROOT	6014-00:0d:ec:a3:8f:bc	9002	P2P
EXT24 (pc53)	128	990!+	DISC	ALTN	7014-00:05:9b:7b:84:3c	9001	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

Spanning Tree Group 30: On (PVRST)
VLANs: 30

Current Root: Path-Cost Port Hello MaxAge FwdDel
601e 00:05:9b:7b:84:3c 990 EXT21 2 20 15

Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
61470 2 20 15 300 18

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	FWD	DESG	f01e-08:17:f4:76:78:00	8020	P2P
EXT21 (pc53)	128	990!+	FWD	ROOT	601e-00:05:9b:7b:84:3c	9001	P2P
EXT22 (pc54)	128	990!+	DISC	ALTN	701e-00:0d:ec:a3:8f:bc	9002	P2P
EXT23 (pc54)	128	990!+	DISC	ALTN	701e-00:0d:ec:a3:8f:bc	9002	P2P
EXT24 (pc53)	128	990!+	FWD	ROOT	601e-00:05:9b:7b:84:3c	9001	P2P

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

```
-----
Spanning Tree Group 40: On (PVRST)
VLANs: 40
```

```
Current Root:          Path-Cost  Port Hello MaxAge FwdDel
6028 00:0d:ec:a3:8f:bc    990  EXT22   2    20    15
```

```
Parameters: Priority Hello MaxAge FwdDel Aging Topology Change Counts
             61480     2     20    15    300          20
```

```

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
-----
EXT4           128    20000! FWD    DESG f028-08:17:f4:76:78:00    8020        P2P
EXT21 (pc53)   128     990!+ DISC  ALTN 7028-00:05:9b:7b:84:3c    9001        P2P
EXT22 (pc54)   128     990!+ FWD   ROOT 6028-00:0d:ec:a3:8f:bc    9002        P2P
EXT23 (pc54)   128     990!+ FWD   ROOT 6028-00:0d:ec:a3:8f:bc    9002        P2P
EXT24 (pc53)   128     990!+ DISC  ALTN 7028-00:05:9b:7b:84:3c    9001        P2P
```

```
! = Automatic path cost.
+ = Portchannel cost, not the individual port cost.
```

```
-----
Spanning Tree Group 128: Off (PVRST), FDB aging timer 300
VLANs: 4095
```

```

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
-----
MGT1           0         0  FWD *
*= STP turned off for this port.
```

As shown in Example 4-15 on page 116, Ethernet interfaces EXT21-24 are bundled to channels, in which EXT21 and EXT24 form portchannel 53 and EXT22 and EXT23 form portchannel 54.

4.3.4 Bridge priority field in the show spanning tree output

When STP was first used, there was only one spanning tree per physical switch in which the bridge priority was stored as a 16-bit value (0-65535). With the introduction of per VLAN spanning tree, the need to carry the VLAN ID within the bridge priority field became apparent.

The top 4 bits were still used for the bridge priority value, but the remaining 12 bits were used to carry the VLAN ID (1-1046).

Table 4-1 lists the 16 bits translated to decimal.

Table 4-1 Bridge priority field

Usage	Bridge priority: 4 bits				VLAN ID: 12 bit											
Bit value	32768	16384	8192	4096	2048	1024	512	256	128	64	32	16	8	4	2	1

If you configure the bridge priority value at Cisco IOS, you must enter a multiple of 4096 or use the keywords root primary or root secondary. If you configure the bridge priority at IBM OS, you can enter any value and the switch changes it to the next lower value that is divisible by 4096. The output of the show spanning tree command is shown in Example 4-16 on page 119 and Example 4-17 on page 120. The important parameters and details are highlighted in red.

Example 4-16 Output of show spanning-tree command

str# show spanning-tree

VLAN0001

Spanning tree enabled protocol rstp
Root ID Priority 32769
 Address 0005.9b7b.843c
 This bridge is the root
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

VLAN0010

Spanning tree enabled protocol rstp
Root ID Priority 24586
 Address 0005.9b7b.843c
 This bridge is the root

For VLAN 10 and other odd vlans, this bridge is the root

Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 24586 (priority 24576 sys-id-ext 10)
Address 0005.9b7b.843c
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po1	Desg	FWD	1	128.4096	P2p
Po2	Desg	FWD	1	128.4097	P2p
Eth1/16	Desg	FWD	2	128.144	P2p

VLAN0020

Spanning tree enabled protocol rstp
Root ID Priority 24596
 Address 000d.eca3.8fbc
 Cost 1
 Port 4096 (port-channel1)

For VLAN 20 and other even vlans, Po1 leads to the rootbridge (Nexus 5000 Vie)

Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 28692 (priority 28672 sys-id-ext 20)
Address 0005.9b7b.843c
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po1	Root	FWD	1	128.4096	P2p
Po2	Desg	FWD	1	128.4097	P2p

VLAN0030

Spanning tree enabled protocol rstp
Root ID Priority 24606
 Address 0005.9b7b.843c
 This bridge is the root
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

Bridge ID Priority 24606 (priority 24576 sys-id-ext 30)
Address 0005.9b7b.843c
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po1	Desg	FWD	1	128.4096	P2p
Po2	Desg	FWD	1	128.4097	P2p

VLAN0040

```

Spanning tree enabled protocol rstp
Root ID Priority 24616
Address 000d.eca3.8fbc
Cost 1
Port 4096 (port-channel1)
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

```

Bridge ID Priority 28712 (priority 28672 sys-id-ext 40)
Address 0005.9b7b.843c
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po1	Root	FWD	1	128.4096	P2p
Po2	Desg	FWD	1	128.4097	P2p

Example 4-17 Output from show spanning-tree on VIE switch

```
vie# show spanning-tree
```

VLAN0001

```

Spanning tree enabled protocol rstp
Root ID Priority 32769
Address 0005.9b7b.843c
Cost 1
Port 4096 (port-channel1)
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

```

Bridge ID Priority 32769 (priority 32768 sys-id-ext 1)
Address 000d.eca3.8fbc
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po1	Root	FWD	1	128.4096	P2p

VLAN0010

```

Spanning tree enabled protocol rstp
Root ID Priority 24586
Address 0005.9b7b.843c
Cost 1
Port 4096 (port-channel1)
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

```

Bridge ID Priority 28682 (priority 28672 sys-id-ext 10)
Address 000d.eca3.8fbc

```


Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po1	Root	FWD	1	128.4096	P2p
Po3	Desg	FWD	1	128.4098	P2p
Eth1/16	Desg	FWD	2	128.144	P2p

VLAN0020

Spanning tree enabled protocol rstp

Root ID Priority 24596
 Address 000d.eca3.8fbc
 This bridge is the root
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 24596 (priority 24576 sys-id-ext 20)
 Address 000d.eca3.8fbc
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po1	Desg	FWD	1	128.4096	P2p
Po3	Desg	FWD	1	128.4098	P2p

VLAN0030

Spanning tree enabled protocol rstp

Root ID Priority 24606
 Address 0005.9b7b.843c
 Cost 1
 Port 4096 (port-channel1)
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 28702 (priority 28672 sys-id-ext 30)
 Address 000d.eca3.8fbc
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po1	Root	FWD	1	128.4096	P2p
Po3	Desg	FWD	1	128.4098	P2p

VLAN0040

Spanning tree enabled protocol rstp

Root ID Priority 24616
 Address 000d.eca3.8fbc
 This bridge is the root
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 24616 (priority 24576 sys-id-ext 40)
 Address 000d.eca3.8fbc
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po1	Desg	FWD	1	128.4096	P2p
Po3	Desg	FWD	1	128.4098	P2p

4.3.5 Show running-config of all switches in Use Case 2

In the following configuration print outs of the IBM Flex Switch and the Cisco Nexus switches, you can see the necessary configuration steps that we performed during our test. The important parameters and details are highlighted in red.

The commands that were run on the three switches produced the following outputs:

- Flex System EN2029: Example 4-18
- G8264 STR switch: Example 4-19 on page 124
- G8264 VIE switch: Example 4-20 on page 126

Example 4-18 Output of show running-config command: Flex System switch

```
Flex#sh run
Current configuration:
!
version "7.2.2.2"
switch-type "IBM Flex System EN2092 1Gb Ethernet Scalable Switch"
!
!

...
hostname "Flex"
system idle 60
!
!
access http enable
access telnet enable
!
...
interface port EXT4
    name "TEST_PC"
    tagging
    exit
!
...
interface port EXT21
    name "TO_STR_ETH1/2"
    tagging
    pvid 10
    exit
!
interface port EXT22
    name "TO_VIE_ETH1/1"
    tagging
    pvid 10
    exit
!
interface port EXT23
    name "TO_VIE_ETH1/2"
    tagging
    pvid 10
    exit
!
interface port EXT24
    name "TO_STR_ETH1/1"
    tagging
    pvid 10
    exit
```

```

!
vlan 1
    member INTA1-EXT20
    no member EXT21-EXT24
!
!
vlan 10
    enable
    name "Server"
    member EXT4,EXT21-EXT24
!
!
vlan 20
    enable
    name "Data20"
    member EXT4,EXT21-EXT24
!
!
vlan 30
    enable
    name "Data30"
    member EXT4,EXT21-EXT24
!
!
vlan 40
    enable
    name "Data40"
    member EXT4,EXT21-EXT24
!
!
!
spanning-tree stp 10 vlan 10

spanning-tree stp 20 vlan 20

spanning-tree stp 30 vlan 30

spanning-tree stp 40 vlan 40

!
! This configures the LACP portchannels in the IBM PureFlex switch
!
interface port EXT21
    lacp mode active
    lacp key 2
!
interface port EXT22
    lacp mode active
    lacp key 3
!
interface port EXT23
    lacp mode active
    lacp key 3
!
interface port EXT24
    lacp mode active
    lacp key 2
!
!
!

```

```
!  
!  
!  
lldp enable  
!  
!  
!  
!  
...  
end
```

Example 4-19 Output from the show running-config command: STR switch

```
str# show run  
version 5.1(3)N2(1)  
hostname str  
  
feature telnet  
no feature http-server  
  
! Enables LACP  
feature lACP  
feature lldp  
  
username admin password 5 $1$0c8ULbm7$bRaCJLmRCrkJRUIDcNaaJO role network-admin  
no password strength-check  
  
...  
  
vrf context management  
  ip route 0.0.0.0/0 192.168.240.1  
vlan 1  
vlan 10  
  name Server  
vlan 20  
  name Data20  
vlan 30  
  name Data30  
vlan 40  
  name Data40  
spanning-tree vlan 10,30 priority 24576  
spanning-tree vlan 20,40 priority 28672  
  
interface port-channel1  
  description TO_VIE_P01  
  switchport mode trunk  
  switchport access vlan 10  
!  
!  
! Configure Portchannel  
!  
interface port-channel2  
  description TO_FLEX_EXT21,EXT24  
  switchport mode trunk  
  switchport trunk native vlan 10  
  switchport trunk allowed vlan 10,20,30,40
```

```

!
!Configure interface and add it to portchannel2 by use of LACP (keyword = active)
!
interface Ethernet1/1
  description TO_FLEX_EXT24
  switchport mode trunk
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  channel-group 2 mode active

interface Ethernet1/2
  description TO_FLEX_EXT21
  switchport mode trunk
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  channel-group 2 mode active

...

interface Ethernet1/19
  description TO_VIE_ETH1/39
  switchport mode trunk
  switchport access vlan 10
  channel-group 1 mode active

interface Ethernet1/20
  description TO_VIE_ETH1/40
  switchport mode trunk
  switchport access vlan 10
  channel-group 1 mode active

interface mgmt0
  ip address 192.168.240.30/24
  clock timezone MESZ 2 0
  line console
  line vty
  boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N2.1.bin
  boot system bootflash:/n5000-uk9.5.1.3.N2.1.bin

```

Example 4-20 Output of show running-config command: VIE switch

```
vie# show run

version 5.1(3)N2(1)
hostname vie

feature telnet
feature lacp
feature

username admin password 5 $1$3QkdUbKB$s1Ytem8Ty6FfYtQc9Zs0k1 role network-admin
no password strength-check

...

vrf context management
  ip route 0.0.0.0/0 192.168.240.1
vlan 1
vlan 10
  name Server
vlan 20
  name Data20
vlan 30
  name Data30
vlan 40
  name Data40
spanning-tree vlan 10,30 priority 28672
spanning-tree vlan 20,40 priority 24576

interface port-channel1
  description T0_STR_P01
  switchport mode trunk
  switchport access vlan 10

!Configure interface and add it to portchannel3 by use of LACP (keyword = active)

interface port-channel3
  description T0_FLEX_EXT22,EXT23
  switchport mode trunk
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40

interface Ethernet1/1
  description T0_FLEX_EXT22
  switchport mode trunk
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  channel-group 3 mode active

interface Ethernet1/2
  description T0_FLEX_EXT23
  switchport mode trunk
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  channel-group 3 mode active

...

interface Ethernet1/39
```

```

description TO_STR_ETH1/19
switchport mode trunk
switchport access vlan 10
channel-group 1 mode active

interface Ethernet1/40
description TO_STR_ETH1/19
switchport mode trunk
switchport access vlan 10
channel-group 1 mode active

...

interface mgmt0
no snmp trap link-status
vrf member management
ip address 192.168.240.20/24
clock timezone MESZ 2 0
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N2.1.bin
boot system bootflash:/n5000-uk9.5.1.3.N2.1.bin

```

4.4 Use Case 3: MST with LACP Channeling

In this use case, we configured MST instead of PVRST as the spanning tree option with LACP channeling, as shown in Figure 4-7.

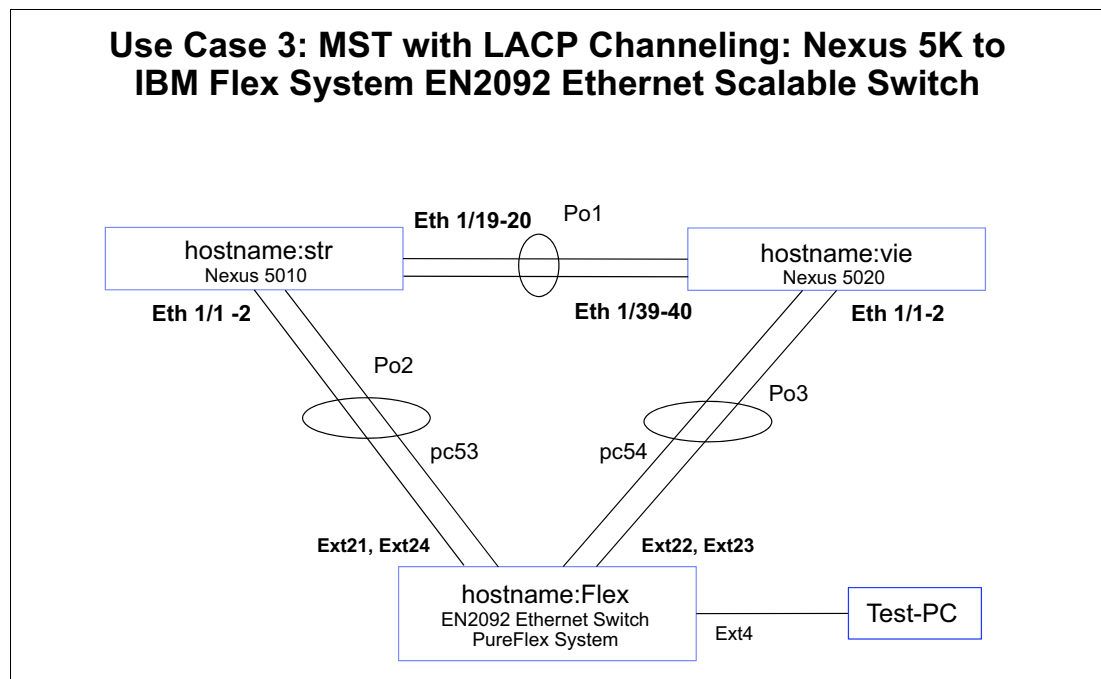


Figure 4-7 Use Case 3: MST with LACP Channeling

4.4.1 Verifying the topology used by using lldp

As in the other use cases, we verified the configurations by using several show commands on the IBM and on the Cisco switches.

A best practice to verify the topology is the use of the **show lldp remote-device** command on the IBM Flex switch and the **show lldp neighbors** command on the Cisco Nexus switch. First, we verified the topology after the configuration changes were made, as shown in Example 4-21.

Example 4-21 Verifying the configurations

```
Flex#sh lldp remote-device
LLDP Remote Devices Information
```

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
EXT22	1	00 0d ec a3 8f 88	Eth1/1	vie
EXT24	2	00 05 9b 7b 84 08	Eth1/1	str
INTA1	3	5c f3 fc 5f 43 9d	5c-f3-fc-5f-43-9d	
EXT21	4	00 05 9b 7b 84 09	Eth1/2	str
EXT23	5	00 0d ec a3 8f 89	Eth1/2	vie

```
str# show lldp neighbors
```

Capability codes:

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

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

Device ID	Local Intf	Hold-time	Capability	Port ID
Flex	Eth1/1	120	BR	52
Flex	Eth1/2	120	BR	49
vie	Eth1/19	120	B	Eth1/39
vie	Eth1/20	120	B	Eth1/40

Total entries displayed: 4

```
vie# show lldp neighbors
```

Capability codes:

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

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

Device ID	Local Intf	Hold-time	Capability	Port ID
Flex	Eth1/1	120	BR	50
Flex	Eth1/2	120	BR	51
str	Eth1/39	120	B	Eth1/19
str	Eth1/40	120	B	Eth1/20

Total entries displayed: 4

4.4.2 Verifying trunks

To verify which VLANs are active on which trunk, we used the **show interface trunk** command on the IBM Flex System switch and the Cisco Nexus switch, as shown in Example 4-22. Important parameters and details are highlighted in red.

Example 4-22 Reviewing active VLANs and trunks

```
Flex#show interface trunk
```

Alias	Port	Tag	RMON	Ln	Fld	PVID	NAME	VLAN(s)
...								
EXT4	32	y	d	e	e	1	TEST_PC	1 10 20 30 40
...								


```

EXT21  49  y  d  e  e  10  TO_STR_ETH1/2  10 20 30 40
EXT22  50  y  d  e  e  10  TO_VIE_ETH1/1  10 20 30 40
EXT23  51  y  d  e  e  10  TO_VIE_ETH1/2  10 20 30 40
EXT24  52  y  d  e  e  10  TO_STR_ETH1/1  10 20 30 40
MGT1   53  y  d  e  e  4095 MGT1          4095

```

str# **show interface trunk**

```

-----
Port          Native  Status      Port
              Vlan                Channel
-----
Eth1/1        10      trnk-bndl   Po2
Eth1/2        10      trnk-bndl   Po2
Eth1/19       1       trnk-bndl   Po1
Eth1/20       1       trnk-bndl   Po1
Po1           1       trunking    --
Po2           10      trunking    --

```

```

-----
Port          Vlans Allowed on Trunk
-----
Eth1/1        10,20,30,40
Eth1/2        10,20,30,40
Eth1/19       1-3967,4048-4093
Eth1/20       1-3967,4048-4093
Po1           1-3967,4048-4093
Po2           10,20,30,40

```

```

-----
Port          Vlans Err-disabled on Trunk
-----
Eth1/1        none
Eth1/2        none
Eth1/19       none
Eth1/20       none
Po1           none
Po2           none

```

```

-----
Port          STP Forwarding
-----
Eth1/1        none
Eth1/2        none
Eth1/19       none
Eth1/20       none
Po1           1,10,20,30,40
Po2           10,30

```

```

-----
Port          Vlans in spanning tree forwarding state and not pruned
-----
Eth1/1        --
Eth1/2        --
Eth1/19       --
Eth1/20       --
Po1           --
Po2           --

```

```
-----
Port          Vlans Forwarding on FabricPath
-----
```

```
vie# show interface trunk
```

```
-----
Port          Native  Status      Port
              Vlan                Channel
-----
Eth1/1        10      trnk-bndl   Po3
Eth1/2        10      trnk-bndl   Po3
Eth1/39       1       trnk-bndl   Po1
Eth1/40       1       trnk-bndl   Po1
Po1           1       trunking    --
Po3           10      trunking    --
-----
```

```
-----
Port          Vlans Allowed on Trunk
-----
Eth1/1        10,20,30,40
Eth1/2        10,20,30,40
Eth1/39       1-3967,4048-4093
Eth1/40       1-3967,4048-4093
Po1           1-3967,4048-4093
Po3           10,20,30,40
-----
```

```
-----
Port          Vlans Err-disabled on Trunk
-----
Eth1/1        none
Eth1/2        none
Eth1/39       none
Eth1/40       none
Po1           none
Po3           none
-----
```

```
-----
Port          STP Forwarding
-----
Eth1/1        none
Eth1/2        none
Eth1/39       none
Eth1/40       none
Po1           1,10,20,30,40
Po3           20,40
-----
```

```
-----
Port          Vlans in spanning tree forwarding state and not pruned
-----
Eth1/1        --
Eth1/2        --
Eth1/39       --
Eth1/40       --
Po1           --
Po3           --
-----
```

Figure 4-8 shows the odd-numbered VLANs. Figure 4-9 shows the even-numbered VLANs.

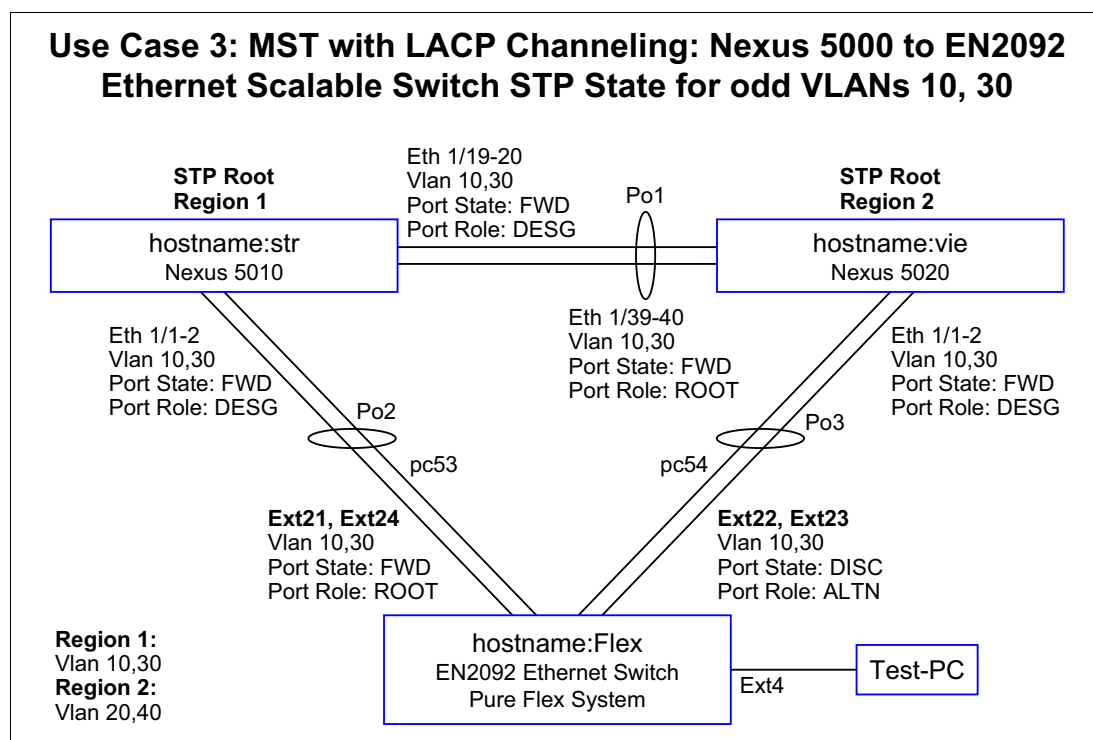


Figure 4-8 Use Case 3: VLANs 10, 30

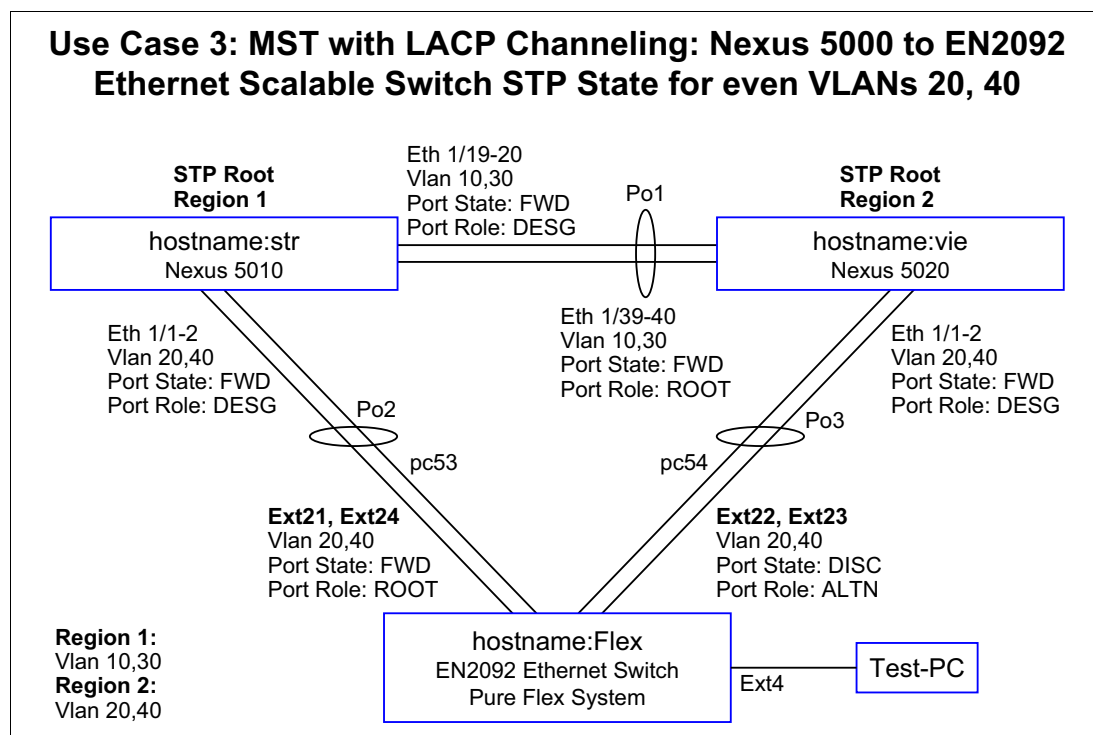


Figure 4-9 Use Case 3: VLANs 20, 40

4.4.3 Verifying MST spanning tree configuration

In the following configuration print outs of the IBM Flex System switch and the Cisco Nexus switches, you can see the necessary configuration steps that we performed during our test. We also add some remarks to help explain the configuration that was used.

The commands that were run on the three switches produced the following outputs:

- ▶ Flex System EN2029: Example 4-23
- ▶ G8264 STR switch: Example 4-24 on page 133
- ▶ G8264 VIE switch: Example 4-25 on page 134

The important parameters and details are highlighted in red. As you can see highlighted in red, MST is enabled on all MST instances on both Nexus switches.

Example 4-23 Output of show spanning-tree command: Flex System switch

```
Flex#sh spanning-tree
-----
Pvst+ compatibility mode enabled

Mstp Digest: 0xe821ccee7501115289b37c79a72e07c9

-----
Spanning Tree Group 1: On (MSTP)
VLANs MAPPED: 10 30
VLANs: 10 30

! Now spanning-tree protocol is MST for odd vlans 10 and 30

Current Root:          Path-Cost  Port
6000 00:05:9b:7b:84:3c    990   EXT21

Parameters:  Priority  Aging  Topology Change Counts
              61440    300           2

      Port      Prio    Cost    State  Role Designated Bridge      Des Port    Type
-----
EXT4           128    20000!  FWD    DESG f000-08:17:f4:76:78:00    8020
EXT21 (pc53)   128    990!+  FWD    ROOT 6000-00:05:9b:7b:84:3c    9001      P2P
EXT22 (pc54)   128    990!+  FWD    DESG f000-08:17:f4:76:78:00    806b      P2P
EXT23 (pc54)   128    990!+  FWD    DESG f000-08:17:f4:76:78:00    806b      P2P
EXT24 (pc53)   128    990!+  FWD    ROOT 6000-00:05:9b:7b:84:3c    9001      P2P

! EXT 21 - 24 are portchannels. EXT21 and EXT24 formed portchannel 53, EXT 21 and EXT23
! formed portchannel 54.

! = Automatic path cost.
+ = Portchannel cost, not the individual port cost.

-----
Spanning Tree Group 2: On (MSTP)
VLANs MAPPED: 20 40
VLANs: 20 40

! Now spanning-tree protocol is MST for even vlans 20 and 40

Current Root:          Path-Cost  Port
6000 00:0d:ec:a3:8f:bc    990   EXT22

Parameters:  Priority  Aging  Topology Change Counts
```

61440 300 1

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	FWD	DESG	f000-08:17:f4:76:78:00	8020	P2P
EXT21 (pc53)	128	990!+	FWD	DESG	f000-08:17:f4:76:78:00	806a	P2P
EXT22 (pc54)	128	990!+	FWD	ROOT	6000-00:0d:ec:a3:8f:bc	9002	P2P
EXT23 (pc54)	128	990!+	FWD	ROOT	6000-00:0d:ec:a3:8f:bc	9002	P2P
EXT24 (pc53)	128	990!+	FWD	DESG	f000-08:17:f4:76:78:00	806a	P2P
! EXT 21 – 24 are portchannels. EXT21 and EXT24 formed portchannel 53, EXT 21 and EXT23 formed portchannel 54.							

! = Automatic path cost.
 + = Portchannel cost, not the individual port cost.

Example 4-24 Output of show spanning-tree command: STR switch

str# **show spanning-tree**

MST0000

Spanning tree enabled protocol mstp

Root ID Priority 32768
 Address 0005.9b7b.843c
 This bridge is the root
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 32768 (priority 32768 sys-id-ext 0)
 Address 0005.9b7b.843c
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po1	Desg	FWD	1000	128.4096	P2p
Po2	Desg	FWD	1000	128.4097	P2p
Eth1/16	Desg	FWD	2000	128.144	P2p

MST0001

Spanning tree enabled protocol mstp

Root ID Priority 24577
 Address 0005.9b7b.843c
 This bridge is the root
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 24577 (priority 24576 sys-id-ext 1)
 Address 0005.9b7b.843c
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po1	Desg	FWD	1000	128.4096	P2p
Po2	Desg	FWD	1000	128.4097	P2p
Eth1/16	Desg	FWD	2000	128.144	P2p

MST0002

Spanning tree enabled protocol mstp

```

Root ID    Priority    24578
          Address    000d.eca3.8fbc
          Cost       1000
          Port       4096 (port-channel1)
          Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID  Priority    32770 (priority 32768 sys-id-ext 2)
          Address    0005.9b7b.843c
          Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface    Role Sts Cost      Prio.Nbr Type
-----
Po1           Root FWD 1000      128.4096 P2p
Po2           Altn BLK 1000      128.4097 P2p

```

Example 4-25 Output of show spanning-tree command: VIE switch

```
vie# show spanning-tree
```

```

MST0000
Spanning tree enabled protocol mstp
Root ID    Priority    32768
          Address    0005.9b7b.843c
          Cost       0
          Port       4096 (port-channel1)
          Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID  Priority    32768 (priority 32768 sys-id-ext 0)
          Address    000d.eca3.8fbc
          Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

```

Interface    Role Sts Cost      Prio.Nbr Type
-----
Po1           Root FWD 1000      128.4096 P2p
Po3           Altn BLK 1000      128.4098 P2p
Eth1/16       Desg FWD 2000      128.144  P2p

```

```

MST0001
Spanning tree enabled protocol mstp
Root ID    Priority    24577
          Address    0005.9b7b.843c
          Cost       1000
          Port       4096 (port-channel1)
          Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID  Priority    32769 (priority 32768 sys-id-ext 1)
          Address    000d.eca3.8fbc
          Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

```

Interface    Role Sts Cost      Prio.Nbr Type
-----
Po1           Root FWD 1000      128.4096 P2p
Po3           Altn BLK 1000      128.4098 P2p
Eth1/16       Desg FWD 2000      128.144  P2p

```

```

MST0002
  Spanning tree enabled protocol mstp
  Root ID    Priority    24578
            Address     000d.eca3.8fbc
            This bridge is the root
            Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24578 (priority 24576 sys-id-ext 2)
            Address     000d.eca3.8fbc
            Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

Interface      Role Sts Cost      Prio.Nbr Type
-----
Po1            Desg FWD 1000      128.4096 P2p
Po3            Desg FWD 1000      128.4098 P2p

```

4.4.4 Show running-config of all switches in Use Case 3

In the following configuration print outs of the IBM Flex System switch and the Cisco Nexus switches, you can see the necessary configuration steps that we performed during our test.

The commands that were run on the three switches produced the following outputs:

- ▶ Flex System EN2029: Example 4-26
- ▶ G8264 STR switch: Example 4-27 on page 137
- ▶ G8264 VIE switch: Example 4-28 on page 139

The important parameters and details are highlighted in red.

Example 4-26 Output of show running-config command: Flex System switch

```

Flex#sh running-config
Current configuration:
!
version "7.2.2.2"
switch-type "IBM Flex System EN2092 1Gb Ethernet Scalable Switch"
!
...

hostname "Flex"
system idle 60
!
!
access http enable
access telnet enable
!
...
interface port EXT4
    name "TEST_PC"
    tagging
    exit
!
...
interface port EXT21
    name "TO_STR_ETH1/2"
    tagging
    pvid 10
    exit

```

```

!
interface port EXT22
    name "TO_VIE_ETH1/1"
    tagging
    pvid 10
    exit
!
interface port EXT23
    name "TO_VIE_ETH1/2"
    tagging
    pvid 10
    exit
!
interface port EXT24
    name "TO_STR_ETH1/1"
    tagging
    pvid 10
    exit
!
vlan 1
    member INTA1-EXT20
    no member EXT21-EXT24
!
!
vlan 10
    enable
    name "Server"
    member EXT4,EXT21-EXT24
!
!
vlan 20
    enable
    name "Data20"
    member EXT4,EXT21-EXT24
!
!
vlan 30
    enable
    name "Data30"
    member EXT4,EXT21-EXT24
!
!
vlan 40
    enable
    name "Data40"
    member EXT4,EXT21-EXT24
!
! Configuration Part to enable MST on the PureFlex Switch
!
spanning-tree mstp version 10
spanning-tree mstp name "PureFlex"
spanning-tree mode mst
spanning-tree mstp cist-add-vlan 1
spanning-tree mstp cist-add-vlan 4095
!
! For odd vlans 10 and 30 we had to configure stp group 1
!
spanning-tree stp 1 vlan 10
spanning-tree stp 1 vlan 30
!

```


! For even vlans 20 and 40 we had to configure stp group 2

```
!  
spanning-tree stp 2 vlan 20  
spanning-tree stp 2 vlan 40
```

```
!  
interface port EXT21  
    lacp mode active  
    lacp key 2  
!  
interface port EXT22  
    lacp mode active  
    lacp key 3  
!  
interface port EXT23  
    lacp mode active  
    lacp key 3  
!  
interface port EXT24  
    lacp mode active  
    lacp key 2  
!  
!  
!  
!  
!  
!  
lldp enable  
!  
!  
!  
!  
!  
...  
end
```

Example 4-27 Output of show running-config command: STR switch

```
str# show run  
  
version 5.1(3)N2(1)  
hostname str  
  
feature telnet  
no feature http-server  
feature lacp  
feature lldp  
  
username admin password 5 $1$0c8ULbm7$bRaCJLmRCrkJRUIDcNaaJ0 role network-admin  
no password strength-check  
  
vrf context management  
    ip route 0.0.0.0/0 192.168.240.1  
vlan 1  
vlan 10  
    name Server
```

```

vlan 20
  name Data20
vlan 30
  name Data30
vlan 40
  name Data40
!
! On the Cisco Nexus switch configuration is slightly different. One the str Nexus
spanning-tree ! priority for odd vlan 10 and 10 are lower than for the even vlan 20 and 40.
This has to be vice ! versa on the vie Nexus Switch. Furthermore you have to define a name
for the MST domain.
!
spanning-tree mode mst
spanning-tree mst 1 priority 24576
spanning-tree vlan 10,30 priority 24576
spanning-tree vlan 20,40 priority 28672
spanning-tree mst configuration
  name PureFlex
  revision 10
  instance 1 vlan 10,30
  instance 2 vlan 20,40

interface port-channel1
  description T0_VIE_P01
  switchport mode trunk
  switchport access vlan 10

interface port-channel2
  description T0_FLEX_EXT21,EXT24
  switchport mode trunk
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40

interface Ethernet1/1
  description T0_FLEX_EXT24
  switchport mode trunk
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  channel-group 2 mode active

interface Ethernet1/2
  description T0_FLEX_EXT21
  switchport mode trunk
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  channel-group 2 mode active

...

interface Ethernet1/19
  description T0_VIE_ETH1/39
  switchport mode trunk
  switchport access vlan 10
  channel-group 1 mode active

interface Ethernet1/20
  description T0_VIE_ETH1/40
  switchport mode trunk
  switchport access vlan 10

```

```

channel-group 1 mode active

interface mgmt0
    ip address 192.168.240.30/24
    clock timezone MESZ 2 0
    line console
    line vty
    boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N2.1.bin
    boot system bootflash:/n5000-uk9.5.1.3.N2.1.bin

```

Example 4-28 Output of show running-config command: VIE switch

```

vie# show run

version 5.1(3)N2(1)
hostname vie

feature telnet
feature lacp
feature lldp

username admin password 5 $1$3QkdUbKB$s1Ytem8Ty6FfYtQc9Zs0k1 role network-admin
no password strength-check
...

vrf context management
    ip route 0.0.0.0/0 192.168.240.1
vlan 1
vlan 10
    name Server
vlan 20
    name Data20
vlan 30
    name Data30
vlan 40
    name Data40
!
! On the Cisco Nexus switch configuration is slightly different. One the vie Nexus
spanning-tree ! priority for even vlan 20 and 40 are lower than for odd vlan 10 and 30.
This has to be vice ! versa on the vie Nexus Switch. Furthermore you have to define a
name for the MST domain.
!
spanning-tree mode mst
spanning-tree mst 2 priority 24576
spanning-tree vlan 10,30 priority 28672
spanning-tree vlan 20,40 priority 24576
spanning-tree mst configuration
    name PureFlex
    revision 10
    instance 1 vlan 10,30
    instance 2 vlan 20,40

interface port-channel1
    description TO_STR_P01
    switchport mode trunk
    switchport access vlan 10

```

```

interface port-channel3
  description T0_FLEX_EXT22,EXT23
  switchport mode trunk
  switchport access vlan 10
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40

interface Ethernet1/1
  description T0_FLEX_EXT22
  switchport mode trunk
  switchport access vlan 10
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  channel-group 3 mode active

interface Ethernet1/2
  description T0_FLEX_EXT23
  switchport mode trunk
  switchport access vlan 10
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  channel-group 3 mode active

interface Ethernet1/39
  description T0_STR_ETH1/19
  switchport mode trunk
  switchport access vlan 10
  channel-group 1 mode active

interface Ethernet1/40
  description T0_STR_ETH1/19
  switchport mode trunk
  switchport access vlan 10
  channel-group 1 mode active

...

interface mgmt0
  no snmp trap link-status
  vrf member management
  ip address 192.168.240.20/24
clock timezone MESZ 2 0
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N2.1.bin
boot system bootflash:/n5000-uk9.5.1.3.N2.1.bin

```

4.5 Use Case 4: MST with LACP Channeling and vPC

To reach our goal of eliminating the spanning tree, we configured vPC on the Nexus 5000 switches. In this case, MST is still enabled. Multiple physical connections between the switches are still channelled by using LACP, as shown in Figure 4-10.

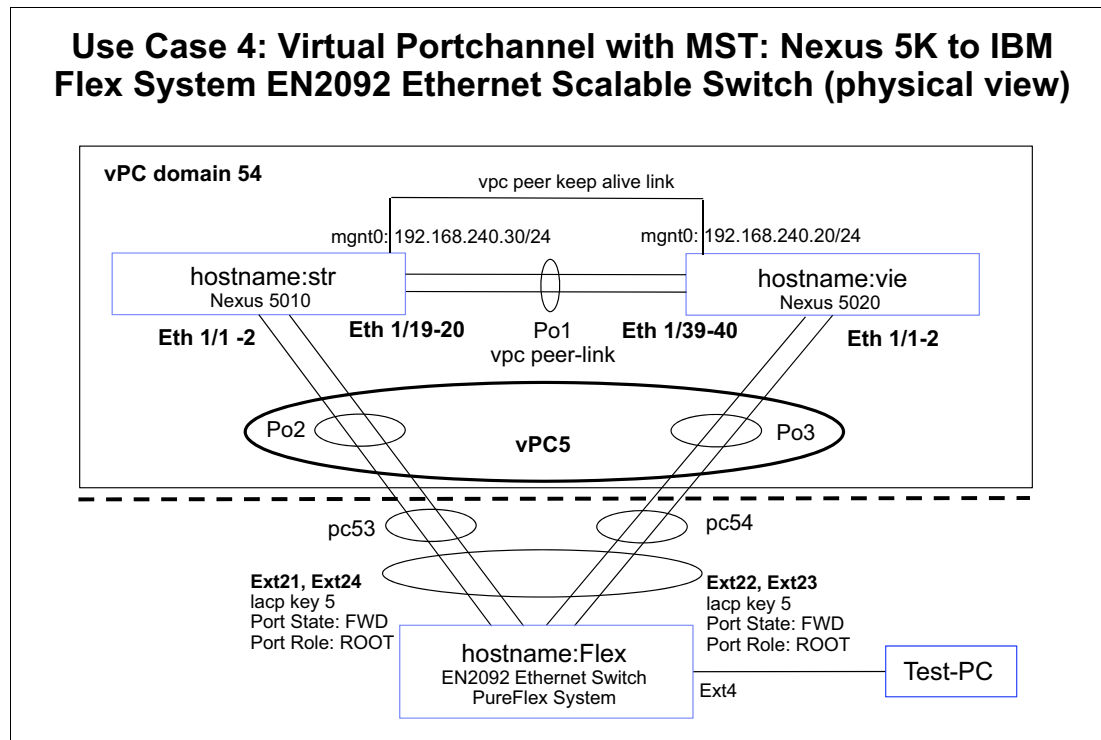


Figure 4-10 Use Case 4: MST with LACP Channeling and vPC

4.5.1 Configuring vPC on STR

To configure vPC, the two Nexus 5000 switches are configured with a vPC peer link in between.

To avoid an active-active scenario if there is a failure, a vPC peer keep-alive link is configured. The MGMT Interfaces are directly connected to the out-of-band keep-alive link. The interface that forms the channel across the Nexus 5000 switches must use the same vPC number on both Nexus 5000 switches (vPC 5 in this case), as shown in Example 4-29 on page 142. The important parameters and details are highlighted in red.

Example 4-29 Use Case 4: vPC Config on STR

```
vpc domain 54
peer-keepalive destination 192.168.240.20 source 192.168.240.30

interface port-channel1
  description T0_VIE_P01
  switchport mode trunk
  switchport access vlan 10
  spanning-tree port type network
  vpc peer-link

interface port-channel2
  description T0_FLEX_EXT21,EXT24
  switchport mode trunk
  switchport access vlan 10
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  vpc 5

interface Ethernet1/1
  description T0_FLEX_EXT24
  switchport mode trunk
  switchport access vlan 10
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  channel-group 2 mode active

interface Ethernet1/2
  description T0_FLEX_EXT21
  switchport mode trunk
  switchport access vlan 10
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  channel-group 2 mode active

...

interface Ethernet1/19
  description T0_VIE_ETH1/39
  switchport mode trunk
  switchport access vlan 10
  channel-group 1 mode active

interface Ethernet1/20
  description T0_VIE_ETH1/40
  switchport mode trunk
  switchport access vlan 10
  channel-group 1 mode active

interface mgmt0
  ip address 192.168.240.30/24
```

4.5.2 Configuring MST on the STR

The commands that are shown in Example 4-30 were used to configure MST on the STR switch.

Example 4-30 Use Case 4: MST Config STR

```
spanning-tree mode mst
spanning-tree mst 1 priority 24576
spanning-tree vlan 10,30 priority 24576
spanning-tree vlan 20,40 priority 28672
spanning-tree mst configuration
  name PureFlex
  revision 10
  instance 1 vlan 10,30
  instance 2 vlan 20,40
```

4.5.3 Configuring vPC on VIE

The commands that are shown Example 4-31 were used to configure vPC on the VIE switch. The important parameters and details are highlighted in red.

Example 4-31 Use Case 4: vPC Config VIE

```
vpc domain 54
  peer-keepalive destination 192.168.240.30 source 192.168.240.20

interface port-channel1
  description TO_STR_P01
  switchport mode trunk
  switchport access vlan 10
  spanning-tree port type network
  vpc peer-link

interface port-channel3
  description TO_FLEX_EXT22,EXT23
  switchport mode trunk
  switchport access vlan 10
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  vpc 5

interface Ethernet1/1
  description TO_FLEX_EXT22
  switchport mode trunk
  switchport access vlan 10
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  channel-group 3 mode active

interface Ethernet1/2
  description TO_FLEX_EXT23
  switchport mode trunk
  switchport access vlan 10
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  channel-group 3 mode active
```

```
...  
  
interface Ethernet1/39  
  description TO_STR_ETH1/19  
  switchport mode trunk  
  switchport access vlan 10  
  channel-group 1 mode active  
  
interface Ethernet1/40  
  description TO_STR_ETH1/19  
  switchport mode trunk  
  switchport access vlan 10  
  channel-group 1 mode active  
  
...  
interface mgmt0  
  no snmp trap link-status  
  vrf member management  
  ip address 192.168.240.20/24
```

4.5.4 Configuring MST on VIE

The commands that are shown Example 4-32 were used to configure MST on the VIE switch.

Example 4-32 Use Case 4: MST Config VIE

```
spanning-tree mode mst  
spanning-tree mst 2 priority 24576  
spanning-tree vlan 10,30 priority 28672  
spanning-tree vlan 20,40 priority 24576  
spanning-tree mst configuration  
  name PureFlex  
  revision 10  
  instance 1 vlan 10,30  
  instance 2 vlan 20,40
```

4.5.5 Reviewing the Flex System switch configuration

The Flex System switch is unaware of vPC. The EN2092, like any end system, sees only one Nexus switch, as shown in Example 4-33.

Example 4-33 Use Case 4: Flex System switch

```
interface port EXT21
    name "TO_STR_ETH1/2"
    tagging
    pvid 10
    exit
!
interface port EXT22
    name "TO_VIE_ETH1/1"
    tagging
    pvid 10
    exit
!
interface port EXT23
    name "TO_VIE_ETH1/2"
    tagging
    pvid 10
    exit
!
interface port EXT24
    name "TO_STR_ETH1/1"
    tagging
    pvid 10
    exit
!
...
!
interface port EXT21
    lacp mode active
    lacp key 5
!
interface port EXT22
    lacp mode active
    lacp key 5
!
interface port EXT23
    lacp mode active
    lacp key 5
!
interface port EXT24
    lacp mode active
    lacp key 5
!
```

4.5.6 Configuring MST on the Flex System switch

The commands that are shown Example 4-34 were used to configure MST on the Flex System switch.

Example 4-34 Use Case 4: MST Config Flex

```
spanning-tree mstp version 10
spanning-tree mstp name "PureFlex"
spanning-tree mode mst
spanning-tree mstp cist-add-vlan 1
spanning-tree mstp cist-add-vlan 4095
!
spanning-tree stp 1 vlan 10
spanning-tree stp 1 vlan 30

spanning-tree stp 2 vlan 20
spanning-tree stp 2 vlan 40
```

4.5.7 Logical view

Figure 4-11 shows the logical view of the setup. To the end system (the IBM Flex System switch), the two Cisco Nexus 5000 switches looks like one switch.

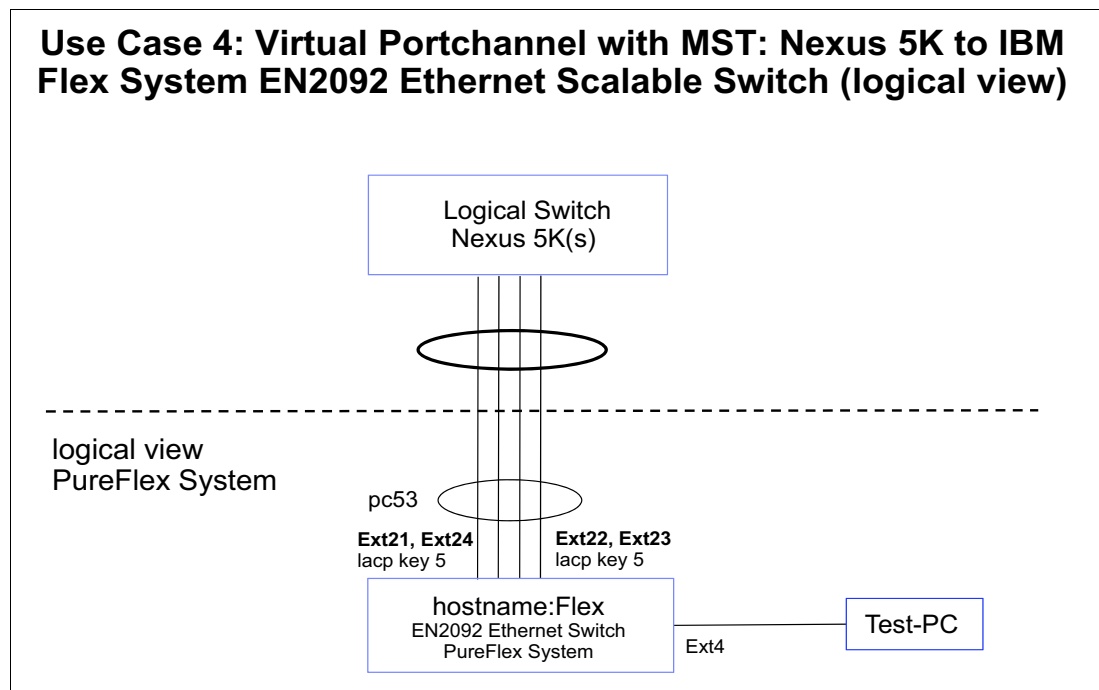


Figure 4-11 Use Case 4: Logical view

4.5.8 Verifying the configuration

We used the **show** commands that are shown in Example 4-35 to verify the vPC configuration that was used on the Nexus 5000 switches. The output helps visualize the setup. The important parameters and details are highlighted in red.

Example 4-35 Use Case 4: Verify the configuration

```
str# show vpc peer-keepalive
```

```
vPC keep-alive status      : peer is alive
--Peer is alive for       : (3417) seconds, (551) msec
--Send status             : Success
--Last send at           : 2012.05.23 19:14:17 134 ms
--Sent on interface       : mgmt0
--Receive status          : Success
--Last receive at         : 2012.05.23 19:14:16 992 ms
--Received on interface   : mgmt0
--Last update from peer   : (0) seconds, (753) msec
```

```
vPC Keep-alive parameters
--Destination             : 192.168.240.20
--Keepalive interval      : 1000 msec
--Keepalive timeout       : 5 seconds
--Keepalive hold timeout  : 3 seconds
--Keepalive vrf           : management
--Keepalive udp port      : 3200
--Keepalive tos           : 192
```

```
str# show vpc brief
```

Legend:

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

```
vPC domain id             : 54
Peer status                : peer adjacency formed ok
vPC keep-alive status      : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status  : success
vPC role                   : primary
Number of vPCs configured  : 1
Peer Gateway               : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
```

```
vPC Peer-link status
```

id	Port	Status	Active vlans
1	Po1	up	1,10,20,30,40

```
vPC status
```

id	Port	Status	Consistency	Reason	Active vlans
5	Po2	up	success	success	10,20,30,40

```
str# show vpc consistency-parameters global
```

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
-----	----	-----	-----
QoS	2	([], [3], [], [], [], [])	([], [3], [], [], [], [])
Network QoS (MTU)	2	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network Qos (Pause)	2	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	MST	MST
STP Disabled	1	None	None
STP MST Region Name	1	PureFlex	PureFlex
STP MST Region Revision	1	10	10
STP MST Region Instance to VLAN Mapping	1		
STP Loopguard	1	Disabled	Disabled
STP Bridge Assurance	1	Enabled	Enabled
STP Port Type, Edge BPDUGuard	1	Normal, Disabled, Disabled	Normal, Disabled, Disabled
STP MST Simulate PVST	1	Enabled	Enabled
Allowed VLANs	-	1,10,20,30,40	1,10,20,30,40
Local suspended VLANs	-	-	-

str# **show vpc consistency-parameters interface po1**

Note: **** Global type-1 parameters will be displayed for peer-link ****

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
-----	----	-----	-----
QoS	2	([], [3], [], [], [], [])	([], [3], [], [], [], [])
Network QoS (MTU)	2	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network Qos (Pause)	2	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	MST	MST
STP Disabled	1	None	None
STP MST Region Name	1	PureFlex	PureFlex
STP MST Region Revision	1	10	10
STP MST Region Instance to VLAN Mapping	1		
STP Loopguard	1	Disabled	Disabled
STP Bridge Assurance	1	Enabled	Enabled
STP Port Type, Edge BPDUGuard	1	Normal, Disabled, Disabled	Normal, Disabled, Disabled
STP MST Simulate PVST	1	Enabled	Enabled
Allowed VLANs	-	1,10,20,30,40	1,10,20,30,40

Local suspended VLANs - - -

str# show vpc consistency-parameters vlan

Name	Type	Reason Code	Pass Vlans
-----	---	-----	-----
STP Mode	1	success	0-4095
STP Disabled	1	success	0-4095
STP MST Region Name	1	success	0-4095
STP MST Region Revision	1	success	0-4095
STP MST Region Instance to	1	success	0-4095
VLAN Mapping			
STP Loopguard	1	success	0-4095
STP Bridge Assurance	1	success	0-4095
STP Port Type, Edge	1	success	0-4095
BPDUFILTER, Edge BPDUGuard			
STP MST Simulate PVST	1	success	0-4095
Pass Vlans	-		0-4095

str# show vpc consistency-parameters vpc 5

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
-----	---	-----	-----
Shut Lan	1	No	No
STP Port Type	1	Default	Default
STP Port Guard	1	None	None
STP MST Simulate PVST	1	Default	Default
lag-id	1	[(7f9b, 0-23-4-ee-be-36, 8005, 0, 0), (8000, 8-17-f4-76-78-0, 5, 0, 0)]	[(7f9b, 0-23-4-ee-be-36, 8005, 0, 0), (8000, 8-17-f4-76-78-0, 5, 0, 0)]
mode	1	active	active
Speed	1	10 Gb/s	10 Gb/s
Duplex	1	full	full
Port Mode	1	trunk	trunk
Native Vlan	1	10	10
MTU	1	1500	1500
Admin port mode	1		
Allowed VLANs	-	10,20,30,40	10,20,30,40
Local suspended VLANs	-	-	-

str# show vpc

Legend:

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

```

vPC domain id          : 54
Peer status             : peer adjacency formed ok
vPC keep-alive status   : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role                : primary
Number of vPCs configured : 1
Peer Gateway            : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled

```

```
vPC Peer-link status
-----
id   Port   Status Active vlans
--   -
1    Po1    up    1,10,20,30,40

vPC status
-----
id   Port   Status Consistency Reason           Active vlans
--   -
5    Po2    up    success  success           10,20,30,40

str# show vpc 5

vPC status
-----
id   Port   Status Consistency Reason           Active vlans
--   -
5    Po2    up    success  success           10,20,30,40

str#
```

4.5.9 Verifying the vPC configuration on VIE

The commands that are shown Example 4-36 were used to verify the vPC configuration of the VIE switch.

Example 4-36 Output of show commands on VIE

```
vie# show vpc peer-keepalive

vPC keep-alive status           : peer is alive
--Peer is alive for             : (3289) seconds, (742) msec
--Send status                   : Success
--Last send at                  : 2012.05.23 19:12:07 422 ms
--Sent on interface             : mgmt0
--Receive status                : Success
--Last receive at               : 2012.05.23 19:12:07 780 ms
--Received on interface         : mgmt0
--Last update from peer        : (0) seconds, (559) msec

vPC Keep-alive parameters
--Destination                   : 192.168.240.30
--Keepalive interval            : 1000 msec
--Keepalive timeout              : 5 seconds
--Keepalive hold timeout        : 3 seconds
--Keepalive vrf                 : management
--Keepalive udp port            : 3200
--Keepalive tos                 : 192

vie# show vpc brief
Legend:
      (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                   : 54
Peer status                     : peer adjacency formed ok
vPC keep-alive status           : peer is alive
Configuration consistency status: success
```

```

Per-vlan consistency status      : success
Type-2 consistency status       : success
vPC role                        : secondary
Number of vPCs configured       : 1
Peer Gateway                    : Disabled
Dual-active excluded VLANs      : -
Graceful Consistency Check      : Enabled

```

vPC Peer-link status

id	Port	Status	Active vlans
1	Po1	up	1,10,20,30,40

vPC status

id	Port	Status	Consistency	Reason	Active vlans
5	Po3	up	success	success	10,20,30,40

vie# **show vpc**

Legend:

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

```

vPC domain id                  : 54
Peer status                    : peer adjacency formed ok
vPC keep-alive status          : peer is alive
Configuration consistency status: success
Per-vlan consistency status     : success
Type-2 consistency status      : success
vPC role                       : secondary
Number of vPCs configured      : 1
Peer Gateway                   : Disabled
Dual-active excluded VLANs     : -
Graceful Consistency Check     : Enabled

```

vPC Peer-link status

id	Port	Status	Active vlans
1	Po1	up	1,10,20,30,40

vPC status

id	Port	Status	Consistency	Reason	Active vlans
5	Po3	up	success	success	10,20,30,40

vie# **show vpc consistency-parameters global**

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
QoS	2	([], [3], [], [], [], [])	([], [3], [], [], [], [])
Network QoS (MTU)	2	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network Qos (Pause)	2	(F, T, F, F, F, F)	(F, T, F, F, F, F)

Input Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	MST	MST
STP Disabled	1	None	None
STP MST Region Name	1	PureFlex	PureFlex
STP MST Region Revision	1	10	10
STP MST Region Instance to VLAN Mapping	1		
STP Loopguard	1	Disabled	Disabled
STP Bridge Assurance	1	Enabled	Enabled
STP Port Type, Edge BPDUGuard	1	Normal, Disabled, Disabled	Normal, Disabled, Disabled
STP MST Simulate PVST	1	Enabled	Enabled
Allowed VLANs	-	1,10,20,30,40	1,10,20,30,40
Local suspended VLANs	-	-	-

vie# **show vpc consistency-parameters interface port-channel 1**

Note: **** Global type-1 parameters will be displayed for peer-link ****

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
-----	----	-----	-----
QoS	2	([], [3], [], [], [], [])	([], [3], [], [], [], [])
Network QoS (MTU)	2	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network Qos (Pause)	2	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	MST	MST
STP Disabled	1	None	None
STP MST Region Name	1	PureFlex	PureFlex
STP MST Region Revision	1	10	10
STP MST Region Instance to VLAN Mapping	1		
STP Loopguard	1	Disabled	Disabled
STP Bridge Assurance	1	Enabled	Enabled
STP Port Type, Edge BPDUGuard	1	Normal, Disabled, Disabled	Normal, Disabled, Disabled
STP MST Simulate PVST	1	Enabled	Enabled
Allowed VLANs	-	1,10,20,30,40	1,10,20,30,40
Local suspended VLANs	-	-	-

vie# **show vpc consistency-parameters vlan**

Name	Type	Reason Code	Pass Vlans
-----	----	-----	-----
STP Mode	1	success	0-4095
STP Disabled	1	success	0-4095
STP MST Region Name	1	success	0-4095
STP MST Region Revision	1	success	0-4095


```

STP MST Region Instance to 1      success      0-4095
VLAN Mapping
STP Loopguard                1      success      0-4095
STP Bridge Assurance         1      success      0-4095
STP Port Type, Edge          1      success      0-4095
BPDUFilter, Edge BPDUGuard
STP MST Simulate PVST        1      success      0-4095
Pass Vlans                    -

```

```
vie# show vpc consistency-parameters vpc 5
```

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
Shut Lan	1	No	No
STP Port Type	1	Default	Default
STP Port Guard	1	None	None
STP MST Simulate PVST	1	Default	Default
lag-id	1	[(7f9b, 0-23-4-ee-be-36, 8005, 0, 0), (8000, 8-17-f4-76-78-0, 5, 0, 0)]	[(7f9b, 0-23-4-ee-be-36, 8005, 0, 0), (8000, 8-17-f4-76-78-0, 5, 0, 0)]
mode	1	active	active
Speed	1	10 Gb/s	10 Gb/s
Duplex	1	full	full
Port Mode	1	trunk	trunk
Native Vlan	1	10	10
MTU	1	1500	1500
Admin port mode	1		
Allowed VLANs	-	10,20,30,40	10,20,30,40
Local suspended VLANs	-	-	-

```
vie# show vpc 5
```

vPC status

id	Port	Status	Consistency	Reason	Active vlans
5	Po3	up	success	success	10,20,30,40

```
vie#
```

The Flex System switch now includes one port channel that consists of four links to the Cisco switches, instead of two port channels that consist of two links each to two N5000 switches, as shown in Example 4-37 on page 154. The vPC 5 on Cisco does not need to be the same vPC 5 that was used in the Flex System configuration. These values are significant only to Cisco and the IBM Flex System switch.

Example 4-37 Output of show lacp command

```
Flex#sh lacp aggregator 5
Aggregator Id 5
-----
Aggregator MAC address - 08:17:f4:76:78:89
Actor System Priority - 32768
Actor System ID - 08:17:f4:76:78:00
Individual - FALSE
Actor Oper Key - 5
Partner System Priority - 32667
Partner System ID - 00:23:04:ee:be:36
Partner Oper Key - 32773
ready - TRUE
Min-Links - 1
Number of Ports in aggr - 4
index 0 port EXT24
index 1 port EXT21
index 2 port EXT22
index 3 port EXT23
```

MST spanning tree is still configured, as shown in Example 4-38. In contrast to the configurations that do not include vPC, all four ports are in spanning tree status forwarding because they all belong to the same LCAP channel.

Example 4-38 Output of show spanning-tree commands

```
Flex#sh spanning-tree
-----
Pvst+ compatibility mode enabled

Mstp Digest: 0xe821ccee7501115289b37c79a72e07c9

-----
Spanning Tree Group 1: On (MSTP)
VLANs MAPPED: 10 30
VLANs: 10 30

Current Root:          Path-Cost Port
6000 00:05:9b:7b:84:3c 490 EXT21

Parameters: Priority Aging Topology Change Counts
             61440    300             21

Port      Prio    Cost    State    Role Designated Bridge    Des Port    Type
-----
EXT4      128      20000! FWD      DESG f000-08:17:f4:76:78:00 8020        P2P
EXT21 (pc53) 128      490!+ FWD      ROOT 6000-00:23:04:ee:be:36 9001        P2P
EXT22 (pc53) 128      490!+ FWD      ROOT 6000-00:23:04:ee:be:36 9001        P2P
EXT23 (pc53) 128      490!+ FWD      ROOT 6000-00:23:04:ee:be:36 9001        P2P
EXT24 (pc53) 128      490!+ FWD      ROOT 6000-00:23:04:ee:be:36 9001        P2P
! = Automatic path cost.
+ = Portchannel cost, not the individual port cost.

-----
Spanning Tree Group 2: On (MSTP)
VLANs MAPPED: 20 40
VLANs: 20 40

Current Root:          Path-Cost Port
```

6000 00:0d:ec:a3:8f:bc 1490 EXT21

Parameters: Priority Aging Topology Change Counts
61440 300 18

Port	Prio	Cost	State	Role	Designated Bridge	Des Port	Type
EXT4	128	20000!	FWD	DESG	f000-08:17:f4:76:78:00	8020	P2P
EXT21 (pc53)	128	490!+	FWD	ROOT	8000-00:23:04:ee:be:36	9001	P2P
EXT22 (pc53)	128	490!+	FWD	ROOT	8000-00:23:04:ee:be:36	9001	P2P
EXT23 (pc53)	128	490!+	FWD	ROOT	8000-00:23:04:ee:be:36	9001	P2P
EXT24 (pc53)	128	str#	show spanning-tree				

No spanning tree instance exists.

str# show vpc peer-keep

vPC keep-alive status : peer is alive
--Peer is alive for : (5012) seconds, (175) msec
--Send status : Success
--Last send at : 2012.05.23 19:40:51 754 ms
--Sent on interface : mgmt0
--Receive status : Success
--Last receive at : 2012.05.23 19:40:51 941 ms
--Received on interface : mgmt0
--Last update from peer : (0) seconds, (428) msec

vPC Keep-alive parameters
--Destination : 192.168.240.20
--Keepalive interval : 1000 msec
--Keepalive timeout : 5 seconds
--Keepalive hold timeout : 3 seconds
--Keepalive vrf : management
--Keepalive udp port : 3200
--Keepalive tos : 192

str# show vpc brief

Legend:

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

vPC domain id : 54
Peer status : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role : primary
Number of vPCs configured : 1
Peer Gateway : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled

vPC Peer-link status

id	Port	Status	Active vlans
1	Po1	up	1,10,20,30,40

vPC status

id	Port	Status	Consistency Reason	Active vlans
----	------	--------	--------------------	--------------

```
-----
5      Po2      up      success      success      10,20,30,40
```

```
str# show vpc consistency-parameters global
```

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
QoS	2	([], [3], [], [], [], [])	([], [3], [], [], [], [])
Network QoS (MTU)	2	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network Qos (Pause)	2	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	Rapid-PVST	Rapid-PVST
STP Disabled	1	VLANs 1,10,20,30,40	VLANs 1,10,20,30,40
STP MST Region Name	1	PureFlex	PureFlex
STP MST Region Revision	1	10	10
STP MST Region Instance to VLAN Mapping	1		
STP Loopguard	1	Disabled	Disabled
STP Bridge Assurance	1	Enabled	Enabled
STP Port Type, Edge BPDUGuard	1	Normal, Disabled, Disabled	Normal, Disabled, Disabled
STP MST Simulate PVST	1	Enabled	Enabled
Allowed VLANs	-	1,10,20,30,40	1,10,20,30,40
Local suspended VLANs	-	-	-

```
str# show vpc consistency-parameters interface po1
```

Note: **** Global type-1 parameters will be displayed for peer-link ****

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
QoS	2	([], [3], [], [], [], [])	([], [3], [], [], [], [])
Network QoS (MTU)	2	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network Qos (Pause)	2	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	Rapid-PVST	Rapid-PVST
STP Disabled	1	VLANs 1,10,20,30,40	VLANs 1,10,20,30,40
STP MST Region Name	1	PureFlex	PureFlex
STP MST Region Revision	1	10	10
STP MST Region Instance to VLAN Mapping	1		
STP Loopguard	1	Disabled	Disabled

STP Bridge Assurance	1	Enabled	Enabled
STP Port Type, Edge	1	Normal, Disabled,	Normal, Disabled,
BPDUFILTER, Edge BPDUGuard		Disabled	Disabled
STP MST Simulate PVST	1	Enabled	Enabled
Allowed VLANs	-	1,10,20,30,40	1,10,20,30,40
Local suspended VLANs	-	-	-

str# show vpc consistency-parameters vlan

Name	Type	Reason Code	Pass Vlans
-----	----	-----	-----
STP Mode	1	success	0-4095
STP Disabled	1	success	0-4095
STP MST Region Name	1	success	0-4095
STP MST Region Revision	1	success	0-4095
STP MST Region Instance to	1	success	0-4095
VLAN Mapping			
STP Loopguard	1	success	0-4095
STP Bridge Assurance	1	success	0-4095
STP Port Type, Edge	1	success	0-4095
BPDUFILTER, Edge BPDUGuard			
STP MST Simulate PVST	1	success	0-4095
Pass Vlans	-		0-4095

str# show vpc consistency-parameters vpc 5

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
-----	----	-----	-----
Shut Lan	1	No	No
STP Port Type	1	Default	Default
STP Port Guard	1	None	None
STP MST Simulate PVST	1	Default	Default
lag-id	1	[(7f9b, 0-23-4-ee-be-36, 8005, 0, 0), (8000, 8-17-f4-76-78-0, 5, 0, 0)]	[(7f9b, 0-23-4-ee-be-36, 8005, 0, 0), (8000, 8-17-f4-76-78-0, 5, 0, 0)]
mode	1	active	active
Speed	1	10 Gb/s	10 Gb/s
Duplex	1	full	full
Port Mode	1	trunk	trunk
Native Vlan	1	10	10
MTU	1	1500	1500
Admin port mode	1		
Allowed VLANs	-	10,20,30,40	10,20,30,40
Local suspended VLANs	-	-	-

str# show vpc

Legend:

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

```

vPC domain id          : 54
Peer status             : peer adjacency formed ok
vPC keep-alive status   : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role                : primary

```

```

Number of vPCs configured      : 1
Peer Gateway                   : Disabled
Dual-active excluded VLANs     : -
Graceful Consistency Check     : Enabled

```

vPC Peer-link status

```

-----
id   Port   Status Active vlans
--   -
1    Po1    up     1,10,20,30,40

```

vPC status

```

-----
id   Port   Status Consistency Reason           Active vlans
-----
5    Po2    up     success    success           10,20,30,40

```

str# show spanning-tree

No spanning tree instance exists.

str# show vpc peer-keep

```

vPC keep-alive status      : peer is alive
--Peer is alive for       : (5012) seconds, (175) msec
--Send status             : Success
--Last send at            : 2012.05.23 19:40:51 754 ms
--Sent on interface       : mgmt0
--Receive status          : Success
--Last receive at         : 2012.05.23 19:40:51 941 ms
--Received on interface   : mgmt0
--Last update from peer   : (0) seconds, (428) msec

```

vPC Keep-alive parameters

```

--Destination              : 192.168.240.20
--Keepalive interval       : 1000 msec
--Keepalive timeout        : 5 seconds
--Keepalive hold timeout   : 3 seconds
--Keepalive vrf            : management
--Keepalive udp port       : 3200
--Keepalive tos            : 192

```

str# show vpc brief

Legend:

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

```

vPC domain id              : 54
Peer status                 : peer adjacency formed ok
vPC keep-alive status      : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status  : success
vPC role                   : primary
Number of vPCs configured  : 1
Peer Gateway               : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled

```

vPC Peer-link status

```

-----
id   Port   Status Active vlans
--   -

```

```
1    Po1    up    1,10,20,30,40
```

```
vPC status
```

```
-----
id      Port      Status Consistency Reason      Active vlans
-----
5       Po2       up      success      success      10,20,30,40
```

```
str# show vpc consistency-parameters global
```

```
Legend:
```

```
Type 1 : vPC will be suspended in case of mismatch
```

Name	Type	Local Value	Peer Value
QoS	2	([], [3], [], [], [], [])	([], [3], [], [], [], [])
Network QoS (MTU)	2	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network QoS (Pause)	2	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	Rapid-PVST	Rapid-PVST
STP Disabled	1	VLANs 1,10,20,30,40	VLANs 1,10,20,30,40
STP MST Region Name	1	PureFlex	PureFlex
STP MST Region Revision	1	10	10
STP MST Region Instance to VLAN Mapping	1		
STP Loopguard	1	Disabled	Disabled
STP Bridge Assurance	1	Enabled	Enabled
STP Port Type, Edge	1	Normal, Disabled,	Normal, Disabled,
BPDUGuard	1	Disabled	Disabled
STP MST Simulate PVST	1	Enabled	Enabled
Allowed VLANs	-	1,10,20,30,40	1,10,20,30,40
Local suspended VLANs	-	-	-

```
str# show vpc consistency-parameters interface po1
```

```
Note: **** Global type-1 parameters will be displayed for peer-link ****
```

```
Legend:
```

```
Type 1 : vPC will be suspended in case of mismatch
```

Name	Type	Local Value	Peer Value
QoS	2	([], [3], [], [], [], [])	([], [3], [], [], [], [])
Network QoS (MTU)	2	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network QoS (Pause)	2	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	Rapid-PVST	Rapid-PVST
STP Disabled	1	VLANs 1,10,20,30,40	VLANs 1,10,20,30,40
STP MST Region Name	1	PureFlex	PureFlex

```

STP MST Region Revision      1      10      10
STP MST Region Instance to 1
  VLAN Mapping
STP Loopguard                1      Disabled      Disabled
STP Bridge Assurance         1      Enabled      Enabled
STP Port Type, Edge          1      Normal, Disabled, Normal, Disabled,
BPDUFilter, Edge BPDUGuard   Disabled      Disabled
STP MST Simulate PVST        1      Enabled      Enabled
Allowed VLANs                -      1,10,20,30,40 1,10,20,30,40
Local suspended VLANs        -      -      -
str# show vpc consistency-parameters vlan

```

Name	Type	Reason Code	Pass Vlans
STP Mode	1	success	0-4095
STP Disabled	1	success	0-4095
STP MST Region Name	1	success	0-4095
STP MST Region Revision	1	success	0-4095
STP MST Region Instance to	1	success	0-4095
VLAN Mapping			
STP Loopguard	1	success	0-4095
STP Bridge Assurance	1	success	0-4095
STP Port Type, Edge	1	success	0-4095
BPDUFilter, Edge BPDUGuard			
STP MST Simulate PVST	1	success	0-4095
Pass Vlans	-		0-4095

```
str# show vpc consistency-parameters vpc 5
```

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
Shut Lan	1	No	No
STP Port Type	1	Default	Default
STP Port Guard	1	None	None
STP MST Simulate PVST	1	Default	Default
lag-id	1	[(7f9b, 0-23-4-ee-be-36, 8005, 0, 0), (8000, 8-17-f4-76-78-0, 5, 0, 0)]	[(7f9b, 0-23-4-ee-be-36, 8005, 0, 0), (8000, 8-17-f4-76-78-0, 5, 0, 0)]
mode	1	active	active
Speed	1	10 Gb/s	10 Gb/s
Duplex	1	full	full
Port Mode	1	trunk	trunk
Native Vlan	1	10	10
MTU	1	1500	1500
Admin port mode	1		
Allowed VLANs	-	10,20,30,40	10,20,30,40
Local suspended VLANs	-	-	-

```
str# show vpc
```

Legend:

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

```

vPC domain id      : 54
Peer status        : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status: success

```



```

Per-vlan consistency status      : success
Type-2 consistency status       : success
vPC role                        : primary
Number of vPCs configured       : 1
Peer Gateway                    : Disabled
Dual-active excluded VLANs      : -
Graceful Consistency Check      : Enabled

```

vPC Peer-link status

```

-----
id  Port  Status Active vlans
--  ---  -
1   Po1   up      1,10,20,30,40

```

vPC status

```

-----
id    Port      Status Consistency Reason          Active vlans
-----
5     Po2       up      success    success                    10,20,30,40

```

```

str#      490!+ FWD  ROOT 8000-00:23:04:ee:be:36    9001    P2P

```

! = Automatic path cost.

+ = Portchannel cost, not the individual port cost.

4.6 Use Case 5: LACP Channeling and vPC without spanning tree

We can switch off spanning tree because we now have two switches that are connected with one cable. The physical setup still consists of two Nexus 5000 switches and four 10 GE links, as shown in Figure 4-12.

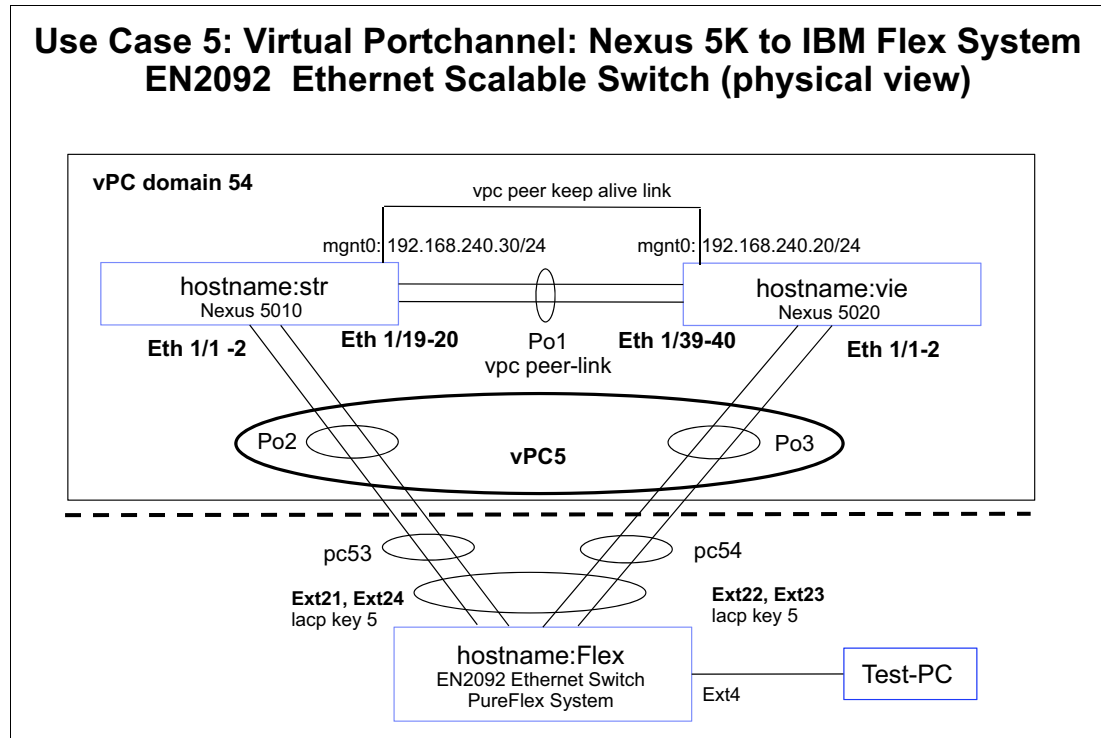


Figure 4-12 Use Case 5

We disabled STP for VLANs 10, 20, 30, and 40.

After STP is switched off and LACP and vPC are used, the logical setup looks like two switches that are connected by one cable. Because of this configuration, there is no need for an STP to run to block redundant links, as shown in Figure 4-13 on page 163.

Use Case 5: Virtual Portchannel, no STP: Nexus 5K to Flex System EN2092 Ethernet Scalable Switch (logical view)

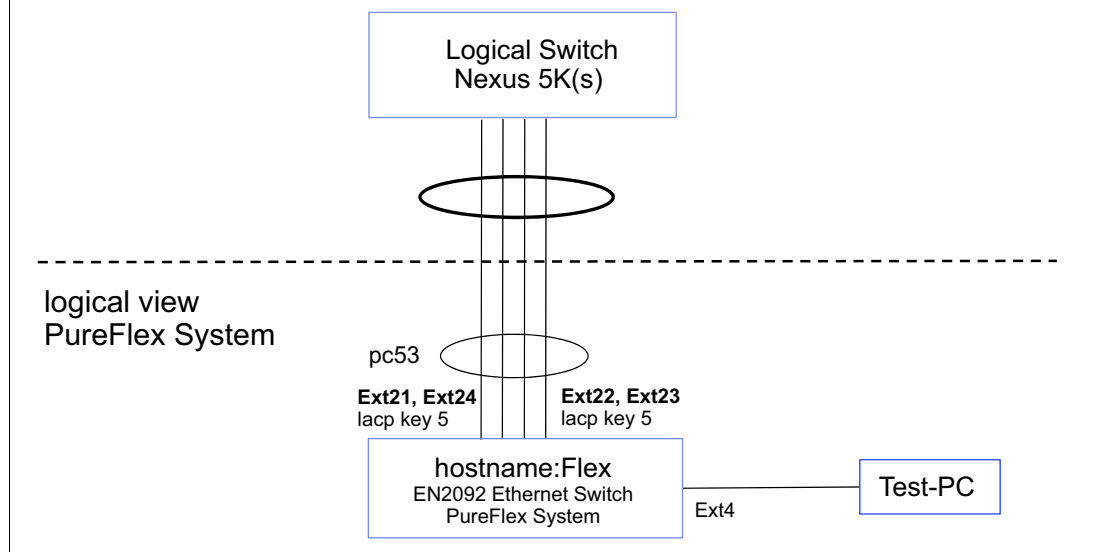


Figure 4-13 Use Case 5: Logical view

4.6.1 Configuring vPC on STR

The commands that are shown Example 4-39 were used to configure vPC on STR. The important parameters and details are highlighted in red.

Example 4-39 Use Case 5

```
str# show spanning-tree
No spanning tree instance exists.

str# show vpc peer-keepalive

vPC keep-alive status          : peer is alive
--Peer is alive for           : (5012) seconds, (175) msec
--Send status                  : Success
--Last send at                 : 2012.05.23 19:40:51 754 ms
--Sent on interface            : mgmt0
--Receive status               : Success
--Last receive at              : 2012.05.23 19:40:51 941 ms
--Received on interface        : mgmt0
--Last update from peer        : (0) seconds, (428) msec

vPC Keep-alive parameters
--Destination                   : 192.168.240.20
--Keepalive interval            : 1000 msec
--Keepalive timeout             : 5 seconds
--Keepalive hold timeout        : 3 seconds
--Keepalive vrf                 : management
--Keepalive udp port            : 3200
--Keepalive tos                 : 192
str# show vpc brief
Legend:
```

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

```
vPC domain id          : 54
Peer status            : peer adjacency formed ok
vPC keep-alive status  : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role               : primary
Number of vPCs configured : 1
Peer Gateway           : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
```

vPC Peer-link status

```
-----
id   Port   Status Active vlans
--   -
1    Po1    up     1,10,20,30,40
```

vPC status

```
-----
id   Port   Status Consistency Reason           Active vlans
-----
5    Po2    up     success    success           10,20,30,40
```

str# show vpc consistency-parameters global

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
QoS	2	([], [3], [], [], [], [])	([], [3], [], [], [], [])
Network QoS (MTU)	2	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network Qos (Pause)	2	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	Rapid-PVST	Rapid-PVST
STP Disabled	1	VLANs 1,10,20,30,40	VLANs 1,10,20,30,40
STP MST Region Name	1	PureFlex	PureFlex
STP MST Region Revision	1	10	10
STP MST Region Instance to	1		
VLAN Mapping			
STP Loopguard	1	Disabled	Disabled
STP Bridge Assurance	1	Enabled	Enabled
STP Port Type, Edge	1	Normal, Disabled,	Normal, Disabled,
BPDUGuard	1	Disabled	Disabled
STP MST Simulate PVST	1	Enabled	Enabled
Allowed VLANs	-	1,10,20,30,40	1,10,20,30,40
Local suspended VLANs	-	-	-

str# show vpc consistency-parameters int po1

Note: **** Global type-1 parameters will be displayed for peer-link ****

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
-----	----	-----	-----
QoS	2	([], [3], [], [], [], [])	([], [3], [], [], [], [])
Network QoS (MTU)	2	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network QoS (Pause)	2	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	Rapid-PVST	Rapid-PVST
STP Disabled	1	VLANs 1,10,20,30,40	VLANs 1,10,20,30,40
STP MST Region Name	1	PureFlex	PureFlex
STP MST Region Revision	1	10	10
STP MST Region Instance to VLAN Mapping	1		
STP Loopguard	1	Disabled	Disabled
STP Bridge Assurance	1	Enabled	Enabled
STP Port Type, Edge BPDUGuard	1	Normal, Disabled, Disabled	Normal, Disabled, Disabled
STP MST Simulate PVST	1	Enabled	Enabled
Allowed VLANs	-	1,10,20,30,40	1,10,20,30,40
Local suspended VLANs	-	-	-

str# **show vpc consistency-parameters vlan**

Name	Type	Reason Code	Pass Vlans
-----	----	-----	-----
STP Mode	1	success	0-4095
STP Disabled	1	success	0-4095
STP MST Region Name	1	success	0-4095
STP MST Region Revision	1	success	0-4095
STP MST Region Instance to VLAN Mapping	1	success	0-4095
STP Loopguard	1	success	0-4095
STP Bridge Assurance	1	success	0-4095
STP Port Type, Edge BPDUGuard	1	success	0-4095
STP MST Simulate PVST	1	success	0-4095
Pass Vlans	-		0-4095

str# **show vpc consistency-parameters vpc 5**

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
-----	----	-----	-----
Shut Lan	1	No	No
STP Port Type	1	Default	Default
STP Port Guard	1	None	None
STP MST Simulate PVST	1	Default	Default
lag-id	1	[(7f9b, 0-23-4-ee-be-36, 8005, 0, 0), (8000, 8-17-f4-76-78-0, 5, 0, 0)]	[(7f9b, 0-23-4-ee-be-36, 8005, 0, 0), (8000, 8-17-f4-76-78-0, 5, 0, 0)]

```

mode                1      active          active
Speed               1      10 Gb/s      10 Gb/s
Duplex              1      full           full
Port Mode           1      trunk          trunk
Native Vlan         1      10           10
MTU                 1      1500        1500
Admin port mode     1
Allowed VLANs       -      10,20,30,40  10,20,30,40
Local suspended VLANs -      -           -

```

str# **show vpc**

Legend:

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

```

vPC domain id       : 54
Peer status          : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role             : primary
Number of vPCs configured : 1
Peer Gateway         : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled

```

vPC Peer-link status

id	Port	Status	Active vlans
1	Po1	up	1,10,20,30,40

vPC status

id	Port	Status	Consistency	Reason	Active vlans
5	Po2	up	success	success	10,20,30,40

str#

4.6.2 Configuring vPC on VIE

The commands that are shown Example 4-40 were used to configure vPC on the VIE switch.

Example 4-40 Use Case 5: vPC config VIE

vie# **show vpc peer-keepalive**

```

vPC keep-alive status      : peer is alive
--Peer is alive for       : (5140) seconds, (176) msec
--Send status              : Success
--Last send at             : 2012.05.23 19:42:58 751 ms
--Sent on interface        : mgmt0
--Receive status           : Success
--Last receive at          : 2012.05.23 19:42:58 563 ms
--Received on interface    : mgmt0
--Last update from peer    : (0) seconds, (210) msec

```

```

vPC Keep-alive parameters
--Destination              : 192.168.240.30

```

```
--Keepalive interval      : 1000 msec
--Keepalive timeout       : 5 seconds
--Keepalive hold timeout  : 3 seconds
--Keepalive vrf           : management
--Keepalive udp port      : 3200
--Keepalive tos           : 192
```

```
vie# show vpc brief
```

Legend:

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

```
vPC domain id      : 54
Peer status        : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role           : secondary
Number of vPCs configured : 1
Peer Gateway       : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
```

vPC Peer-link status

id	Port	Status	Active vlans
1	Po1	up	1,10,20,30,40

vPC status

id	Port	Status	Consistency	Reason	Active vlans
5	Po3	up	success	success	10,20,30,40

```
vie# show vpc consistency-parameters global
```

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
QoS	2	([], [3], [], [], [], [])	([], [3], [], [], [], [])
Network QoS (MTU)	2	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network Qos (Pause)	2	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	Rapid-PVST	Rapid-PVST
STP Disabled	1	VLANs 1,10,20,30,40	VLANs 1,10,20,30,40
STP MST Region Name	1	PureFlex	PureFlex
STP MST Region Revision	1	10	10
STP MST Region Instance to VLAN Mapping	1		
STP Loopguard	1	Disabled	Disabled
STP Bridge Assurance	1	Enabled	Enabled

```

STP Port Type, Edge      1      Normal, Disabled,      Normal, Disabled,
BPDUFilter, Edge BPDUGuard      Disabled      Disabled
STP MST Simulate PVST      1      Enabled      Enabled
Allowed VLANs      -      1,10,20,30,40      1,10,20,30,40
Local suspended VLANs      -      -      -

```

vie# **show vpc consistency-parameters int po 1**

Note: **** Global type-1 parameters will be displayed for peer-link ****

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
QoS	2	([], [3], [], [], [], [])	([], [3], [], [], [], [])
Network QoS (MTU)	2	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0, 0)
Network Qos (Pause)	2	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	2	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute Priority)	2	(F, F, F, F, F, F)	(F, F, F, F, F, F)
STP Mode	1	Rapid-PVST	Rapid-PVST
STP Disabled	1	VLANs 1,10,20,30,40	VLANs 1,10,20,30,40
STP MST Region Name	1	PureFlex	PureFlex
STP MST Region Revision	1	10	10
STP MST Region Instance to VLAN Mapping	1		
STP Loopguard	1	Disabled	Disabled
STP Bridge Assurance	1	Enabled	Enabled
STP Port Type, Edge	1	Normal, Disabled,	Normal, Disabled,
BPDUFilter, Edge BPDUGuard		Disabled	Disabled
STP MST Simulate PVST	1	Enabled	Enabled
Allowed VLANs	-	1,10,20,30,40	1,10,20,30,40
Local suspended VLANs	-	-	-

vie# **show vpc consistency-parameters vlan**

Name	Type	Reason Code	Pass Vlans
STP Mode	1	success	0-4095
STP Disabled	1	success	0-4095
STP MST Region Name	1	success	0-4095
STP MST Region Revision	1	success	0-4095
STP MST Region Instance to VLAN Mapping	1	success	0-4095
STP Loopguard	1	success	0-4095
STP Bridge Assurance	1	success	0-4095
STP Port Type, Edge	1	success	0-4095
BPDUFilter, Edge BPDUGuard			
STP MST Simulate PVST	1	success	0-4095
Pass Vlans	-		0-4095

vie# **show vpc consistency-parameters vpc 5**

Legend:

Type 1 : vPC will be suspended in case of mismatch

Name	Type	Local Value	Peer Value
Shut Lan	1	No	No

STP Port Type	1	Default	Default
STP Port Guard	1	None	None
STP MST Simulate PVST	1	Default	Default
lag-id	1	[(7f9b, 0-23-4-ee-be-36, 8005, 0, 0), (8000, 8-17-f4-76-78-0, 5, 0, 0)]	[(7f9b, 0-23-4-ee-be-36, 8005, 0, 0), (8000, 8-17-f4-76-78-0, 5, 0, 0)]
mode	1	active	active
Speed	1	10 Gb/s	10 Gb/s
Duplex	1	full	full
Port Mode	1	trunk	trunk
Native Vlan	1	10	10
MTU	1	1500	1500
Admin port mode	1		
Allowed VLANs	-	10,20,30,40	10,20,30,40
Local suspended VLANs	-	-	-

4.6.3 Disabling STP on the Flex System switch

The commands that are shown Example 4-41 were used to disable STP on the Flex System switch. The important parameters and details are highlighted in red.

Example 4-41 Use Case 5: Flex System switch

```
!
spanning-tree mstp version 10
spanning-tree mstp name "PureFlex"
spanning-tree mode disable
!
spanning-tree stp 1 vlan 1
spanning-tree stp 1 vlan 10
spanning-tree stp 1 vlan 20
spanning-tree stp 1 vlan 30
spanning-tree stp 1 vlan 40

!
```

Show spanning tree on Flex

```
-----
Spanning Tree is shut down.
```

```
-----
MSTP is not on.
-----
```



Cisco Catalyst 6500 switch connectivity

Many customers still use the Cisco Catalyst 6500 switch in their data center. This chapter describes the use case that we performed with the IBM Flex System chassis and the Catalyst 6500 switch.

5.1 Use Case 1: LACP channeling and vPC without spanning tree

We had only one Catalyst 6500 switch available for this use case. We connected the one Flex System switch to one Catalyst 6500 switch by using four parallel links, as shown in Figure 5-1.

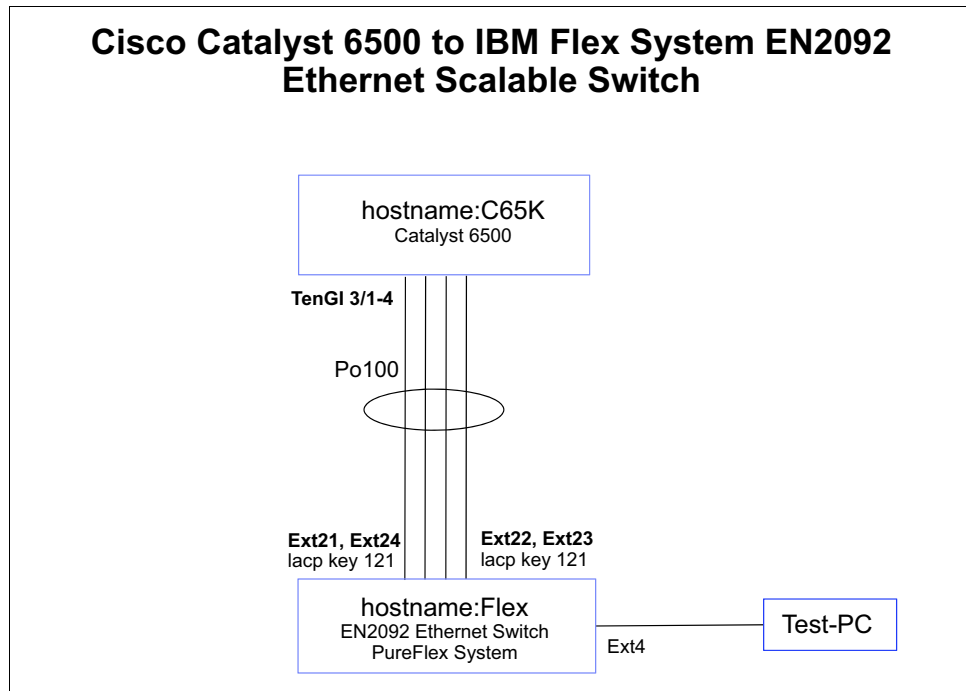


Figure 5-1 Catalyst 6500 Use Case

5.1.1 Catalyst 6500 switch configuration

The Catalyst 6500 switch configuration that was used in this use case is shown in Example 5-1. The important parameters and details are highlighted in red.

Example 5-1 Catalyst 6500 switch configuration

```
lldp run

interface Port-channel100
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  switchport mode trunk
  !
  ...

interface TenGigabitEthernet3/1
  description T0_Flex_EXT21
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10,20,30,40
  switchport mode trunk
  channel-group 100 mode active
```

```

!
interface TenGigabitEthernet3/2
description T0_Flex_EXT22
switchport
switchport trunk encapsulation dot1q
switchport trunk native vlan 10
switchport trunk allowed vlan 10,20,30,40
switchport mode trunk
channel-group 100 mode active
!
interface TenGigabitEthernet3/3
description T0_Flex_EXT23
switchport
switchport trunk encapsulation dot1q
switchport trunk native vlan 10
switchport trunk allowed vlan 10,20,30,40
switchport mode trunk
channel-group 100 mode active
!
interface TenGigabitEthernet3/4
description T0_Flex_EXT24
switchport
switchport trunk encapsulation dot1q
switchport trunk native vlan 10
switchport trunk allowed vlan 10,20,30,40
switchport mode trunk
channel-group 100 mode active

```

C6K#sh lldp neighbors

Capability codes:

(R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
(W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

Device ID	Local Intf	Hold-time	Capability	Port ID
Flex	Te3/1	120	B,R	49
Flex	Te3/4	120	B,R	52
Flex	Te3/3	120	B,R	51
Flex	Te3/2	120	B,R	50
Total entries displayed: 4Te3/2		120	B,R	50

5.1.2 Flex System switch configuration

The Flex System switch configuration that was used in this use case is shown in Example 5-2. The important parameters and details are highlighted in red.

Example 5-2 Flex System switch configuration

```
!  
spanning-tree mstp version 10  
spanning-tree mstp name "PureFlex"  
spanning-tree mode disable  
!  
spanning-tree stp 1 vlan 1  
spanning-tree stp 1 vlan 10  
spanning-tree stp 1 vlan 20  
spanning-tree stp 1 vlan 30  
spanning-tree stp 1 vlan 40  
  
!  
interface port EXT21  
    name "TO_C6K_TEN3/1"  
    tagging  
    pvid 10  
    exit  
!  
interface port EXT22  
    name "TO_C6K_TEN3/2"  
    tagging  
    pvid 10  
    exit  
!  
interface port EXT23  
    name "TO_C6K_TEN3/3"  
    tagging  
    pvid 10  
    exit  
!  
interface port EXT24  
    name "TO_C6K_TEN3/4"  
    tagging  
    pvid 10  
    exit  
!  
...  
!  
interface port EXT21  
    lacp mode active  
    lacp key 121  
!  
interface port EXT22  
    lacp mode active  
    lacp key 121  
!  
interface port EXT23  
    lacp mode active  
    lacp key 121  
!  
interface port EXT24  
    lacp mode active
```

lACP key 121

!
!

Flex#sh lldp remote-device

LLDP Remote Devices Information

LocalPort	Index	Remote Chassis ID	Remote Port	Remote System Name
EXT21	6	00 1a 2f 00 a0 d6	T0_Flex_EXT21	C6K.cisco.com
EXT22	7	00 1a 2f 00 a0 d7	T0_Flex_EXT22	C6K.cisco.com
EXT23	8	00 1a 2f 00 a0 d8	T0_Flex_EXT23	C6K.cisco.com
EXT24	9	00 1a 2f 00 a0 d9	T0_Flex_EXT24	C6K.cisco.com

Flex#sh lacp aggregator

Aggregator Id 49

```
-----
Aggregator MAC address - 08:17:f4:76:78:86
Actor System Priority   - 32768
Actor System ID        - 08:17:f4:76:78:00
Individual              - FALSE
Actor Oper Key          - 121
Partner System Priority - 32768
Partner System ID      - 00:19:07:a9:07:00
Partner Oper Key        - 100
ready                   - TRUE
Min-Links               - 1
Number of Ports in aggr - 4
index 0  port EXT21
index 1  port EXT22
index 2  port EXT23
index 3  port EXT24
```

Flex#sh int status

Alias	Port	Speed	Duplex	Flow Ctrl		Link	Name
-----	----	-----	-----	--TX--	---RX--	-----	-----
...							
EXT21	49	10000	full	no	no	up	T0_C6K_TEN3/1
EXT22	50	10000	full	no	no	up	T0_C6K_TEN3/2
EXT23	51	10000	full	no	no	up	T0_C6K_TEN3/3
EXT24	52	10000	full	no	no	up	T0_C6K_TEN3/4
MGT1	53	1000	full	no	no	up	MGT1

Flex#sh lacp information

port	mode	adminkey	operkey	selected	prio	aggr	trunk	status	minlinks

...									
EXT21	active	121	121	yes	32768	49	53	up	1
EXT22	active	121	121	yes	32768	49	53	up	1
EXT23	active	121	121	yes	32768	49	53	up	1
EXT24	active	121	121	yes	32768	49	53	up	1

Flex#sh lacp aggregator

Aggregator Id 49

Aggregator MAC address - 08:17:f4:76:78:86
Actor System Priority - 32768
Actor System ID - 08:17:f4:76:78:00
Individual - FALSE
Actor Oper Key - 121
Partner System Priority - 32768
Partner System ID - 00:19:07:a9:07:00
Partner Oper Key - 100
ready - TRUE
Min-Links - 1
Number of Ports in aggr - 4
index 0 port EXT21
index 1 port EXT22
index 2 port EXT23
index 3 port EXT24

Flex#**show spanning-tree**

Spanning Tree is shut down.



Troubleshooting

The methodology and commands that are used for troubleshooting connectivity problems are described in this appendix. A sample of network documentation also is provided.

In this Redpaper, the focus thus far has been placed on Layer 2. Therefore, the focus of this appendix is on problems about Ethernet, VLANs, and spanning tree.

In the first part, we describe a useful troubleshooting methodology. In the second part, you find the most common commands to show and verify the status of the configuration, which help you to track down the root cause of your problem. The last part of the appendix shows a sample of network documentation, which is the information you need with which to troubleshoot.

This appendix includes the following topics:

- ▶ Basic troubleshooting for connectivity problems
- ▶ Baseline documentation
- ▶ Firmware update of IBM Flex System network switches

Nexus 5000 switch upgrades: For more information about how to upgrade NX-OS for the Nexus 5000 Series switches, see this website:

http://www.cisco.com/en/US/products/ps9670/products_configuration_example09186a0080b4b9dd.shtml

Basic troubleshooting for connectivity problems

This section describes basic troubleshooting techniques.

Approach

This basic Layer 2 troubleshooting guideline supports you when you are looking for connectivity problems of adjacent devices. These devices are devices that should be able to communicate with each other on Layer 2. This configuration might be two hosts in the same VLAN or a host and its default gateway.

The following symptoms often indicate a problem:

- ▶ Failing application or failing pings between adjacent devices.
- ▶ Address resolution protocol (ARP) failures (missing or “incomplete” ARP entry).
- ▶ Missing packets on the receiving host that are shown with a packet sniffer.

Verify connectivity

Before you start troubleshooting on Layer 2, you should verify the following connectivity configurations on Layer 3:

- ▶ Ping the two devices from each other. Do you receive an Internet Control Message Protocol (ICMP) echo in one or other direction?

If you do not receive an echo, the following causes for a ping failure are possible:

- A Firewall or personal firewall on a host
 - Wrong or missing default gateway (DGW)
 - Wrong IP subnet mask
- ▶ Verify that the ARP caches on the devices. Even if a ping does not work, it is possible that the address resolution protocol (ARP) did work. This status indicates a working Layer 2 link and a problem on the IP level (Layer 3). Even if the ping fails, the ARP entries should be verified.

Determine the Layer 2 path

When you are at the point that your problem seems to be a Layer 2 or Layer 1 problem, you want to reduce the scope of the potential failures. This common troubleshooting method might help you to diagnose your problem.

In the first step, it is useful to determine the expected Layer 2 path that is based on documentation, baselines, and general knowledge of the network. Determining the Layer 2 path shows the path that the traffic is expected to take between the two affected hosts. The analysis results indicate a good starting point for the next steps of gathering information about what is happening in the network, and make it easier to detect abnormal behavior.

Track the traffic flow across the Layer 2 path

The second step is to follow the expected path and verify that the links are up and forwarding traffic. If the actual path is different from the expected path, this conflict can indicate where to proceed with troubleshooting, what links and protocols are involved, and might cause the failure. Often included in this process is comparing the spanning tree topology against the expected Layer 2 topology. If the actual topology differs from the expected, this difference might give some clue about the cause of the problem.

Verification of traffic flows can be done by showing MAC address tables, interface statistics, and so on.

Analyze links

After you find a divergence between the expected and the current traffic path, you should examine the links to determine where the expected path is broken. You can start to target troubleshooting commands to narrow down the root cause of the problem. Even if you cannot figure out on yourself the root cause, you can establish a good base of information and documentation for problem escalation.

Figure A-1 shows an overview of the troubleshooting steps.

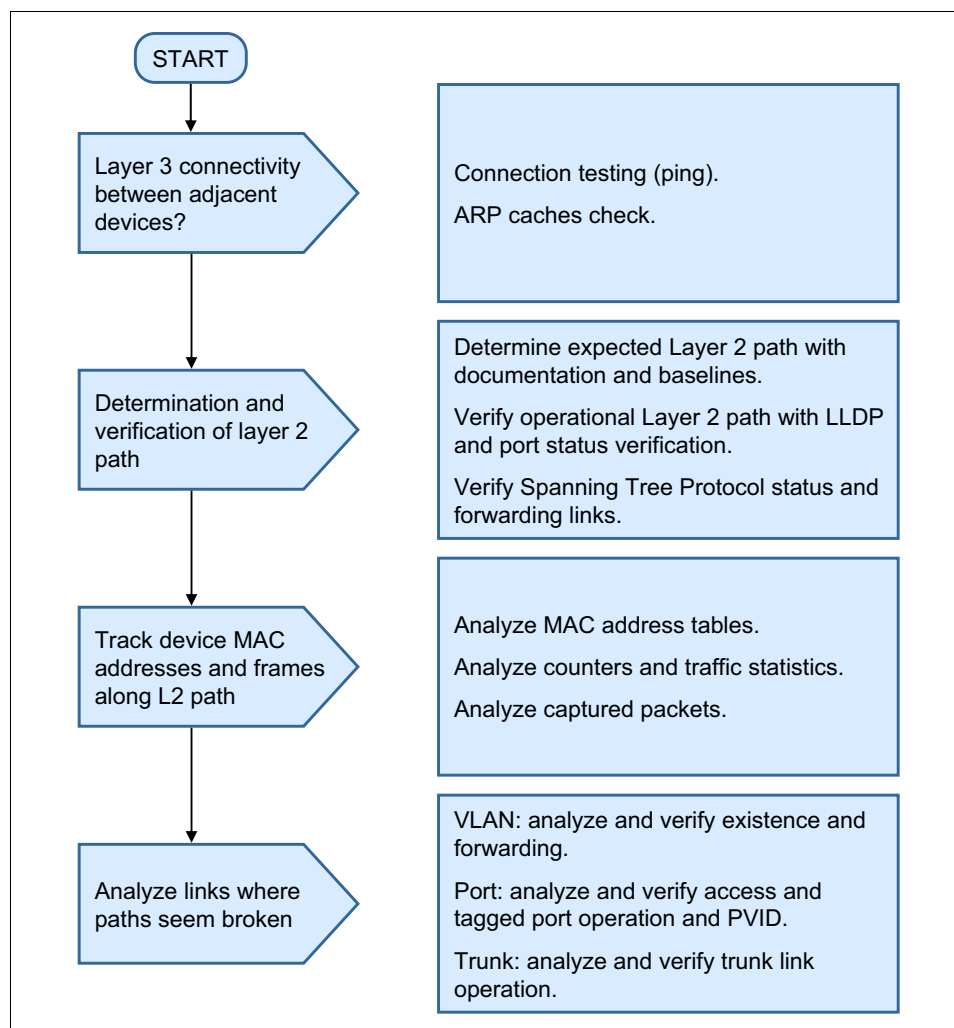


Figure A-1 Troubleshooting flowchart

Layer 2 troubleshooting commands

The following commands are listed according to the workflow that we described in the previous section.

Verify connectivity

Verify the connectivity by using the following ping:

```
ping 10.1.1.1
```

Verify the ARP cache

When you start a ping, the host needs to know the destination MAC address first so it can address the Ethernet frame properly. To determine the destination MAC address, the host sends an ARP request frame, which is responded to with an ARP reply. The ARP reply contains the destination IP and MAC address. This information is stored in the ARP cache, often for a few minutes.

If the ping failed and you can find the destination MAC address in the ARP cache, this result is a strong indication that your Layer 2 connectivity is working. You might experience problems with a firewall or other security measures on a device.

Use the following commands to display the ARP cache:

- ▶ On a Windows host: `arp -a`
- ▶ On the switch: `show ip arp`

Determination of Layer 2 path

You use the existing network documentation and compare the current network condition against it. If the documentation is missing, you document the current network situation by using the following command results as input:

- ▶ Use the following commands to verify which interfaces are up, duplex, speed, and so on:
 - IBM: `show interface link`
 - Cisco: `show interfaces status`
- ▶ Use the following commands to verify the mapping of ports and VLANs:
 - IBM: `show interface information`
 - Cisco: `show interface trunk`
- ▶ Use the following commands to verify the interconnection of switches and routers:
 - IBM: `show lldp remote-device`
 - Cisco: `show lldp neighbors`
- ▶ Use the following commands to verify the forwarding of traffic on links:
 - IBM and Cisco: `show spanning-tree`
 - IBM and Cisco: `show interface counters`
- ▶ Use the following commands to verify the LACP trunks:
 - IBM: `show portchannel information`
 - Cisco: `show etherchannel summary`

Tracking traffic along L2 path

After you know what your actual network looks like, you can track the flow of traffic across it. This tracking is best done by tracking MAC addresses. Every switch holds a table of MAC addresses. The table is built and updated with every new Ethernet frame that crosses the switch by putting the source MAC address and the switchport ID where the frame entered the switch into the MAC address table. This information is needed by the switch when an Ethernet frame is forwarded to the specific MAC address. Any destination MAC address can be mapped to a switchport.

If a frame is to be forwarded but there is no valid entry in the MAC address table, the frame is broadcasted on all ports, except the port where the frame entered the switch. There are instances in which this configuration makes sense to clear the table, initiate some traffic, and verify it again.

Reviewing this table shows you where the switch sees the device with that specific MAC address connected.

Use the following commands to show the current content of the table:

- ▶ IBM: `show mac-address-table`
- ▶ Cisco: `show mac address-table`

Use the following commands to clear the current content of the table:

- ▶ IBM: `clear mac-address-table`
- ▶ Cisco: `clear mac address-table`

Analyze links where path seems broken

When you find a path that seems to be broken, the following commands can help to analyze the root cause of the problem:

- ▶ Use the following commands to verify the existence and the correct forwarding of the VLANs:
 - IBM and Cisco: `show vlan`
 - IBM: `show interface information`
 - Cisco: `show interface switchport`
 - IBM and Cisco: `show spanning-tree`
- ▶ Use the following commands to verify the correct membership and tagging on the switch ports and interswitch links:
 - IBM: `show interface information`
 - Cisco: `show interface trunk`
 - Cisco: `show interface status`

Baseline documentation

Experience shows that documenting a network is a difficult task. Often there is too much or not enough information, or the information is not what you need.

To simplify the effort of creating and reading the documentation of a network, it might make sense to separate the documentation by OSI Layers 1, 2 and 3. Each of these layers is reflected by its own configuration in the network devices. You also can troubleshoot the layers individually. The following drawings shall show a simple network:

Figure A-2 on page 182 shows the cabling, devices, naming convention that is used, and ports of OSI Layer 1.




Figure A-2 OSI Layer 1

Figure A-3 shows the VLANs, ports, VLAN membership, tagging, and PVID of OSI Layer 2.

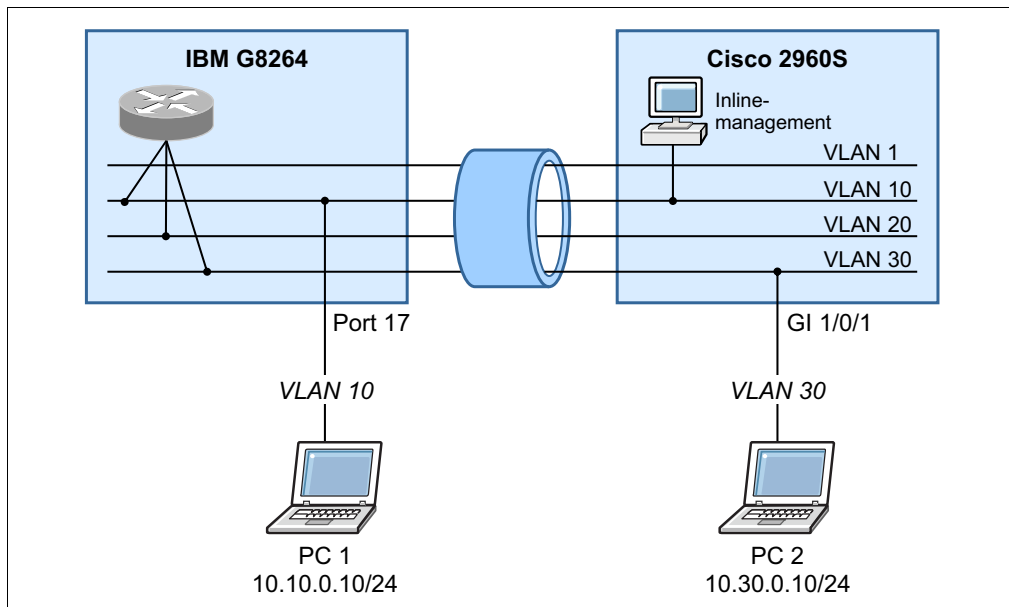


Figure A-3 OSI Layer 2

Figure A-4 shows the IP subnets, routes, and default gateway of OSI Layer 3.




Figure A-4 OSI Layer 3

Additional useful information for baseline documentation

The following useful information also is used in baseline documentation:

- ▶ Average and peak bandwidth for switch-to-switch links and switch-to-server links.
- ▶ Average rate of broadcasts and multicasts in the network.
- ▶ Software version that is used and the date of last firmware update.

Firmware update of IBM Flex System network switches

The Ethernet switch firmware can be updated by using one of the following methods:

- ▶ The use of a graphical user interface (GUI)
- ▶ Through Flex System Manager (FSM) by using the Update Manager
- ▶ The use of the Command-line Interface (CLI)

If there an FSM module is not installed, you can use one of the following ways to update the firmware of the integrated network switches.

For more information, see the IBM Flex System Information Center at this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.net.workdevices.doc/network_iomodule.html

Update the switch by using the web-based GUI

Complete the following steps to update the switch by using the browser-based GUI:

1. Go to the IBM Fix Central website: <http://ibm.com/support/fixcentral/options>
2. Select the choices as shown in Figure A-5 on page 184 and click **Continue**.

Fix Central

Fix Central provides fixes and updates for your system's software, hardware, and operating system.

For additional information, click on the following link.
[Getting started with Fix Central](#)

Select product

Find product

Select the product below.

When using the keyboard to navigate the page, use the **Alt** and **down arrow** keys to navigate the selection lists.

Product Group
 PureSystems

Select from PureSystems
 PureFlex System

Select from PureFlex System
 Chassis

Select from Chassis
 Enterprise Chassis

Select from Enterprise Chassis
 7893

Operating system
 Operating system independent / None

[Continue](#)

Figure A-5 Fix Central window

3. Select the products that you want to install and click **Continue**, as shown in Figure A-6.

<input type="checkbox"/>	IBM Flex System Fabric EN4093 10Gb Scalable Switch Firmware Update →	May
	ibm_fw_scsw_en4093-7.2.2.2_anyos_noarch	
	Change History	Readme
<input type="checkbox"/>	IBM Flex System EN4091 10Gb Ethernet Passthru →	May
	ibm_fw_scsw_en4091-1.0.6.0_anyos_noarch	
	Change History	Readme
<input checked="" type="checkbox"/>	IBM Flex EN2092 1Gb Ethernet Scalable Switch →	May
	ibm_fw_scsw_en2092-7.2.2.2_anyos_noarch	
	Change History	Readme
		Back
Continue	Clear selections	Back
		Show fix details Hide fix details

Figure A-6 Selecting fixes

4. Log in by using your IBM ID and select your preferred download, as shown in Figure A-7 on page 185.

Figure A-7 Download options

5. Accept the terms and conditions.
6. Download the Firmware package.
7. Check the readme file for updates of the update process.
8. Extract the boot and OS image files into a directory.

The compressed file that contains the following files and directories:

- Boot image: `ibm_fw_scsw_en2092-7.2.2.2_anyos_noarch_Boot.img`
- OS image: `ibm_fw_scsw_en2092-7.2.2.2_anyos_noarch_OS.img`
- A directory that contains the MIB files

9. Establish a connection between the Ethernet port of the Chassis Management Module (CMM) and the machine that is running the browser.

For more information about how to configure an IP address on a Switch module, see the CMM documentation.

10. Enter the IP address of the Switch and log in to the browser-based user interface (BBI) by using the following credentials:
 - Username (default): `admin` (or `USERID`)
 - Password (default): `admin` (or `PASSWORD`)

11. Click the **Configure** tab, as shown in Figure A-8 on page 186.

12. From the left-tree view, click **IBM Flex System EN2092 10 Gb Switch** → **System** → **Config/Image Control**.

13. Scroll down to the Image Settings group, as shown in Figure A-8 on page 186, and complete the following steps:

- a. In the **Image for Transfer** menu, select the wanted OS image bank.
- b. Click **Browse** and browse to your local file system to select the OS image file: `ibm_fw_scsw_en2092-7.x.x.x_anyos_noarch_OS.img`.
- c. Click **Download via Browser**.

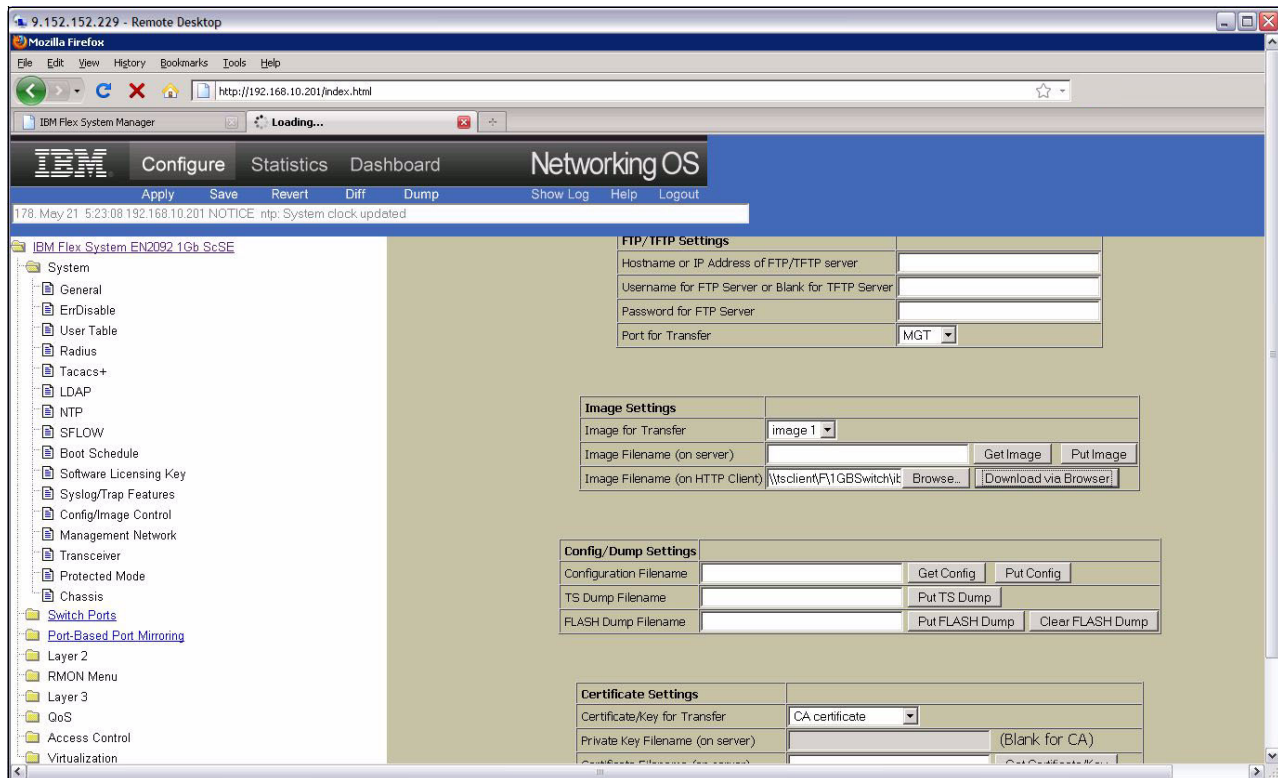


Figure A-8 Updating the firmware

The file transfer begins, followed by flashing non-volatile memory on the Switch. When the operation completes, the browser window returns and you see the following message at the bottom of the page:

Status of Previous Transfer ...

... Image downloaded via Browser ibm_fw_scsw_en2092-7.x.x.x_anyos_noarch_OS.img - Successful

***If you want to update both image banks, repeat step e above for the second image bank before updating the boot image below.

Do not reset: Do not reset or boot the switch between the OS and boot upgrades.

14. Repeat step 13 on page 185 and select the boot image from the menu and select the `ibm_fw_scsw_en2092-7.x.x.x_anyos_noarch_Boot.img` file.

The file transfer begins, followed by flashing non-volatile memory on the Switch. When the operation completes, the browser window returns and you see the following message at the bottom of the page:

Status of Previous Transfer ...

... Image downloaded via Browser `ibm_fw_scsw_en2092-7.x.x.x_anyos_noarch_Boot.img` - Successful

15. Set the **Next Boot Image Selection** to the image bank (1 or 2) that contains the new firmware, as shown in Figure A-9 on page 187.
16. Click **Submit** at the bottom of the page.
17. Click **REBOOT!** at the bottom of the page.
18. Wait for the switch to reboot.

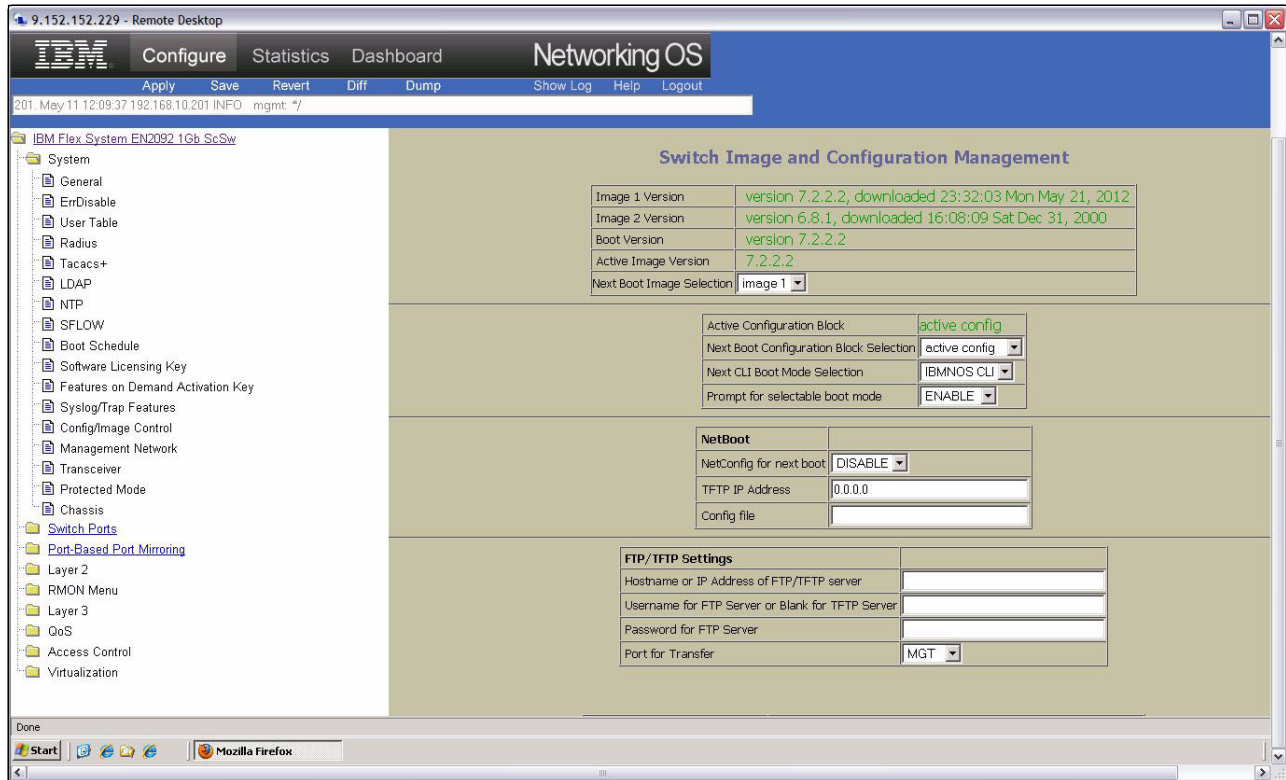


Figure A-9 Completing the firmware update

Using SSHv2 or Telnet

This method uses a Trivial File Transfer Protocol (TFTP) or File Transfer Protocol (FTP) server to update the switch firmware. Often, this server is installed on a machine that is reachable from the switch through the management module. However, when the switch is appropriately configured, the server can be attached to the external management port or an external or internal data port.

Important: Telnet is disabled by default. Unless you previously enabled telnet, use SSHv2.

Complete the following steps to use SSHv2 or Telnet:

1. Download the compressed VFSS software package file to the machine where the TFTP (or FTP) server is located.
2. Extract the boot and OS image files into a directory. Enable the server and set its default directory to the one in which the image files is located.
3. Establish a connection between the Ethernet port of the Management Module and the TFTP Server. For more information about configuring an IP address on a Switch module, see the CMM documentation.
4. Open a session by using the IP address of the Switch and log in to the VFSS Command Line Interface (CLI) by using the following credentials:
 - Username (default): admin (or USERID)
 - Password (default): admin (or PASSWORD)

5. Upgrade the OS image by entering the following command:

```
/boot/gtimg X TADDR Ibm_fw_scsw_en2092-7.x.x.x_anyos_noarch_OS.img
```

Where:

- X = 1 or 2 (determined by the image bank you want to use)
- TADDR = IP address of the TFTP Server

It is recommended that you retain the previous OS version by loading the upgrade into the other image block and then reset the switch by using the new image. Use the `/boot/image` command to select the preferred image. Leave the user name blank for TFTP (press Enter and answer “Y” to the confirmation question). Wait for the upgrade to complete successfully.

Important: Do not reset the switch between the OS and boot upgrades.

6. Upgrade the boot image by entering the following command:

```
/boot/gtimg boot TADDR Ibm_fw_scsw_en2092-7.x.x.x_anyos_noarch_Boot.img
```

Leave the user name blank for TFTP (press Enter and answer 'Y' to the confirmation question). Wait for the upgrade to complete successfully.

7. After the boot upgrade completes, reset the switch by using the following command:

```
/boot/reset
```

You must reset the switch to activate the new image.

When you reset the switch, it boots by using the selected image (1 or 2). Ensure that you are booting from the upgraded image by running the `/boot/cur` command.

A switch reset completes in approximately 60 seconds.

8. After rebooting, you can verify the firmware version by using the `show version` command, as shown in Example A-1 on page 189.

Example: A-1 Verifying the firmware version

Router>**show version**

System Information at 23:48:16 Mon May 21, 2012

Time zone: America/US/Pacific

Daylight Savings Time Status: Disabled

IBM Flex System EN2092 1Gb Ethernet Scalable Switch

Switch has been up for 0 days, 0 hours, 4 minutes and 9 seconds.

Last boot: 23:46:05 Mon May 21, 2012 (reset from Telnet/SSH)

MAC address: 08:17:f4:76:78:00 IP (If 1) address: 0.0.0.0

Management Port MAC Address: 08:17:f4:76:78:ef

Management Port IP Address (if 128): 192.168.10.201

Software Version 7.2.2.2 (FLASH image1), active configuration.

Hardware Part Number : 49Y4295
Hardware Revision : 00
Serial Number : Y050VT16E0AK
Manufacturing Date (WWYY) : 3711
PCBA Part Number : BAC-00079-00
PCBA Revision : 0
PCBA Number : 00
Board Revision : 00
PLD Firmware Version : 1.3

Temperature Warning : 36 C (Warn at 60 C/Recover at 55 C)
Temperature Shutdown : 36 C (Shutdown at 65 C/Recover at 60 C)
Temperature Inlet : 33 C
Temperature Exhaust : 36 C
Temperature Local : 35 C
Temperature Remote 1 : 54 C
Temperature Remote 2 : 42 C
Temperature Remote 3 : 42 C
Temperature Phy 0x01 : 54 C
Temperature Phy 0x09 : 45 C
Temperature Phy 0x11 : 45 C

Power Consumption : 37.980 W (12.408 V, 3.061 A)

Switch is in I/O Module Bay 1

Router>

Abbreviations and acronyms

ARP	Address Resolution Protocol	OUI	organizationally unique identifier
BBI	browser-based interface	PC	personal computer
BPDU	Bridge protocol data unit	PDU	power distribution unit
CDP	Cisco Discovery Protocol	PVRST	Per VLAN Rapid Spanning Tree
CLI	command-line interface	PVST	Per-VLAN Spanning Tree
CMM	Chassis Management Module	RMON	Remote Monitoring
DA	destination address	RSS	Receive-side scaling
DGW	default gateway	RSTP	Rapid Spanning Tree Protocol
DOCSIS	Data Over Cable Service Interface Specification	SA	source address
FDB	forwarding database	STP	Spanning Tree Protocol
FSM	Flex System Manager	TCA	Target Channel Adapter
FTP	File Transfer Protocol	TCN	Topology Change Notification
GE	Gigabit Ethernet	TFTP	Trivial File Transfer Protocol
GUI	graphical user interface	TTL	time to live
ICMP	Internet control message protocol	VLAG	Virtual Link Aggregation Groups
ID	identifier	VLAN	virtual LAN
IEEE	Institute of Electrical and Electronics Engineers		
IGMP	Internet Group Management Protocol		
IP	Internet Protocol		
ISCLI	industry standard command line interface		
ISL	Inter-Switch Link		
ITSO	International Technical Support Organization		
LACP	Link Aggregation Control Protocol		
LACPDU	LACP Data Units		
LAG	link aggregate group		
LAN	local area network		
LCAP	Link Aggregation Control Protocol		
LLDP	Link Layer Discovery Protocol		
MAC	media access control		
MEC	Multichassis Ether Channel		
MIB	management information base		
MLT	Master Latency Timer		
MST	Multiple Spanning Tree		
MSTP	Multiple Spanning Tree Protocol		
MTU	maximum transmission unit		
OS	operating system		
OSI	Open Systems Interconnect		

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 paper.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topics in this document. Note that some publications referenced in this list might be available in softcopy only:

- ▶ *Implementation of IBM j-type Ethernet Switches and Routers*, SG24-7882
- ▶ *IBM Flex System Networking in an Enterprise Data Center*, REDP-4834
- ▶ *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984
- ▶ *IBM Flex System EN2092 1Gb Ethernet Scalable Switch*, TIPS0861
- ▶ *IBM Flex System Fabric EN4093 10Gb Scalable Switch*, TIPS0864

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft, and additional materials at the following website:

<http://www.ibm.com/redbooks>

Other publications

The following publications are also relevant as further information sources:

- ▶ IBM RackSwitch G8264 Application Guide (6.8):
<http://ibm.com/support/docview.wss?uid=isg3T7000464>
- ▶ Virtual PortChannel Quick Configuration Guide:
http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/configuration_guide_c07-543563.html
- ▶ Cisco Nexus 5000 Series NX-OS Software Configuration Guide, Configuring Multiple Spanning Tree:
http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/configuration_guide/cli_rel_4_0_1a/MST.html

Online resources

The following websites are also relevant as further information sources:

- ▶ ProCurve & Cisco Spanning Tree Interoperability
<http://cdn.procurve.com/training/Manuals/ProCurve-and-Cisco-STP-Interoperability.pdf>
- ▶ Best Practice for configuring HP procurve with Cisco switch forum
<http://h30499.www3.hp.com/t5/Switches-Hubs-Modems-Legacy-ITRC/Best-Practice-for-configuring-HP-procurve-with-Cisco-switch/td-p/4701340>
- ▶ Radia Perlman, Intel Labs, Donald Eastlake, Huawei Technologies, Introduction to Trill, The Internet Protocol Journal, Volume 14, No. 3:
http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_14-3/143_trill.html

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



Deploying IBM Flex System into a Cisco Network



**Learn how to
integrate IBM Flex
System into your
network**

**See real life Layer 2
configurations with
Flex System switches**

**Find out how easy it is
to connect network
devices**

This IBM Redpaper publication provides information on how to integrate IBM Flex System into an existing customer network. It focuses on interoperability and seamless integration from the network perspective.

The paper describes the complete configuration of the most common scenarios. It guides you through several setups, and shows in detail how to configure the network switches, and how to verify the functionality and proper operation.

This paper can help you to easily configure and monitor your Layer 2 setup. Typical well established Layer 2 Network setups use combinations of Spanning Tree Protocol, VLANs and link aggregation.

Scenarios described in this paper includes the use of these switching products:

- ▶ Cisco Nexus 5000 (including vPC)
- ▶ Cisco Catalyst 6500
- ▶ IBM RackSwitch (including VLAG)
- ▶ IBM Flex System Ethernet Scalable Switch (including VLAG)

We describe the use of these switches with each of the following Spanning Tree Protocol (STP) configurations:

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

The paper is aimed at network administrators familiar with Cisco network products. It uses the industry standard command-line interface (isCLI) as management interface and we assume the reader is familiar with Cisco products and the use of isCLI.

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