



Interconnecting Data Centers Using VPLS

Ensure Business Continuance on Virtualized Networks
by Implementing Layer 2 Connectivity Across Layer 3

Interconnecting Data Centers Using VPLS

Nash Darukhanawalla, Patrice Bellagamba

Copyright© 2009 Cisco Systems, Inc.

Published by:
Cisco Press
800 East 96th Street
Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America

First Printing June 2009

Library of Congress Cataloging-in-Publication Data:

Darukhanawalla, Nash, 1963-

Interconnecting data centers using VPLS / Nash Darukhanawalla and Patrice Bellagamba.

p. cm.

ISBN-13: 978-1-58705-992-6 (pbk.)

ISBN-10: 1-58705-992-4 (pbk.)

1. Extranets (Computer networks) 2. Internetworking (Telecommunication) 3. Electronic data processing departments. 4. Cisco IOS. I. Bellagamba, Patrice, 1959- II. Title.

TK5105.875.E87D37 2010

004.6—dc22

2009020207

ISBN-13: 978-1-58705-992-6

ISBN-10: 1-58705-992-4

Warning and Disclaimer

This book is designed to provide information about interconnecting data centers using Virtual Private LAN Service (VPLS). Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc., shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Corporate and Government Sales

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact: **U.S. Corporate and Government Sales** 1-800-382-3419 corpsales@pearsontechgroup.com

For sales outside the United States please contact: **International Sales** international@pearsoned.com

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher: Paul Boger

Associate Publisher: Dave Dusthimer

Executive Editor: Mary Beth Ray

Managing Editor: Patrick Kanouse

Senior Development Editor: Christopher Cleveland

Project Editor: Jennifer Gallant

Editorial Assistant: Vanessa Evans

Book Designer: Louisa Adair

Composition: Mark Shirar

Cisco Representative: Eric Ullanderson

Cisco Press Program Manager: Anand Sundaram

Copy Editor: Keith Cline

Technical Editors: Cesar Carballes and Yves Louis

Indexer: Brad Herriman

Proofreader: Apostrophe Editing Services



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aronnet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, Gigamon, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanel, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0812R)

Introduction

This book presents Virtual Private LAN Service (VPLS)-based solutions that provide a high-speed, low-latency network and Spanning Tree Protocol (STP) isolation between data centers. The book also includes detailed information about issues that relate to large Layer 2 bridging domains and offers guidance for extending VLANs over Layer 3 networks using VPLS technology.

The solutions presented in this book have been validated under the Cisco Validated Design System Assurance program. All solutions were validated with a wide range of system tests, including system integration, fault and error handling, and redundancy. Testing also verified the end-to-end flow of both unicast and multicast unidirectional traffic. Voice, using components of the Cisco Unified Communications solution, was also implemented and verified.

The solutions in this book were developed because globalization, security, and disaster recovery considerations are driving divergence of business locations across multiple regions. In addition, organizations are looking to distribute workload between computers, share network resources effectively, and increase the availability of applications. With the ultimate goal of eliminating all downtime and sharing data across regions, enterprises are deploying geographically dispersed data centers to minimize planned or unplanned downtime, whether it is caused by a device failure, security breach, or natural disaster.

As data centers grow in size and complexity, enterprises are adopting server virtualization technologies to achieve increased efficiency and use of resources. In addition to providing resource optimization, virtualization strategies offer data protection, which enables enterprises to build disaster recovery solutions and provide high availability, scalability, flexibility, and business continuity.

Server virtualization technologies include the following:

- **VMotion:** Allows an individual virtual machine (such as Windows Server) to be dynamically moved to another VMware server. A dedicated VLAN is required for VMotion traffic so that virtual machines can be moved without affecting users. In addition, the group of servers that VMs are balanced between must be in the same Layer 2 domain, because attributes such as an IP address cannot change when a virtual machine moves. Therefore, all VMware servers, including the source VMware server, must have connections to the same VLAN.
- **NIC teaming:** Servers that include only one network interface card (NIC) are susceptible to many single points of failure, such as a failure of the NIC, its network cable, or the access switch to which it connects. A solution developed by NIC vendors, NIC teaming eliminates this single point of failure. In this solution, special drivers allow two NICs to be connected to separate access switches or to separate line cards on the same access switch. If one NIC fails, the other NIC assumes the IP address of the server and takes over operation without disruption. Types of NIC teaming solutions include active/standby and active/active. All solutions require Layer 2 adjacency for the teamed NICs.

- **Server clusters:** High-availability (HA) server clusters have become key components of IT strategies as organizations need to increase processing power, distribute workloads between computers, share network resources effectively, and increase the availability of applications. HA clusters typically are built with two separate networks: a public network to access the active node of the cluster from outside the data center, and a private network to interconnect the nodes of the cluster for private communications. The private network connection is also used to monitor the health and status of each node in the HA cluster using the heartbeat system. This solution requires that the network is capable of handling any kind of failure without causing a split-brain condition that could lead to duplicate instances of services and even the corruption of data on shared storage devices. The private network is a nonrouted network that shares the same Layer 2 VLAN between the nodes of the cluster even when extended between multiple sites.

These virtualization technologies have resulted in an expansion of Layer 2 domains, which in turn has increased the spanning-tree domain at the network level. STP was developed to handle a network with a small diameter, so an enterprise network with geographically dispersed data centers needs an effective solution for Layer 2 connectivity between multiple sites.

Also, during the process of migrating physical servers from one location to another, it is much easier to extend the Layer 2 VLAN and maintain the original configuration of the systems, thus avoiding IP address renumbering. Even during a phased migration period, when just part of the server farm is being relocated, the Layer 2 adjacency is often required across the entire server farm to ensure business continuity.

As data center resources and security requirements continue to grow, organizations must connect multiple data centers over larger distances. As a result, organizations are facing additional challenges such as maintaining the high availability of applications and dealing with complex multisite interconnections.

Objective of This Book

This book provides design guidance, configuration examples, and best practices for deploying a single IP/MPLS-based network to interconnect data centers ensuring high availability Layer 2 connectivity with STP isolation in the core. Customers who have already deployed a separate optical network for Layer 2 extension can also take advantage of these solutions to reduce infrastructure and maintenance costs.

This book addresses issues that are related to large Layer 2 bridging domains and provides guidance for extending VLANs using VPLS technology.

This book also examines in detail the technologies that offer such solutions, explains the benefits and trade-offs of various solutions, and describes a variety of deployment options. The deployment model that an organization chooses depends on the complexity of the requirements, the protocols currently deployed in the data centers, the scalability required, and many other factors.

Who Should Read This Book

This book is intended for systems professionals and system engineers who are designing solutions for interconnecting data centers that ensure high availability Layer 2 connectivity and STP isolation. Service providers that offer metro Ethernet leased-line aggregation and Layer 2 transport services can also benefit from these solutions that provide large-scale Layer 2 extension.

Cisco Validated Design Program

The Cisco Validated Design (CVD) program designs, tests, and documents systems and solutions to facilitate faster, more reliable, and more predictable customer deployments. The program includes Cisco Validated Design and CVD System Assurance.

Cisco Validated Design

Cisco Validated Designs are systems or solutions that have been validated through architectural review and proof-of-concept testing in a Cisco lab. Cisco Validated Design provides guidance for deploying new technologies or for applying enhancements to existing infrastructure.

CVD System Assurance

The Cisco Validated Design System Assurance is a program that identifies systems that have undergone architectural- and customer-relevant testing.

A CVD certified design is a highly validated and customized solution that meets the following criteria:

- Reviewed and updated for general deployment
- Achieves the highest levels of consistency and coverage within the Cisco Validated Design program
- Solution requirements successfully tested and documented with evidence to function as detailed within a specific design in a scaled, customer representative environment
- Zero observable operation impacting defects within the given test parameters; that is, no defects that have not been resolved either outright or through software change, redesign, or workaround

For more information about Cisco CVD program, refer to <http://tinyurl.com/6gxuk2>.

How This Book Is Organized

The material in this book is presented in a building-block fashion that takes you from the legacy deployment models for data center interconnect (DCI) and problems associated with extending Layer 2 networks, through VPN technologies, to various Multiple Spanning Tree (MST)-, Embedded Event Manager (EEM)- and generic routing encapsulation (GRE)-based deployment models, and beyond. Although this book is intended to be read cover to cover, it is designed to be flexible and allow you to easily find information that applies to your needs.

The chapters cover the following topics:

- **Chapter 1, “Data Center Layer 2 Interconnect”:** This chapter provides an overview of high availability clusters. It also explains DCI legacy deployment models and problems associated with extending Layer 2 networks.
- **Chapter 2, “Appraising Virtual Private LAN Service”:** This chapter discusses Layer 2 and Layer 3 VPN technologies and provides introductions to VPLS, pseudowires, EEM, and MPLS.
- **Chapter 3, “High Availability for Extended Layer 2 Networks”:** This chapter focuses on design components such as maximum transmission unit (MTU), core routing protocols, and convergence optimization techniques to achieve high availability.
- **Chapter 4, “MPLS Traffic Engineering”:** This chapter explains the implementation of MPLS-TE for load repartition of Layer 2 VPN traffic over parallel links. It also introduces Fast Reroute (FRR) for faster convergence.
- **Chapter 5, “Data Center Interconnect: Architecture Alternatives”:** This chapter highlights several options for implementing DCI. In addition, this chapter provides guidance for selecting an appropriate solution based on your requirements (such as scalability and ease of implementation).
- **Chapter 6, “Case Studies for Data Center Interconnect”:** This chapter provides case studies that relate to the DCI solutions that this book describes.
- **Chapter 7, “Data Center Multilayer Infrastructure Design”:** This chapter highlights the Cisco data center multitier model and provides information about network topology, hardware, software, and traffic profiles used for validating the designs in this book.
- **Chapter 8, “MST-Based Deployment Models”:** This chapter covers “MST in pseudowire” and “isolated MST in N-PE” solutions and provides configuration details for implementing these solutions.

- **Chapter 9, “EEM-Based Deployment Models”:** This chapter explains “EEM semaphore protocol,” which was developed to achieve N-PE redundancy in the absence of ICCP. In addition, this chapter describes various EEM-based VPLS and Hierarchical VPLS (H-VPLS) solutions, provides in-depth theory about the operation of each solution, and provides configuration details.
- **Chapter 10, “GRE-Based Deployment Models”:** This chapter provides VPLS and H-VPLS solutions over an IP network using VPLSoGRE (VPLS over GRE).
- **Chapter 11, “Additional Data Center Interconnect Design Considerations”:** This chapter introduces other technologies or issues that should be considered when designing DCI solutions.
- **Chapter 12, “VPLS PE Redundancy Using Inter-Chassis Communication Protocol”:** This chapter introduces ICCP protocols and provides various redundancy mechanisms and sample configurations.
- **Chapter 13, “Evolution of Data Center Interconnect”:** This chapter provides a brief overview of the emerging technologies and the future of DCI.
- **Glossary:** This element provides definitions for some commonly used terms associated with DCI and the various deployment models discussed in the book.

The authors have also written several documents, including *Interconnecting Geographically Dispersed Data Centers*, *High Availability Clusters*, *Layer 2 Extension Between Remote Data Centers*. These documents cover a few key concepts from this book and are freely available on Cisco.com.

Data Center Layer 2 Interconnect

Many enterprises are making fundamental changes to their business processes by using advanced IT applications to achieve enhanced productivity and operational efficiencies. As a result, the underlying network architecture to support these applications is evolving to better accommodate this new model.

As data availability becomes a critical requirement, many businesses are devoting more resources to ensure continuous operation. Enterprises are provisioning dedicated networks to guarantee performance metrics for applications without compromising security.

Although maintaining uninterrupted access to all data center applications is desirable, the economics of business-continuance require network operators to prioritize applications according to their importance to the business. As a result, data centers need a range of business-continuance solutions to accommodate this goal, from simple tape backup and remote replication to synchronous mirroring and mirrored distributed data centers.

Enterprises can enhance application resilience in several ways, including the following:

- Removing single points of server failure by deploying high-availability clusters or load-balancing technology across web and application servers
- Extending the deployment of these clusters in different data centers to protect against major disruptions

User access is as important as downtime protection and data recovery. Following a disruption, how long can the business afford for users to be without access to applications? Companies are employing technologies such as Global Site Selector that enable users to manually or automatically connect to an alternative site running the application they need.

Businesses run tens and often hundreds of applications, each of which might have differing continuance requirements, measured in a time-to-recovery and data-loss perspective. IT groups need to match the associated characteristics and cost of a business-continuance

solution to the potential business and consider which technologies to deploy where problems impact data, applications, and user access.

Cisco delivers scalable, secure, and cost-effective technology that helps enterprises build end-to-end backup and recovery solutions and disaster recovery solutions. These solutions include the following:

- High-availability data center networking and storage-area networks for nonstop access to applications and data
- Synchronized distributed data centers for continuous service over WANs in the event of site disruptions
- Synchronous disk mirroring and replication over WANs for fast recovery and zero data loss
- Asynchronous data replication over IP networks for remote data protection
- Consolidated backup to tape or near-line disk and remote electronic vaulting over enterprise-wide storage networks for consistent protection of distributed data

Each of these solutions requires the appropriate network infrastructure to help ensure that user-specific availability, performance, distance, and latency requirements are met. In addition, enterprises require a resilient, integrated business-continuance network infrastructure to protect three key areas in the event of a disruption:

- Data
- Applications
- User access

Overview of High-Availability Clusters

High-availability (HA) clusters operate by using redundant computers or nodes that provide services when system components fail. Normally, if a server with a particular application crashes, the application is unavailable until the problem is resolved. HA clustering remedies this situation by detecting hardware/software faults, and immediately providing access to the application on another system without requiring administrative intervention. This process is known as *failover*.

HA clusters are often used for key databases, file sharing on a network, business applications, and customer services such as e-commerce websites. HA cluster implementations attempt to build redundancy into a cluster to eliminate single points of failure. These implementations include multiple network connections and data storage that connects via storage-area networks (SAN).

HA clusters usually are built with two separate networks:

- **The public network:** Used to access the active node of the cluster from outside the data center

- **The private network:** Used to interconnect the nodes of the cluster for private communications within the data center and to monitor the health and status of each node in the cluster

Public Network Attachment

For the public network (facing the nodes cluster), the server often is enabled by a dual-homing mechanism with one network interface card (NIC) configured in active state and one NIC configured in standby state. If a link to the active NIC fails, or the NIC loses connectivity with its default gateway, the operating system performs a failover. A NIC failover for a public network has no affect on cluster availability because the heartbeat mechanism and NICs in active/standby mode for public access are two separate hand-check mechanisms.

The network design must provide the highest availability for the LAN infrastructure. To achieve this goal, the teaming service or dual-homing should be distributed between different access switches, which in turn should be connected to different aggregation switches, as illustrated in Figure 1-1.

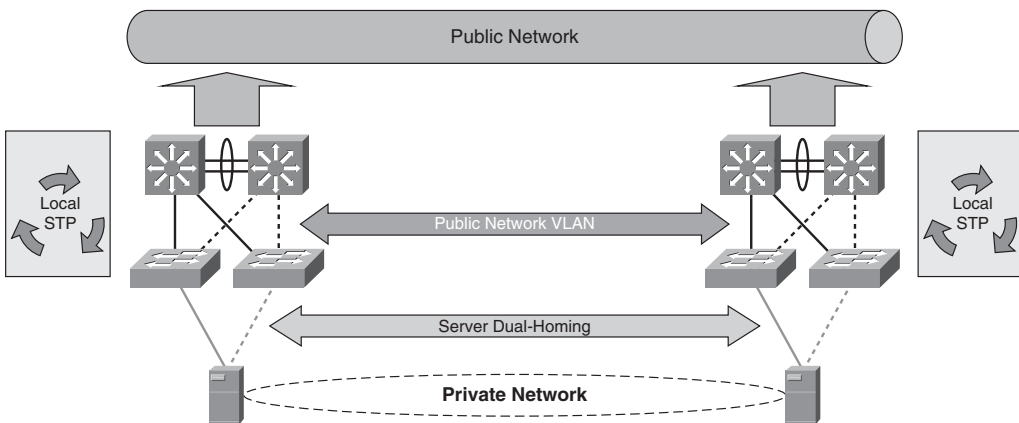


Figure 1-1 *Extending the public network.*

Private Network Attachment

The private network primarily carries cluster heartbeat, or keepalive, messages. Other server-to-server communications that occur on this private network include the following:

- Cluster data
- Cluster file system data
- Application data (back-end)

The private network is a nonrouted network that shares the same Layer 2 VLAN between the nodes of the cluster even when extended across multiple sites. In a campus cluster environment, heartbeats are sent via the private network from node to node of the HA cluster using a proprietary Layer 2 (nonroutable) protocol. The servers manage the I/O by sending traffic over all interfaces and by preventing traffic from being sent over a failing path. This approach provides resiliency in the event of a NIC failure on a server.

The heartbeat is the most important component of the cluster that uses the private network interconnection. However, if all paths go down for more than 10 seconds (applicable for most HA clusters), a *split-brain* situation can occur, which prompts the cluster framework to check the number of votes and decide which server or servers will continue as the members in the cluster. Nodes that lose cluster membership assume that they are isolated, and any applications that run on those nodes terminate. Surviving nodes know that the nonsurviving nodes have stopped, and the cluster will then restart the applications.

Although some HA cluster vendors recommend disabling Spanning Tree Protocol (STP) for the private interconnect infrastructure, such a drastic measure is neither necessary nor recommended when using Cisco Catalyst switches. In fact, Cisco has since provided the *PortFast* feature, which puts an access port into forwarding mode immediately after link up without losing loop-detection capabilities. To avoid connectivity delays, PortFast must be enabled on all access interfaces connecting cluster nodes. This rule also applies to any servers connected to the switch. The IEEE also defines the PortFast concept within the Rapid STP 802.1w standard under the *edge port* designation. In addition, Cisco supports Per-VLAN Spanning Tree, which maintains a spanning-tree instance for each VLAN configured in the network.

Note For detailed information about HA clusters, refer to the Windows HPC Server 2008 site at <http://www.microsoft.com/HPC/>.

For detailed information about STP PortFast configuration to resolve server/workstation startup connectivity delay, refer to the Cisco document Using PortFast and Other Commands to Fix Workstation Startup Connectivity Delays, available at <http://tinyurl.com/2e29bw>.

For detailed information about designing a network for extended HA clusters, refer to the following Cisco white paper A: “Technology and Networking Guide for High Availability Clusters Extended Across Multiple Data Centers,” at <http://tinyurl.com/cb4f3k>.

Data Center Interconnect: Legacy Deployment Models

Several transport technologies are available for interconnecting the data centers, each of which provides various features and allows different distances, depending on factors such as the power budget of the optics, the lambda used for transmission, the type of fiber, and so forth.

Consider the features of the LAN and SAN switches that provide higher availability for the data center interconnect (DCI) before considering some of the available technologies. The convergence time required for the application also is important and should be evaluated.

The list that follows describes common transport options:

- **Dark fiber:** Dark fiber is a viable method for extending VLANs over data center or campus distances. The maximum attainable distance is a function of the optical characteristics (transmit power and receive sensitivity) of the LED or laser that resides in a small form-factor pluggable (SFP) or Gigabit Interface Converter (GBIC) transponder, combined with the number of fiber joins, and the attenuation of the fiber.
- **Coarse wavelength-division multiplexing (CWDM):** CWDM offers a simple solution to carry up to eight channels (1 Gbps or 2 Gbps) on the same fiber. These channels can carry Ethernet or Fiber Channel. CWDM does not offer protected lambdas, but client protection allows rerouting of the traffic on the functioning links when a failure occurs. CWDM lambdas can be added and dropped, allowing the creation of hub-and-spoke, ring, and meshed topologies. The maximum achievable distance is approximately 60 miles (100 km) with a point-to-point physical topology and approximately 25 miles (40 km) with a physical ring topology.
- **Dense wavelength-division multiplexing (DWDM):** DWDM enables up to 32 channels (lambdas), each of which can operate at up to 10 Gbps. DWDM networks can be designed either as multiplexing networks that are similar to CWDM or with a variety of protection schemes to guard against failures in the fiber plant. DWDM also offers more protection mechanisms (splitter protection and Y-cable protection), and the possibility to amplify the channels to reach greater distances.

Note For details about data center transport technologies, refer to Chapter 2 of *Data Center High Availability Clusters Design Guide*, available at <http://tinyurl.com/ct4cw8>.

In nearly all of these deployment models, costs associated with deploying and maintaining a dedicated optical network is one of the biggest concerns. Also, there is no STP isolation. Depending on the nature of the problem, issues in one data center will affect other data centers. Another disadvantage is the lack of load balancing across redundant paths due to blocked links in the core network.

Problems Associated with Extended Layer 2 Networks

A common practice is to add redundancy when interconnecting data centers to avoid split-subnet scenarios and interruption of the communication between servers, as illustrated in Figure 1-2. The split-subnet is not necessarily a problem if the routing metric makes one site preferred over the other. Also, if the servers at each site are part of a cluster and the communication is lost, mechanisms such as the quorum disk avoid a split-brain condition.

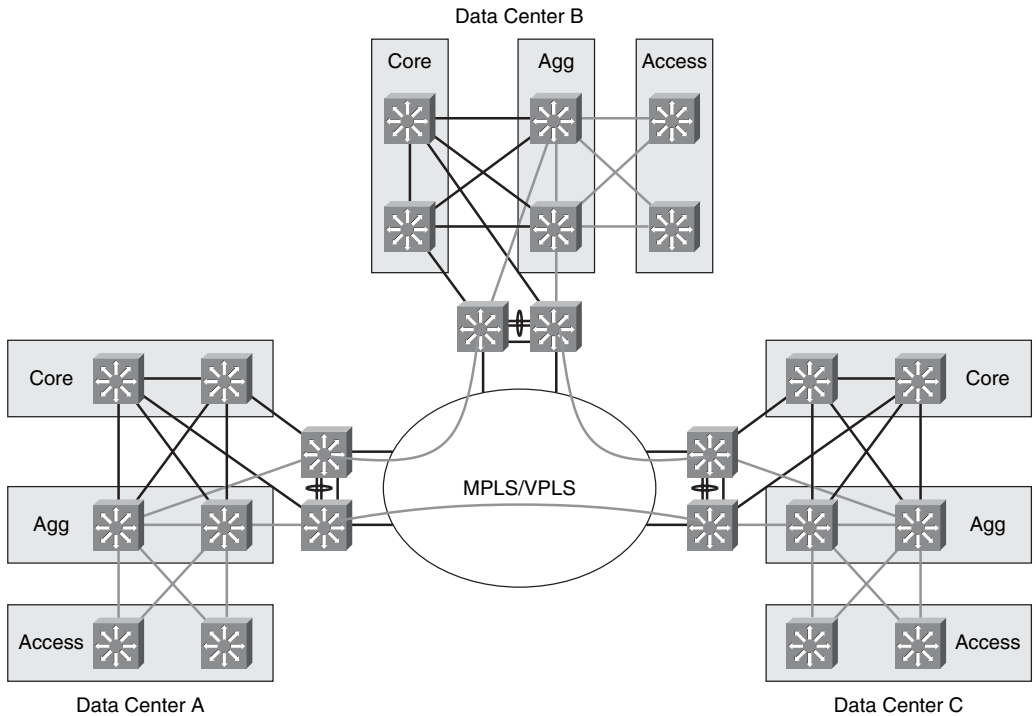


Figure 1-2 *Layout of multiple data center interconnect with redundant N-PEs in each data center.*

Adding redundancy to an extended Ethernet network typically means relying on STP to keep the topology loop free. STP domains should be reduced as much as possible and limited within the data center. Cisco does not recommend deploying the legacy 802.1d because of its old timer-based mechanisms that make the recovery time too slow for most applications, including typical clustering software.

An extended Layer 2 network does introduce some problems to contend with, however.

STP operates at Layer 2 of the Open Systems Interconnection (OSI) model, and the primary function STP is to prevent loops that redundant links create in bridge networks. By exchanging bridge protocol data units (BPDU) between bridges, STP elects the ports that eventually forward or block traffic.

The conservative default values for the STP timers impose a maximum network diameter of seven. Therefore, two bridges cannot be more than seven hops away from each other.

When a BPDU propagates from the root bridge toward the leaves of the tree, the age field increments each time the BPDU goes through a bridge. Eventually, the bridge discards the BPDU when the age field goes beyond maximum age. Therefore, convergence of the spanning tree is affected if the root bridge is too far away from some bridges in the network.

An aggressive value for the max-age parameter and the forward delay can lead to an unstable STP topology. In such cases, the loss of some BPDUs can cause a loop to appear. Take special care if you plan to change STP timers from the default value to achieve faster STP convergence.

Unlike legacy STP, Rapid STP (RSTP) converges faster because it does not depend on the timers to make a rapid transition. However, STP does not provide the required robustness for large-scale Layer 2 deployments:

- Network stability is compromised as a result of slow response to network failures (slow convergence). Even new spanning-tree developments such as RSTP and Multiple Spanning Tree (MST) assume good-quality physical connections such as dark fiber or WDM connections. These STP protocols are not built to accommodate frequent link-flapping conditions, high error rates, unidirectional failures, or nonreport of loss of signal. These typical and frequent behaviors of long- and medium-distance links could lead to STP slow convergence or even instability.
- The primary reason for multisite data centers is disaster recovery. However, because data centers typically require Layer 2 connectivity, failure in one data center can affect other data centers, which could lead to a blackout of all data centers at the same time.
- A broadcast storm propagates to every data center, which, if uncontrolled, could result in network-wide outage.
- STP blocks links, which prevents load balancing of traffic across redundant paths in the core network.

Note For understanding and tuning STP timers and the rules to tune them when absolutely necessary, refer to the Cisco document, “Understanding and Tuning Spanning Tree Protocol Timers,” available at <http://tinyurl.com/7ppqq>.

Summary

This chapter provided an overview of HA clusters, legacy deployment models for interconnecting data centers, and problems related to extending Layer 2 networks. The solutions that this book presents address these issues in more detail and provide guidance for designing and deploying DCI.

Index

Symbols

1RU (one rack unit) servers, 71
802.1Q in 802.1Q, 139

A

AC (attachment circuit) redundancy,
321

interaction, 324

access, VPLS

N-PEs (Network Provider Edges)
using ICCP (Inter-Chassis), 56-57

N-PEs (Network Provider Edges)
using MST (Multiple Spanning
Tree), 56

access layer, data center multilayer
topology, 71

active/standby VPLS node synchroni-
zation, 139-140

Adaptive Security Appliance (ASA),
QinQ bridge domain, 308

Adding FRR to Protect the Primary
Tunnel example (4-8), 52

affinity path (MPLS-TE), 43

affinity paths (MPLS-TE)

creating, 52

FRR (Fast Reroute), adding, 52

aggregation layer, data center
multilayer topology, 69

aggregation switches, 55

algorithms

Cisco Express Forwarding (CEF)-
based hashing algorithms, 69

Dijkstra shortest path first (SPF), 41

aligned design, 306

alternate route computation, routing
designs, 40-42

Any Transport over MPLS over GRE
(AToMoGRE), 277

applications

business applications, HA clusters, 2
resilience, enhancing, 1

architectures, 12

data center Layer 2 interconnect, 55

ASA (Adaptive Security Appliance),
QinQ bridge domain, 308

asymmetric flows, MAC address
flushing, 109

asynchronous data replication over IP networks, 2

asynchronous mode, BFD (Bidirectional Forwarding Detection), 36

attachment circuit (AC) redundancy, 321

attachment circuits (VPLS), 12

autodiscovery, PEs, 12

B

backbones, 29

backup N-PEs, 142

backup tunnels, presetting, 51

backups, consolidated backup, 2

balancing VFI traffic, 48

Balancing VFI Traffic listing (4-3), 48

Best Practices for Catalyst 6500/6000 Series and Catalyst 4500/4000 Series Switches, 310

BFD (Bidirectional Forwarding Detection), 34-35, 63

asynchronous mode, 36

neighbor relationships, establishing, 36

OSPF neighbor relationships, tearing down, 37

BFD (Bidirectional Forwarding Detection), 34-37, 63

ICCP (Inter-Chassis Communication), 323

blade servers, integral switches, 71

BPDUs (bridge protocol data units), 6-7

MST (Multiple Spanning Tree), 78

bridged networks, loops, 12

bridges

BPDUs (bridge protocol data units), 6-7

learning bridge mechanisms, 11

bridging

L2 bridging

DCI (*data center interconnect*), 329-331

Ethernet L2 bridging, 331

improving with STP (*Spanning Tree*), 330-331

IP (*Internet Protocol*), 332-333

MPLS (*Multiprotocol Label Switching*), 332-333

new concepts, 331-332

Rbridging, 332

broadcast packet storms, controlling, 312-313

business applications, HA clusters, 2

business-continuance solutions, 1-2

C

C-DA (customer destination address), 307

C-SA (customer source address), 307

C-TAG (customer VLAN tag), 307

calculation summary, MTU (maximum transmission unit), 23

campus cluster environments, heartbeats, 4

carrier Ethernet, 330

carrier-delay configuration, displaying, 38

Carrier-Delay Timer Configuration example (3-5), 39

carrier-delay timers, 38

configuration, 38-39

- Ethernet interfaces, 38-40
- POS interfaces, 40
- case studies, DCI (data center interconnect)
 - GOV, 61-65
 - OUT, 61, 65-68
- CEF (Cisco Express Forwarding)-based hashing algorithms, 69
- Cisco Catalyst 6500 switches
 - “LINEPROTO-UPDOWN” message, 40
 - QinQ, 291
 - VPLSoGRE, 278
- Cisco Catalyst switches
 - Cisco Catalyst 6500 switches
 - “LINEPROTO-UPDOWN” message, 40
 - QinQ, 291
 - VPLSoGRE, 278
 - PortFast, 4
 - STP (Spanning Tree Protocol), 4
- Cisco Data Center Infrastructure 2.5 Design Guide, 71**
- Cisco IOS Embedded Event Manager Data Sheet, 19**
- Cisco IOS MPLS Virtual Private LAN Service: Application Note, 14**
- classes, PWs (pseudowires), 15
- cluster server tests
 - isolated MST in N-PE deployment model, 134
 - MST in N-PE deployment model, 103
 - VPLS with N-PE redundancy using EEM semaphore protocol, 172
 - VPLSoGRE Multidomain with H-VPLS, 298
 - VPLSoGRE with N-PE redundancy using EEM semaphore, 286, 291

clusters

- HA (high-availability) clusters, 2-4, 7
 - business applications, 2*
 - failover, 2-3*
 - heartbeats, 4*
 - key databases, 2*
 - network file sharing, 2*
 - PortFast, 4*
 - private networks, 3-4*
 - public networks, 2-3*
 - split-brain situations, 4*
 - STP (Spanning Tree Protocol), 4*
- heartbeats, 4
- high-availability clusters, 1
- split-brain situations, 4

commands

- mpls ldp neighbor targeted, 20
- mpls ldp tcp pak-priority, 21
- no mpls ldp advertise-labels, 21
- preferred-path, 48
- process-max-time 50, 18
- show interface, 39
- show interface dampening, 35
- show mpls l2transport vc, 163, 190
- show running-config | begin event manager, 166, 192, 248

components, VPLS, 13-14

configuration

- carrier-delay timers, 38-39
- IP event dampening, 34-35
- MTU (Maximum Transmission Unit), 24
- multidomain H-VPLS with MEC and VLAN load balancing EEM, 265-267
- N-PE routers, distribute lists, 27

- N-PE1, 280
- STP mode MST*, 81-82
- OSPF (Open Shortest Path First) Protocol, N-PE routers, 93, 95
- PE (Provider Edge) redundancy, ICCP (Inter-Chassis Communication), 326-327
- VPLSoGRE, 279-282
- Configuring a Distribute List on an N-PE Router example (3-1)**, 27
- Configuring Explicit-Path Sequences example (4-4)**, 49
- Confirming Carrier-Delay Status on an Interface example (3-6)**, 39
- Congestion Notification 802.1Qau initiative**, 331
- connectivity, convergence, IGP, 32-33
- consolidated backup, 2
- constraint-based routing, 30
- content switching, 70
- control planes
 - H-VPLS with N-PE redundancy using EEM semaphore protocol, 179
 - multidomain H-VPLS with MEC and VLAN load balancing, 233
 - VPLS with N-PE redundancy using EEM semaphore protocol, 151
- Converged Enhanced Ethernet initiative**, 331
- convergence, IGP, 32-33
- convergence optimization, extended Layer 2 networks, HA (high availability), 32-42
- convergence tests
 - data center multilayer topologies, 73-76
 - traffic flow*, 73-76
 - traffic profile*, 74-76
 - traffic rate*, 73
 - H-VPLS with N-PE redundancy using EEM semaphore protocol, 195, 199, 217
 - isolated MST in N-PE deployment model, 130
 - MST in N-PE deployment model, 100
 - multidomain H-VPLS with MEC and VLAN load balancing, 252, 270
 - VPLS with N-PE redundancy using EEM semaphore protocol, 168
 - VPLSoGRE Multidomain with H-VPLS, 296-298
 - VPLSoGRE with N-PE redundancy using EEM semaphore, 284-286
- core layer, data center multilayer topology, 69
- core routing, extended Layer 2 networks, HA (high availability), 25-31
- customer destination address (C-DA), 307
- customer source address (C-SA), 307
- customer VLAN tag (C-TAG), 307
- CWDM (coarse wavelength-division multiplexing), DCI (data center interconnect), 5

D

- dark fiber, DCI (data center interconnect), 5
- data availability, 1
- Data Center Ethernet (DCE) initiative, 331
- Data Center High Availability Clusters Design Guide*, 5, 306
- data center interconnect (DCI), 4, 61
- data center Layer 2 interconnect, 55
- data centers, 1-2

- DCI (data center interconnect), 4-7, 61, 72, 329
 - Case studies*, 61-68
 - CWDM (*coarse wavelength-division*), 5
 - dark fiber*, 5
 - design alternatives*, 57-59
 - DWDM (*dense wavelength-division*), 5
 - HSRP (*Hot Standby Router Protocol*), *aligned design*, 306
 - importance*, 329
 - interdependence, removing*, 55-57
 - IP routing*, 329-330
 - L2 bridging*, 329-333
 - multicast deployment, Layer 2 environment*, 303-305
 - OSPF (*Open Shortest Path First*) *cost*, 318
 - QinQ *MAC address overlapping*, 307-310
 - QoS (*quality of service*), 315-317
 - redundancy*, 5-6
 - router IDs, selecting*, 319
 - routing design*, 306-307
 - services modules, aligning*, 306
 - SSO (*stateful switchover*), 318
 - storm control*, 310-314
 - STP (*Spanning Tree Protocol*), *aligned design*, 306
 - topology*, 72
- interdependence, removing*, 55-57
- loop-free topologies, ensuring*, 55-57
- monotenant data centers, 55
- multilayer topologies
 - convergence tests*, 73-76
 - hardware*, 72
 - software*, 72
- multilayer topology, 69-76
 - access layer*, 71
 - aggregation layer*, 69
 - core layer*, 69
 - networking staging for design validation*, 71-76
- multitenant data centers, 55
- N-PEs, *redundancy*, 15
- RPVST, *implementation*, 110-111, 116-122, 125, 129-130
- data plane**, 151
- data planes**
 - H-VPLS with N-PE *redundancy using EEM semaphore protocol*, 179, 203-204
 - multidomain H-VPLS with MEC and VLAN *load balancing*, 233
 - VPLS with N-PE *redundancy using EEM semaphore protocol*, 151
- data recovery**, 1
- data storm control**, DCI (data center interconnect), 310-314
- DCE (Data Center Ethernet) *initiative*, 331
- DCI (data center Interconnect), 4-7, 61, 72, 329
 - case studies*
 - GOV, 61-65
 - OUT, 61, 65-68
 - CWDM (*coarse wavelength-division multiplexing*), 5
 - dark fiber*, 5
 - design alternatives*, 57-59
 - DWDM (*dense wavelength-division multiplexing*), 5

- HSRP (Hot Standby Router Protocol), aligned design, 306
 - importance, 329
 - interdependence, removing, 55-57
 - IP routing, 329-330
 - L2 bridging, 329
 - improving with STP (Spanning Tree Protocol)*, 330-331
 - IP (Internet Protocol)*, 332-333
 - MPLS (Multiprotocol Label Switching)*, 332-333
 - new concepts*, 331-332
 - multicast deployment, Layer 2 environment, 303-305
 - OSPF (Open Shortest Path First) cost, 318
 - QinQ MAC address overlapping, 307-310
 - QoS (quality of service), 315-317
 - redundancy, adding, 5-6
 - router IDs, selecting, 319
 - routing design, 306-307
 - services modules, aligning, 306
 - SSO (stateful switchover), 318
 - storm control, 310-314
 - STP (Spanning Tree Protocol), aligned design, 306
 - topology, 72
 - debounce timers, link debounce timers, 37-38
 - dense wavelength-division multiplexing (DWDM) protection, 62
 - design alternatives, DCI (data center interconnect), 57-59
 - Dijkstra shortest path first (SPF) algorithm, 41
 - Dijkstra, Edsger, 142
 - Displaying Link Debounce and Carrier-Delay Configuration example (3-4), 38
 - distribute lists, N-PE routers, configuring, 27
 - downtime protection, 1
 - DWDM (dense wavelength-division multiplexing), DCI (data center interconnect), 5
 - DWDM (dense wavelength-division multiplexing) protection, 62
 - dynamic path (MPLS-TE), 43
 - dynamic path with bandwidth reservation (MPLS-TE), 43
-
- ## E
-
- ECMP (equal-cost multipath) balancing, 43
 - ECMP (equal-cost multipath) routing, 67
 - EEM (Embedded Event Manager), 15-18, 71
 - active/standby interchassis communication, 56
 - EEM-based deployment models, 139, 275
 - H-VPLS with N-PE redundancy using EEM*, 176-199
 - multidomain H-VPLS with dedicated U-PE*, 227
 - multidomain H-VPLS with MEC and VLAN load*, 230-270
 - multidomain H-VPLS with multichassis*, 227-230
 - multidomain H-VPLS with N-PE redundancy*, 201-221
 - N-PE redundancy using semaphore protocol*, 139-150

VPLS with N-PE redundancy using EEM, 150-172

isolated MST in N-PE deployment model, 109-110

- EEM (Embedded Event Manager) semaphore protocol, 327**
- EEM semaphore protocol, 71**
 - chassis-level synchronization, 56
 - finite state machine, 142
 - implementing, loopbacks, 148-149
 - normal mode, 143
- EEM-based deployment models, 139, 275**
 - H-VPLS with N-PE redundancy using EEM semaphore, 176-199
 - multidomain H-VPLS with dedicated U-PE, 227
 - multidomain H-VPLS with MEC and VLAN load balancing, 230-270
 - multidomain H-VPLS with multichassis EtherChannel, 227-230
 - multidomain H-VPLS with N-PE redundancy using, 201-221
 - N-PE redundancy using semaphore protocol, 139-150
 - VPLS with N-PE redundancy using EEM semaphore, 150-172
- EIGRP (Enhanced Interior Gateway Routing Protocol), 69**
- EMS (Ethernet multipoint services), 11**
- Enabling MPLS-TE example (4-2), 47**
- encapsulation**
 - frame encapsulation, QinQ, 307
 - GRE (generic routing encapsulation), 71
 - GRE (generic routing encapsulation) tunnels, 277
 - VPLSoGRE, 277-302*
 - H-VPLS, 17

- MACinMAC encapsulation, 309
- MPLS encapsulation, Ethernet frames, 23, 25
- Enhanced Interior Gateway Routing Protocol (EIGRP), 69**
- Enhanced Transmission Selection 802.1Qaz initiative, 331**
- EoMPLS (Ethernet over Multiprotocol Label Switching), 11, 42-43, 277, 332**
 - pure IP core model, 30
- equal-cost multipath (ECMP) balancing, 43**
- equal-cost multipath (ECMP) routing, 67**
- EtherChannel, 67**
- Ethernet**
 - carrier Ethernet, 330
 - EMS (Ethernet multipoint services), 11
 - EoMPLS (Ethernet over MPLS), 11, 30, 42-43, 277, 332
 - EVC (Ethernet virtual circuit), 262-263
 - frames, MPLS encapsulation, 23-25
 - interfaces, carrier-delay timers, 38-40
 - L2 bridging, 331
 - switched Ethernet service provider networks, 10
 - switches, MAC addresses, 303
- EVC (Ethernet virtual circuit), 262-263**
- event logs**
 - H-VPLS with N-PE redundancy using EEM semaphore deployment model, 199-221
 - isolated MST in N-PE, 134
 - MST in PW (pseudowire), 103

multidomain H-VPLS with MEC and VLAN load balancing EEM deployment model, 259

VPLSoGRE Multidomain with H-VPLS, 298

VPLSoGRE with N-PE redundancy using EEM semaphore deployment model, 286-291

examples

3-1 (Configuring a Distribute List on an N-PE Router), 27

3-2 (IP Event Dampening Configuration), 35

3-3 (Summary of the Dampening Parameters and Status), 35

3-4 (Displaying Link Debounce and Carrier-Delay Configuration), 38

3-5 (Carrier-Delay Timer Configuration), 39

3-6 (Confirming Carrier-Delay Status on an Interface), 39

4-1 (Setting Up MPLS-TE Tunnels), 47

4-2 (Enabling MPLS-TE), 47

4-3 (Balancing VFI Traffic), 48

4-4 (Configuring Explicit-Path Sequences), 49

4-5 (Specifying Tunnel-Path Options), 49

4-6 (Presetting a Backup Tunnel), 51

4-7 (Using an Affinity), 52

4-8 (Adding FRR to Protect the Primary Tunnel), 52

10-1 (MPLS, VFI, and SVI Configuration and Verification), 282-283

11-1 (Tuning the Multicast IGMP Query Interval), 305

11-2 (Preventing UU Flooding), 315

12-1 (Sample Configuration), 326

explicit paths (MPLS-TE), 43

creating, 48-51

FRR (Fast Reroute), adding, 50-51

extended Layer 2 networks, 7

HA (high availability), 23

convergence optimization, 32-42

core routing, 25-31

MTU (maximum transmission unit), 23-25

problems, 5-7

F

failover, HA (high-availability) clusters, 2-3

failure, primary N-PEs, 145

failure conditions, VPLS with N-PE redundancy using EEM semaphore protocol, 152-153

failure detection, routing designs, 33-40

Fast Reroute (FRR), 15, 43

fields, Ethernet frames, MPLS encapsulation, 25

file sharing, HA clusters, 2

finite state machine, EEM semaphore protocol, 142

Firewall Service Module (FWSM), QinQ bridge domain, 308

firewalls, 70

flags, semaphores, 141

flexibility, routed cores, 33

flooding

multicast traffic, 303

IGMP (Internet Group Management Protocol) snooping, 303-304

- UU (unknown unicast) flooding, preventing, 314
- flushing MAC addresses, 14
- forwarding frames, 12
- forwarding equivalence classes, labels, 19
- fragmentation, MPLS networks, 24
- frame encapsulation, QinQ, 307
- frames
 - Ethernet frames, MPLS encapsulation, 23-25
 - forwarding, 12
- FRR (Fast Reroute), 15, 43
 - IP rerouting, compared, 50
 - local repairs, 45
 - MPLS, 30
 - MPLS-TE, 44
 - affinity paths*, 52
 - explicit paths*, 50-51
- FWSM (Firewall Service Module), QinQ bridge domain, 308

G

- generic routing encapsulation (GRE), 71
- generic routing encapsulation (GRE) tunnels. *See* GRE (generic routing encapsulation)
- Global Site Selector, 1
- GOV case study (DCI), 61-65
 - challenges, 61-62
 - solutions, 62-65
- GRE (generic routing encapsulation), 71
- GRE (generic routing encapsulation) tunnels, 277
 - creating between N-PEs, 279
- MPLS, enabling on, 279-280
- VPLSoGRE, 277
 - Cisco Catalyst 6500 switches*, 278
 - configuration*, 279-282
 - multidomain with H-VPLS*, 291-302
 - N-PE (Network-facing Provider Edge) redundancy using EEM*, 278-286, 291

H

- H-VPLS (Hierarchical-Virtual Private LAN Service), 12, 17, 64, 307
 - encapsulation, 17
- H-VPLS with N-PE redundancy using EEM semaphore EEM-based deployment model, 176-221
 - control plane, 179
 - convergence tests, 195
 - data plane, 179
 - implementation, 182-192
 - N-PE router hardware, 182
 - normal mode, 179-180
 - primary N-PE node, 180-182
 - Q-link failure, 180-181
 - Q-link recovery, 181-182
 - server cluster tests, 199
- HA (high availability)
 - extended Layer 2 networks, 23
 - convergence optimization*, 32-42
 - core routing*, 25-31
 - MTU (maximum transmission unit)*, 23-25
 - routed cores, 33

- HA (high-availability) clusters, 2-4, 7**
 - business applications, 2
 - failover, 2-3
 - heartbeats, 4
 - key databases, 2
 - network file sharing, 2
 - PortFast, 4
 - private networks, 3-4
 - public networks, 2-3
 - split-brain situations, 4
 - STP (Spanning Tree Protocol), 4
 - hardware, data center multilayer topologies, 72**
 - heartbeats, campus cluster environments, 4**
 - hello messages, MPLS LDP, 22**
 - Hierarchical-Virtual Private LAN Service (H-VPLS), 307**
 - Hierarchical-VPLS (H-VPLS).**
See H-VPLS (Hierarchical-VPLS)
 - HPC Server 2008, 4**
 - HSRP (Hot Standby Router Protocol)**
 - aligning, data center interconnection, 306
 - QinQ bridge domain, 308
-
- ICCP (Inter-Chassis Communication Protocol), 56, 66, 321-323, 327**
 - AC (attachment circuit) redundancy, interaction, 324
 - BFD (Bidirectional Forwarding Detection), 323
 - emulation, N-PEs, 56-57
 - IP reachability monitoring, 323
 - Label Distribution Protocol (LDP) extension, 323
 - PE (Provider Edge) redundancy, configuring, 326-327
 - PE node failure, detection, 323
 - PW (pseudowire) redundancy, interaction, 325-326
 - topologies, support, 323
 - IEEE 802.1 DCB WG (Data Center Bridging working group) initiative, 331**
 - IEEE 802.1Q-in-Q VLAN tag (QinQ) for VLAN scalability, 55**
 - IEEE Shortest Path Bridging 802.1aq initiative, 331**
 - IETF (Internet Engineering Task Force), 321**
 - IETF Transparent Interconnection of Lots of Links (TRILL) initiative, 331**
 - IGMP (Internet Group Management Protocol) snooping, 303-304**
 - IGP (Interior Gateway Protocol), 43, 142**
 - convergence, 33
 - MPLS LDP-IGP synchronization, 21
 - MPLS networks, 27
 - implementation**
 - MPLS-TE, 46-52
 - semaphore protocol, loopbacks, 148-149
 - semaphores, 146
 - Inter-Chassis Communication Protocol (ICCP), 56, 66, 321**
 - interchassis communication, EEM (Embedded Event Manager), 56**
 - interconnecting data centers**
see also DCI (data center interconnect)
 - aligning services modules, 306
 - data center Layer 2 interconnect, 55

HSRP (Hot Standby Router Protocol)
 aligned design, 306
 multicast deployment in Layer 2
 environments, 303-305
 OSPF (Open Shortest Path First)
 cost, 318
 QinQ MAC address overlapping,
 307-310
 QoS (quality of service), 315-317
 redundancy, adding, 5
 router IDs, 319
 routing design, 306-307
 SSO (stateful switchover), 318
 storm control, 310-314
 STP (Spanning Tree Protocol) aligned
 design, 306

interfaces
 carrier-delay status, confirming, 39
 Ethernet interfaces, carrier-delay
 timers, 38-40
 POS interfaces, carrier-delay timers,
 40

**Interior Gateway Protocol (IGP). See
 IGP (Interior Gateway Protocol)**

**Intermediate System-to-Intermediate
 System (IS-IS), 62**

**Internet Engineering Task Force
 (IETF), 321**

**Internet Group Management Protocol
 (IGMP) snooping, 303**

IP (Internet Protocol)
 event dampening, configuring, 34-35
 MPLS (Multiprotocol Label
 Switching), compared, 332-333
 L2 bridging, 332-333
 traffic, 43

**IP Event Dampening Configuration
 example (3-2), 35**

**IP reachability monitoring, ICCP
 (Inter-Chassis Communication
 Protocol), 323**

IP rerouting
 DCI (data center interconnect),
 329-330
 FRR (Fast Reroute), compared, 50
 L2 bridging, compared, 329-330

**IS-IS (Intermediate System-to-
 Intermediate System), 62**

**isolated MST in N-PE deployment
 model, 77, 106-108**
 cluster server tests, 134
 convergence tests, 130
 EEM scripting, 109-110
 implementation, 110-130

K-L

key databases, HA clusters, 2

**known unicast packet storms,
 controlling, 313-314**

L2 bridging
 DCI (data center interconnect), 329
*improving with STP (Spanning
 Tree Protocol), 330-331*
 Ethernet L2 bridging, 331
 IP (Internet Protocol), 332-333
 IP routing, compared, 329-330
 MPLS (Multiprotocol Label
 Switching), 332-333
 new concepts, 331-332

**L2 broadcast packet storms,
 controlling, 312-313**

**L2 control-plane packet storm, N-PE,
 311**

L2 interconnection, STP domain, 15

- L2 known unicast packet storms, controlling, 313-314
- L2 multicast packet storms, controlling, 312-313
- L2 switching, L3 routing, compared, 9
- L2 UU (unknown unicast) packet storms, controlling, 314
- L2 VPNs (virtual private networks), 10-11
 - load balancing, MPLS, 45
- L2VPNoGRE (Layer 2 VPN over GRE), 277
- L3 headers, labels, 19
- L3 routing, L2 switching, compared, 9
- L3 VPNs (virtual private networks), 10
- Label Distribution Protocol (LDP). See LDP (Label Distribution Protocol)
- label switch paths (LSPs), 15
- label switching, MPLS (Multiprotocol Label Switching), 19-20
- labels, 19
 - see also LDP (Label Distribution Protocol)*
 - allocation limits, MPLS LDP, 21
 - forwarding equivalence classes, 19
- LACP (Link Aggregation Control Protocol), 330
- Layer 2 environments
 - data centers
 - multicast deployment, 303-305*
- Layer 2 extension, N-PE redundancy, 310
- Layer 2 networks, extended Layer 2 networks, HA (high availability), 23-42
- Layer 2 traffic integration, Layer 3 core networks
 - MPLS/IP core model, 26-27
 - pure IP core model, 30-31
 - pure MPLS core model, 28-30
- Layer 2 VPN over GRE (L2VPNoGRE), 277
- Layer 3 core networks, Layer 2 traffic integration
 - mixed MPLS/IP core model, 26-27
 - pure IP core model, 30-31
 - pure MPLS core model, 28-30
- LDP (Label Distribution Protocol), 20, 277
 - hello messages, 22
 - label allocation, limiting, 21
 - LDP-IGP synchronization, 21
 - session protection, 22
 - targeted sessions, 20
 - TCP pak priority, 21
 - traffic, 43
- LDP (Label Distribution Protocol) extension, ICCP (Inter-Chassis Communication), 323
- LDP link hello adjacency, 22
- LDP link hello message, 22
- LDP targeted hello message, 22
- learning bridge mechanisms, 11
- legacy L2 bridging, improving, STP (Spanning Tree Protocol), 330-331
- “LINEPROTO-UPDOWN” message, Cisco Catalyst 6500 switches, 40
- Link Aggregation Control Protocol (LACP), 330
- link bundles, 46
- link debounce, displaying, 38
- link debounce timers, 37-38
 - default port (link) delay time, 38

link failures

- H-VPLS with N-PE redundancy
 - using EEM semaphore protocol, 195-217
- isolated MST in N-PE, 130
- multidomain H-VPLS with MEC and VLAN load balancing, 252-270
- VPLSoGRE Multidomain with H-VPLS, 296-298
- VPLSoGRE with N-PE redundancy
 - using EEM semaphore protocol, 284-286

link node failures, MST in PW (pseudowire), 100**link-bundle balancing, MPLS-TE, 68****link-state advertisements (LSAs), 41****links**

- link bundles, 46
- MTU (Maximum Transmission Unit), 24
- naming conventions, 149-150
- parallel links, traffic repartition, 46-52

load balancing, MPLS, 45**load repetition, 67**

- MPLS-TE, 45

load-balancing technology, 1**local repairs, FRR (Fast Reroute), 45****Logical view (VSS), 228****loop-free global topologies, ensuring, 55-57****loopbacks, semaphore protocol, implementing, 148-149****loops, bridged networks, 12****LSAs (link-state advertisements), 41****LSPs (label switch paths), 15****M**

MAC addresses, 11

- Ethernet switches, 303
- flushing, 14, 321
 - asymmetric flows, 109*
 - unidirectional traffic, 109*
- overlapping, QinQ, 307-310

MACinMAC encapsulation, 309**mapping, VFIs, tunnel LSP, 48****MEC (Multichassis EtherChannel), 55, 228, 330****mLACP (multichassis Link Aggregation Control Protocol), 66, 321****monotenant data centers, 55****MPLS (Multiprotocol Label Switching), 19, 56, 62****core**

- control plane, 151*
- PWs (pseudowires), 11*
- split horizon, 11*

encapsulation, Ethernet frames, 23-25**EoMPLS (Ethernet over Multiprotocol Label Switching), 43, 277****fragmentation, 24****FRR (Fast Reroute), 30****GRE tunnels, enabling on, 279-280****IP (Internet Protocol), compared, 332-333****L2 bridging, 332-333****L2 VPN load balancing, 45****label switching, 19-20****LDP (Label Distribution Protocol), 20**
hello messages, 22

- LDP-IGP synchronization*, 21
- limiting label allocation*, 21
- session protection*, 22
- targeted sessions*, 20
- TCP pak priority*, 21
- load balancing, 45
- MPLS-TE (MPLS-Traffic Engineering), 29, 43-45, 53
 - affinity path*, 43, 52
 - constraints*, 43-44
 - dynamic path*, 43
 - dynamic path with bandwidth reservation*, 43
 - enabling*, 47
 - explicit path*, 43, 48-51
 - FRR (Fast Reroute)*, 44
 - implementation*, 46-52
 - load repetition*, 45
 - pre-established backup path*, 43
 - tunnel LSP*, 44, 47-48
- N-PEs, configuring on, 95, 98
- Networks, creating, 27
- PWs, 11
- VPLS, 277
- VRF, 29
- mpls ldp neighbor targeted command**, 20
- mpls ldp tcp pak-priority command**, 21
- MPLS Transport Profile (MPLS-TP)**, 332
- MPLS, VFI, and SVI Configuration and Verification example (10-1)**, 282-283
- MPLS-TE**, 29, 43-45, 53
 - affinity path*, 43
 - affinity paths*
 - creating*, 52
 - FRR (Fast Reroute)*, 52
 - constraints*, 43-44
 - dynamic path*, 43
 - dynamic path with bandwidth reservation*, 43
 - enabling*, 47
 - explicit path*, 43
 - explicit paths*
 - creating*, 48-51
 - FRR (Fast Reroute)*, 50-51
 - FRR (Fast Reroute)*, 44
 - link bundle balancing*, 68
 - load repetition*, 45
 - parallel links, implementation for traffic repartition*, 46-52
 - pre-established backup path*, 43
 - tunnel LSP*, 44
 - balancing*, 48
 - mapping VFIs to*, 48
 - setting up*, 47
 - VFI traffic*, 48
- MPLS-TP (MPLS Transport Profile)**, 332
- MPLS/IP core model, Layer 3 core networks, Layer 2 traffic integration**, 26-27
- MPLSoGRE approach, pure IP core model**, 31
- MST (Multiple Spanning Tree)**, 7, 55
 - BPDU (bridge protocol data units), 78
 - MST-based deployment models
 - isolated MST in N-PE*, 77, 106-134
 - MST in N-PE*, 77-81, 90-103
 - N-PEs (Network Provider Edges), VPLS access, 56
 - PWs (pseudowires), implementation, 80-81

MST (Multiple Spanning Tree)-based deployment models

isolated MST in N, 110-130

isolated MST in N-PE, 77, 106-108

*cluster, 134**EEM, 109-110*

MST on N-PE, 77-80

*cluster server, 103**convergence, 100**implementation, 80-81, 90-99***MST in N-PE deployment model, 77-80**

cluster server tests, 103

convergence tests, 100

implementation, 80-99

MTU (maximum transmission unit)

calculation summary, 23

configuring, 24

evaluation for intersite core transport, 23-25

links, 24

PSNs (packet switched networks), 23

multicast deployment, data centers, Layer 2 environments, 303-305***Multicast Does Not Work in the Same VLAN in Catalyst Switches, 304*****multicast IGMP query intervals, tuning, 304-305****multicast packet storms, controlling, 312-313****multicast traffic, flooding, IGMP (Internet Group Management Protocol) snooping, 303-304****Multichassis EtherChannel (MEC), 55, 330****multichassis Link Aggregation**

Control Protocol (mLACP), 66, 321

multidomain H-VPLS with dedicated U-PE EEM deployment model, 227**multidomain H-VPLS with MEC and VLAN load balancing EEM deployment model, 233-270****multidomain H-VPLS with multichassis EtherChannel EEM deployment model, 227-230****Multidomain with H-VPLS**

VPLSoGRE, 291-293, 298, 302

*cluster server tests, 298**convergence testing, 296-298***multilayer topologies, data centers, 69-76**

access layer, 71

aggregation layer, 69

convergence tests, 73-76

core layer, 69

hardware, 72

network staging for design validation, 71-76

software, 72

multipath VPLS with STP isolation, 15**Multiple Spanning Tree (MST). See MST (Multiple Spanning Tree)****multipoint Ethernet LAN services, 11****Multiprotocol Label Switching (MPLS). See MPLS (Multiprotocol Label Switching)****multisite Layer 2 transport over Layer 3 protected links, 55****multisite virtualized data centers, 66****multitenant data centers, 55**

N

N-PE redundancy

H-VPLS with N-PE redundancy using EEM semaphore protocol, 176-199

Layer 2 extension, 310

MST in NPE deployment model, 78-80

multidomain H-VPLS with dedicated U-PE, 227

multidomain H-VPLS with MEC and VLAN load balancing, 230-270

multidomain H-VPLS with multichassis EtherChannel, 227-230

multidomain H-VPLS with N-PE redundancy using EEM semaphore, 201-221

semaphore protocol, 139-150

VPLS with N-PE redundancy using EEM semaphore protocol, 150-172

N-PE1

configuration, 280

STP mode MST, configuring, 81-82

N-PEs

backup N-PEs, 142

GRE tunnels, creating between, 279

hardware, VPLS with N-PE redundancy using EEM semaphore protocol, 154

ICCP (Inter-Chassis Communication Protocol), 56-57

isolated MST in N-PE deployment model, 77, 106-108

cluster server tests, 134

convergence tests, 130

EEM scripting, 109-110

implementation, 110-130

L2 control-plane packet storm, 311

MPLS, configuring on, 95, 98

MST (Multiple Spanning Tree), VPLS access, 56

MST in N-PE deployment model, 77-80

cluster server tests, 103

convergence tests, 100

implementation, 80-99

nodes, on stick design, 64

OSPF (Open Shortest Path First) Protocol, configuring on, 93, 95

primary N-PEs, 142

failure, 145

recovery, 145-146

Q-link interface, 177

QinQ VLANs, creating on, 291, 293

redundancy, 68

data centers, 15

routers, distribute lists, 27

software, VPLS with N-PE redundancy using EEM semaphore protocol, 154

STP (Spanning Tree Protocol), 56

participation, 311

SVI (switched virtual interface), configuring on, 95-98

VFI (virtual forwarding instance), configuring on, 95-98

naming conventions

links, 149-150

nodes, 149-150

near-line disks, 2

neighbor relationships, BFD (Bidirectional Forwarding Detection), establishing, 36

network file sharing, HA clusters, 2

network staging for design validation, data centers, 71-76

networks

convergence, 32

extended Layer 2 networks

HA (high availability), 23-42

problems, 5-7

private networks

HA (high-availability) clusters, 3-4

server-to-server

communications, 3

public networks, HA (high-availability) clusters, 2-3

Nexus Virtual PortChannel (vPC), 55

no mpls ldp advertise-labels command, 21

node failure

H-VPLS with N-PE redundancy using EEM semaphore protocol, 195-217

isolated MST in N-PE, 130

multidomain H-VPLS with MEC and VLAN load balancing, 252-270

VPLSoGRE Multidomain with H-VPLS, 296-298

VPLSoGRE with N-PE redundancy using EEM semaphore, 284-286

primary N-PEs, 153-154

H-VPLS with N-PE redundancy using EEM semaphore, 180-181, 204-205

multidomain H-VPLS with MEC and VLAN load balancing, 234-235

node recovery

primary N-PEs, 154

H-VPLS with N-PE redundancy using EEM semaphore, 181-182, 205-207

multidomain H-VPLS with MEC and VLAN load, 235

node redundancy, 321

node synchronization, active/standby VPLS node synchronization, 139-140

nodes

naming conventions, 149-150

synchronization, ICCP (Inter-Chassis Communication Protocol), 321

normal mode

H-VPLS with N-PE redundancy using EEM semaphore, 204

H-VPLS with N-PE redundancy using EEM semaphore protocol, 179-180

multidomain H-VPLS with MEC and VLAN load balancing EEM, 233

PWs (pseudowires), 142-145

VPLS with N-PE redundancy using EEM semaphore protocol, 151-152

O

object tracking, configuring, EEM, 98-99

on a stick design, VPLS N-PE, 64

one rack unit (1RU) servers, 71

Open Shortest Path First (OSPF) Protocol. See OSPF (Open Shortest Path First) protocol

Open Systems Adapters (OSA) adapters, 71

Open Systems Interconnection (OSI) model, 6

optimization, convergence optimization, extended Layer 2 network HA (high availability), 32-42

- option 1 design (N-PE redundancy using semaphore protocol), 139-150
- option 1a design (MST in N-PE deployment model), 77-80
 - cluster server tests, 103
 - convergence tests, 100
 - implementation, 80-99
- option 1b design (isolated MST in N-PE), 77, 106-108
 - cluster server tests, 134
 - convergence tests, 130
 - EEM scripting, 109-110
 - implementation, 110-130
- option 2 design (VPLS with N-PE redundancy using EEM semaphore protocol), 150-172
- option 3 design (H-VPLS with N-PE redundancy using EEM semaphore protocol), 176-199
- option 4a design (multidomain H-VPLS with N-PE redundancy using EEM semaphore, 201-221
- option 4b design (multidomain H-VPLS with dedicated U-PE), 227
- option 5a design (multidomain H-VPLS with multichassis EtherChannel), 227-230
- option 5b design (multidomain H-VPLS with MEC and VLAN load balancing), 230-259
- option 5c design (multidomain H-VPLS with MEC and VLAN load balancing), 262-270
- OSA (Open Systems Adapters) adapters, 71
- OSI (Open Systems Interconnection) model, 6
- OSPF (Open Shortest Path First), 62, 69

DCI (data center interconnect), costs, 318

N-PE routers, configuring on, 93, 95
neighbor relationships, tearing down, 37

OUT case study (DCI), 61, 65-68
challenges, 65
solution, 65-68

P

parallel links, traffic repartition, MPLS-TE, 46-52

partially meshed PW topology, 143

partitions, traffic repartition, parallel links, 46-52

PE (Provider Edge) redundancy, configuring, ICCP (Inter-Chassis Communication), 326-327

PE (Provider Edge) routers
node failure, detection, 323
redundancy, 321

PE routers, PWs, 14-15

policies (EEM), 18

PortFast, 4

POS interfaces, carrier-delay timers, 40

preestablished backup path (MPLS-TE), 43

preferred-path command, 48

Presetting a Backup Tunnel example (4-6), 51

Preventing UU Flooding example (11-2), 315

primary N-PE node failure, H-VPLS with N-PE redundancy using EEM semaphore, 204-205

primary N-PE node recovery, H-VPLS with N-PE redundancy using EEM semaphore, 205-207

primary N-PEs, 142

failure, 145

node failure

H-VPLS with N-PE redundancy using EEM semaphore, 180-181

multidomain H-VPLS with MEC and VLAN load balancing, 234-235

VPLS with N-PE redundancy using EEM semaphore, 153-154

node recovery

H-VPLS with N-PE redundancy using EEM semaphore, 181-182

multidomain H-VPLS with MEC and VLAN load, 235

VPLS with N-PE redundancy using EEM semaphore, 154

recovery, 145-146

private networks

HA (high-availability) clusters, 3-4
server-to-server communications, 3

process-max-time 50 command, 18

pseudowire (PW) redundancy, 321

PSNs (packet switched networks), MTU (maximum transmission unit), 23

public networks, HA (high-availability) clusters, 2-3

pure IP core model, Layer 3 core networks, Layer 2 traffic integration, 30-31

pure MPLS core model, Layer 3 core networks, Layer 2 traffic integration, 28-30

PW (pseudowire) redundancy

PWs (pseudowire), 14-15

classes, 15

MPLS core, 11

MPLS PWs, 11

MST (Multiple Spanning Tree), implementing in, 80-81

normal mode, 142-145

partially meshed PW topology, 143

redundancy

interaction, 325-326

standby states, 321

signaling, 12

VCLS, 12

VFIs, 14

xconnect, 14

Q

Q-link failure, H-VPLS with N-PE redundancy using EEM semaphore EEM-based, 180-205

Q-Link interface, N-PEs, 177

Q-link recovery, H-VPLS with N-PE redundancy using EEM semaphore EEM-based, 181-207

QinQ (802.1Q in 802.1Q), 17, 139, 307

bridge domain, 308

Cisco Catalyst 6500 switches, 291

frame encapsulation, 307

MAC address overlapping, data center interconnection, 307-310

VFI (virtual forwarding instance), 307

QinQ VLANs, N-PEs, creating on, 291, 293

QoS (quality of service), DCI (data center interconnect), 315-317

R

R-L2GP (Reverse Layer 2 Gateway Protocol), 330

Rapid Spanning Tree Protocol (RSTP), 77

Rapid STP (RSTP), 7

Rbridging, 332

recovery, primary N-PEs, 145-146

redundancy

- AC (attachment circuit) redundancy, 321
 - interaction, 324*
- DCI (data center interconnect), 5-6
- N-PE redundancy, semaphore protocol, 139-150
- N-PEs, data centers, 15
- node redundancy, 321
- PE (Provider Edge) redundancy, configuring, 326-327
- PE (Provider Edge) routers, 321
- PW (pseudowire) redundancy
 - interaction, 325-326*
 - standby states, 321*
- VPLS N-PE, 68

remote electronic vaulting, 2

replication over WANs, 2

resilience, applications, enhancing, 1

Reverse Layer 2 Gateway Protocol (R-L2GP), 330

RHI (Route Health Injection), 306

route computation, routing designs, 40-42

Route Health Injection (RHI), 306

routed cores, advantages, 33

router IDs, DCI (data center interconnect), selecting, 319

routers

- N-PE routers, distribute lists, 27
- PE routers, PWs, 14-15

routing

- constraint-based routing, 30
- core routing, extended Layer 2 network HA (high availability), 25-31
- data center interconnection, 306-307
- RHI (Route Health Injection), 306

routing designs

- alternate route computation, 40-42
- failure detection, 33-40
- tuning, 33-40

RSTP (Rapid Spanning Tree Protocol), 77

RSTP (Rapid STP), 7

S

S-TAG (service provider VLAN tag), 307

Sample Configuration example (12-1), 326

scalability, routed cores, 33

scripts, EEM scripts, 147-148

semaphore protocol

- H-VPLS with N-PE redundancy, 176-199
 - control plane, 179*
 - convergence tests, 195*
 - data plane, 179*
 - implementation, 182-192*
 - N-PE router hardware and software, 182*
 - normal mode, 179-180*
 - primary N-PE node failure, 180-181*

- primary N-PE node recovery, 181-182*
- Q-link failure, 180-181*
- Q-link recovery, 181-182*
- server cluster tests, 199*
- implementing, loopbacks, 148-149
- multidomain H-VPLS with N-PE redundancy, 201-221
 - control plane, 203*
 - convergence tests, 217*
 - data plane, 203-204*
 - implementation, 207-217*
 - N-PE router hardware, 207*
 - normal mode, 204*
 - primary N-PE node, 204-207*
 - Q-link failure, 204-205*
 - Q-link recovery, 205-207*
 - server cluster tests, 221*
- N-PE redundancy, 139-150
- VPLS with N-PE redundancy, 150-172
 - cluster server tests, 172*
 - control planes, 151*
 - convergence tests, 168*
 - data plane, 151*
 - failure conditions, 152-153*
 - implementation, 154-166*
 - N-PE hardware and software, 154*
 - normal mode, 151-152*
 - primary N-PE node failure, 153-154*
 - primary N-PE node recovery, 154*
- semaphores, 141-142**
 - flags, 141
 - implementing, 146
- server cluster tests**
 - multidomain H-VPLS with N-PE redundancy, 221
 - multidomain H-VPLS with MEC and VLAN load balancing, 259
- server-to-server communications, private networks, 3**
- servers, 71**
- service provider VLAN tag (S-TAG), 307**
- service provider WAN aggregation layer, 330**
- services modules, aligning, data center interconnection, 306**
- session protection, MPLS LDP, 22**
- Setting Up MPLS-TE Tunnels example (4-1), 47**
- show interface command, 39**
- show interface dampening command, 35**
- show mpls l2transport vc command, 163, 190**
- show running-config | begin event manager command, 166, 192, 248**
- signaling, PWs, 12**
- sites, interconnection, link bundles, 46**
- software, data center multilayer topologies, 72**
- spanning tree control plane, 151**
- Spanning Tree Protocol (STP), 62**
- spanning-tree algorithm (STA), 6**
- spanning-tree isolation, 55**
- Specifying Tunnel-Path Options example (4-5), 49**
- SPF (shortest path first)**
 - Dijkstra shortest path first (SPF) algorithm, 41
 - throttle mechanism, 41

split horizon, MPLS core, 11
split-brain situations, clusters, 4
split-subnets, 5
SSL (Secure Sockets Layer), 70
SSO (stateful switchover), DCI (data center interconnect), 318
STA (spanning-tree algorithm), 6
standby states, PW (pseudowire) redundancy, 321
storm control for broadcast and multicast, 65
storm control for DCI (data center interconnect), 310-314
STP (Spanning Tree Protocol), 6, 56, 62
 aligning, data center interconnection, 306
 HA (high-availability) clusters, 4
 L2 bridging, improving, 330-331
 multipath VPLS with STP isolation, 15
 N-PE, participation, 311
STP mode MST, N-PE1, configuring, 81-82
Summary of the Dampening Parameters and Status example (3-3), 35
SVIs (switched virtual interfaces)
 N-PEs, configuring on, 95, 98
 QinQ bridge domain, 309
switched Ethernet service provider networks, 10
switches
 aggregation switches, 55
 Cisco Catalyst 6500 switches
 “LINEPROTO-UPDOWN” message, 40
 PortFast, 4
 QinQ, 291
 STP (Spanning Tree Protocol), 4

 Ethernet switches, MAC addresses, 303
 intregal switches, blade servers, 71
switching
 content switching, 70
 MPLS (Multiprotocol Label Switching), 277
 VSS (Virtual Switching System), 55, 330
synchronization
 MPLS LDP-IGP synchronization, 21
 node synchronization, active/standby VPLS node synchronization, 139-140
synchronous disk mirroring, 2

T

targeted session, MPLS LDP, 20
TCL (Tool Command Language), 18
TCP pak priority, MPLS LDP, 21
TE (traffic engineering) tunnels, 15, 43
 MPLS-TE, 29, 43-45, 53
 affinity path, 43, 52
 constraints, 43-44
 dynamic path, 43
 dynamic path with bandwidth reservation, 43
 enabling, 47
 explicit path, 43, 48-51
 FRR (Fast Reroute), 44
 implementation, 46-52
 load repetition, 45
 pre-established backup path, 43
 tunnel LSP, 44, 47-48
 setting up, 47
 VFI, mapping to, 48

throttle mechanism, SPF (shortest path first), 41

TLS (transparent LAN services), 10-11

Tool Command Language (Tcl), 18

topologies

- data centers, multilayer topologies, 69-76
- ICCP (Inter-Chassis Communication Protocol), support, 323
- loop-free global topologies, ensuring, 55-57
- partially meshed PW topology, 143

traffic

- multicast traffic, flooding, 303
- TE (traffic engineering), 43-53

traffic flow, convergence tests, data centers, 73-76

traffic profile, convergence tests, data centers, 74-76

traffic rate, convergence tests, data centers, 73

traffic repartition, parallel links, MPLS-TE, 46-52

transparent LAN services (TLS), 10-11

TRILL (Transparent Interconnection of Lots of Links) initiative, 331

tuning

- multicast IGMP query intervals, 304-305
- routing designs, 33-40

Tuning the Multicast IGMP Query Interval example (11-1), 305

tunnel LSP, MPLS-TE, 44

- mapping VFIs to, 48
- setting up, 47

tunnels

- see also TE (traffic engineering) tunnels*
- backup tunnels, presetting, 51

- tunnel LPS, 44
- mapping VFIs to, 48*
- setting up, 47*

U

Understanding and Tuning Spanning Tree Protocol Timers, 7

unidirectional traffic, MAC address flushing, 109

user access, 1

Using an Affinity example (4-7), 52

Using PortFast and Other Commands to Fix Workstation Startup Connectivity Delays, 4

UU (unknown unicast) flooding, preventing, 314

UU (unknown unicast) packet storms, controlling, 314

V-Z

VCs (virtual circuits), 12

VFI (virtual forwarding instance)

- N-PEs, configuring on, 95-98
- PWs (pseudowires), 14
- QinQ, 307
- Traffic, balancing, 48
- tunnel LSP, mapping to, 48

VFI for VLAN 7, 282

Virtual PortChannel (vPC), 55, 330

Virtual Private LAN Service (VPLS).

- See VPLS (Virtual Private LAN Service)*

virtual private networks (VPNs), 9

Virtual Routing and Forwarding (VRF). *See VRF (Virtual Routing and Forwarding)*

- virtual switch instances (VSIs), 11-12
- Virtual Switching System (VSS). *See* VSS (Virtual Switching System)
- virtualization, MPLS, 29
- VMotion, 61
- VMware machines, QinQ bridge domain, 308
- vPC (Virtual PortChannel), 55, 330
- VPLS (Virtual Private LAN Service), 11-14, 55
 - access
 - N-PEs (Network Provider Edges) using ICCP (Inter-Chassis)*, 56-57
 - N-PEs (Network Provider Edges) using MST (Multiple Spanning Tree)*, 56
 - architecture, 12
 - attachment circuits, 12
 - components, 13-14
 - control plane, 151
 - data centers, loop-free topologies, 55-57
 - DCI (data center interconnect), 61
 - EMS (Ethernet multipoint services), 11
 - frames, forwarding, 12
 - global VPLS architecture, 63-64
 - H-VPLS (Hierarchical-Virtual Private LAN Service), 307
 - encapsulation*, 17
 - MPLS (Multiprotocol Label Switching), 277
 - N-PE (Network Provider Edge) nodes, 64
 - nodes, on a stick design*, 64
 - redundancy*, 68
 - PEs, autodiscovery, 12
 - VPLS with N-PE redundancy using EEM semaphore EEM-based deployment model, 150-172
 - VPLSoGRE, 277
 - Cisco Catalyst 6500 switches, 278
 - configuration, 279-282
 - Multidomain with H-VPLS, 291-293, 298, 302
 - cluster server tests*, 298
 - convergence testing*, 296-298
 - N-PE (Network-facing Provider Edge) redundancy using EEM semaphore, 278-291
 - pure IP core model, 30
 - VPNs (virtual private networks), 9
 - Layer 2 VPNs, 10-11
 - Layer 3 VPNs, 10
 - VRF (Virtual Routing and Forwarding), 62, 330
 - MPLS, 29
 - VSIs (virtual switch instances), 11-12
 - VSS (Virtual Switching System), 55, 227, 330
 - Logical view, 228
 - WANs (wide area networks), replication over WANs, 2
 - Windows HPC Server 2008, 4
 - xconnect (cross connection), PWs (pseudowires), 14