



Building Service-Aware Networks

The Next-Generation WAN/MAN

A thorough introduction to the ASR 1000 series router

Building Service-Aware Networks: The Next-Generation WAN/MAN

Muhammad Afaq Khan

Copyright© 2010 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 September 2009

Library of Congress Cataloging-in-Publication Data:

Library of Congress Cataloging-in-Publication data is on file.

ISBN-13: 978-1-58705-788-5

ISBN-10: 1-58705-788-3

Warning and Disclaimer

This book is designed to provide information about enterprise WAN/MAN network architectures: design, selecting and qualifying, deploying, and troubleshooting the router. 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 author, 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.

Introduction

Both traditional WANs and next-generation WAN aggregation are an integral part of any network design. This is where you aggregate all your traffic coming from branches of all sizes. This place in the network, due to being the center of all traffic aggregation, becomes extremely crucial and must handle failures both from platform and network architecture perspectives. Therefore, just as we must understand WANs, emerging trends, and how those trends are creating newer requirements for WAN routing products, it is equally important to understand how to select, qualify, and adopt them into your network.

This book covers all topics related to WANs/MANs. It starts with an introduction to WAN/MAN architectures, followed by information specific to Cisco ASR 1000 series routers. The book concludes with five chapters of real-world use cases based on the ASR 1000 platform.

Goals and Methods

This book can help you understand existing and emerging WAN architectures, including the underlying business trends and how they are creating newer demands for the infrastructure. The hardware and software architectures of the Cisco ASR 1000 series routers are discussed as a perfect example of the modern infrastructure, and that concept is reinforced using various real-world use cases in the last five chapters of this book.

The book is organized into four parts. Part I discusses the WAN architectures, Part II introduces Cisco ASR 1000 series routers, Part III goes over system management and troubleshooting, and Part IV concludes the book with many real-world examples.

Who Should Read This Book?

This book does not focus on a single technical topic or area. Instead, it covers a variety of technical topics that are useful within the context of overall WAN/MAN architectures. This book introduces the reader to both the technical and business landscapes surrounding WAN/MAN technologies.

Other objectives, such as qualifying and sizing up routing equipment from any vendor or understanding various system architectural aspects of a modern system, can also be achieved by reading this book.

Hence, this book can serve well the needs of network engineers, architects, and network visionaries, those who must understand the changes happening within the WAN/MAN arena and the infrastructure design imperatives of next-generation WANs/MANs.

How This Book Is Organized

Although this book can be read cover to cover, it is designed to be flexible and enable you to easily move between chapters and sections of chapters to cover just the material that you want to focus on.

Part I covers the details of WAN/MAN architectures, changes that are happening, and the associated business drivers of those changes. It also covers the selection, qualification, and how to size up a routing device for a given deployment.

Part II discusses Cisco ASR 1000 series routers (including hardware/software architectures). Part II also covers IOS XE packaging, releases, initial setup, and extremely important topics such as in-service software upgrades (ISSU).

Part III focuses on system management and troubleshooting, including some common error messages that you may encounter during deployment.

Part IV delves deeper into actual deployment scenarios, via use cases. A wide variety of topics are covered, from pure routing/switching, IP services, WAN optimization using Web Cache Communication Protocol Version 2 (WCCPv2), and unified communication topics such as Cisco Unified Border Element (CUBE) and the WebEx Node module.

Chapters 1 to 16 cover the following topics:

- **Chapter 1, “Introduction to WAN Architectures,”** discusses various types of WAN architectures and associated infrastructure requirements to build them. Areas covered include branch WAN aggregation, private WAN aggregation, the Internet edge, data center interconnect, and large branch WANs.
- **Chapter 2, “Next-Generation WAN Architectures,”** builds upon the basic ideas discussed in Chapter 1. It discusses the evolution of WAN architectures, and the drivers behind them, with references to various research done in the area. It also links these drivers to the newer demands that they are creating for the underlying routing network infrastructure.
- **Chapter 3, “Selecting and Qualifying Enterprise Edge Platforms for Next-Generation WANs,”** covers attributes of a next-generation enterprise WAN/MAN platform, to help you select the appropriate devices. It also discusses how to qualify a routing device by using a methodical approach and a test plan.
- **Chapter 4, “Sizing Up a Router,”** examines essential traits of a modern routing platform for WANs/MANs. It also includes benchmarks that prove useful for testing and sizing up a WAN router.
- **Chapter 5, “System Overview and Carrier-Class Attributes,”** introduces the architecture of Cisco ASR 1000 series routers. This includes an overview of both hardware and software of the platform.
- **Chapter 6, “Cisco ASR 1000 Series Router Hardware and Software Details,”** covers Cisco ASR 1000 series routers in detail. It delves into both hardware and software components and expands on the subject matter that was briefly discussed in Chapter 5.

- **Chapter 7, “Cisco IOS XE Software Packaging, Releases, and Licensing,”** covers the IOS XE operating system (its packaging, releases, and licensing) and examines the concept of software redundancy.
- **Chapter 8, “Cisco ASR 1000 Initial Setup and Configuration,”** goes over the initial ASR 1000 system setup and configuration. This chapter serves as a necessary introduction to prepare you for the use cases in Chapter 12 to Chapter 16.
- **Chapter 9, “In-Service Software Upgrade and Software Modularity,”** discusses only one topic: in-service software upgrades (including various booting options). It also examines the different levels of ISSU support available across various packages, IOS flavors, and ASR 1000 chassis.
- **Chapter 10, “Using the ASR 1000 Embedded Graphical User Interface,”** covers the ASR 1000-specific graphical user interface (GUI) and includes a few use cases.
- **Chapter 11, “Understanding ASR 1000 System Troubleshooting and Error Messages,”** covers basic troubleshooting methodology. The focus of the chapter is on IOS and ASR 1000 **show** and **debug** commands. This chapter also examines some common error messages that you may encounter on ASR 1000.
- **Chapter 12, “IP Routing Use Cases,”** starts off the ASR 1000 use cases by using basic IP routing for example purposes. It discusses high availability (HA), route reflectors, and data center interconnect scenarios. These are explained using all relevant configuration snippets and **show** commands.
- **Chapter 13, “IP Services Use Cases,”** covers IOS services such as quality of service (QoS), Network Address Translation (NAT), multicast, NetFlow, and NetFlow event logging. These are explained using all relevant configuration snippets and **show** commands.
- **Chapter 14, “Security Services Use Cases,”** discusses IOS security services available in IOS XE. It uses various IPsec solutions such as Group Encrypted Transport VPNs (GET VPN), IOS zone-based firewalls, and their interaction with other features such as multicast and QoS.
- **Chapter 15, “WAN Optimization Services Use Cases,”** covers fundamental router-based WAN optimization solutions such as compressed Real-Time Protocol (cRTP) but also includes the integration with Cisco Wide Area Application Services (WAAS) and Iron Port’s Web Security Appliances (WSA) with ASR 1000.
- **Chapter 16, “Unified Communication Services Use Cases,”** rounds off the discussion with Cisco unified communication. This chapter discusses the Cisco Unified Border Element (CUBE) SP (service provider) Edition and various associated use cases. It also discusses the Cisco WebEx Node module and provides complete integration details. Chapters 12 to 16 evidence how ASR 1000 offers flexibility without compromising scale and performance.

IP Routing Use Cases

This chapter focuses on routing and switching with ASR 1000 series routers. The chapter starts with a brief discussion of the capabilities of the router family, and then reviews how those strengths can be used to address relevant problems.

This approach, with detailed configuration examples where applicable, will allow you to understand the problems, the challenges they represent, and how you can use the ASR 1000 to address them.

Introduction to the Scalable and Modular Control Plane on the ASR 1000

The control plane is a logical concept that defines the part of the router architecture responsible for building and drawing the network topology map (also known as the routing table) and manifesting it to the forwarding plane (where actual packet forwarding takes place) in the form of the Forwarding Information Base (FIB).

While the forwarding capacity of the routers has continuously scaled throughout the years (the Cisco CRS-1, for example; the forwarding capacity for which boosted up to 92 terabits per second [Tbps]), the control-plane scale is given less attention. When routing products are compared, the focus is usually on the forwarding capacity (packets per second or bits per second).

Contrary to this popular notion, the control-plane scale is equally critical to ensure that the platform has the compute cycles in the form of Route Processor (RP) CPU to perform the following (among other things):

- CLI, and similar external management functions performed via Simple Network Management Protocol (SNMP) or Extensible Markup Language (XML)
- Routing protocols and their associated keepalives (including crypto functions in the control plane)
- Link-layer protocols and their associated keepalives
- Services such as RADIUS, TACACS+, DHCP, Session Border Controller (SBC), and Performance-based Routing (PfR) Master Controller function
- All other traffic that cannot be handled at the data plane (for example, legacy protocols such as IPX), including punt traffic

The Cisco ASR 1000 router series delivers complete separation of the control and data plane, which enables the infrastructure's control plane to scale independently of the data plane. The ASR 1000 has two RPs on the market today: ASR1000-RP1 (first generation) and ASR1000-RP2 (second generation). ASR1000-RP1 is based on a 1.5-GHz RP CPU, whereas ASR1000-RP2 hosts a dual-core Intel 2.66-GHz processor, literally increasing the scale many times over the ASR1000-RP1.

The central benefit of physically separating the forwarding and control planes is that if the traffic load becomes very heavy (the forwarding plane gets overwhelmed), it simply doesn't affect the control plane's capability to process new routing information.

Another way of looking at it is if the routing plane gets very busy because of any of the relevant tasks, causing the control plane to be busy, perhaps because of a flood of new route information (even worse, peer or prefix flaps), busy-ness doesn't adversely affect the capability of the forwarding plane to continue forwarding packets. This is a common problem that plagues all software-based routers (due to a single general-purpose CPU running both control and data planes).

Note ASR1000-RP1 and ASR1000-RP2 use different binaries (RP packages), because of the different processor architectures, where ASR1000-RP1 uses PowerPC and ASR1000-RP2 uses an Intel-based architecture. Hence, in-service software upgrades (ISSU) between those two images is not possible without the RP reboot (for a single-RP platform). However, there is no difference in IOS feature content present in either of them.

Key applications that benefit the most, from a big picture perspective, are network virtualization, infrastructure consolidation, and rapid rollout of various network-based services.

Before delving further and discussing the actual use cases from real-world networks, a quick refresher is in order on some commonly used terms.

NSF/SSO, NSR, Graceful Restart to Ensure Robust Routing

Nonstop forwarding (NSF) refers to the capability of the data plane to continue to function hitless when the routing plane disappears (momentarily, that is) and most likely fails over to a standby RP. Of course, the routing information and topology might change during this time and result in an invalid FIB, and therefore the switchover times should be as small as possible. The Cisco ASR 1000 provides switchover times of less than 50 ms RP to RP (or IOS daemon [IOSD] to IOSD for the ASR 1002-F/ASR 1002/ASR 1004).

Stateful switchover (SSO) refers to the capability of the control plane to hold configuration and various states during this switchover, and to thus effectively reduce the time to utilize the newly failed-over control plane. This is also handy when doing scheduled hitless upgrades within the ISSU execution path. The time to reach SSO for the newly active RP may vary depending on the type and scale of the configuration.

Graceful restart (GR) refers to the capability of the control plane to delay advertising the absence of a peer (going through control-plane switchover) for a “grace period,” and thus help minimize disruption during that time (assuming the standby control plane comes up). GR is based on extensions per routing protocol, which are interoperable across vendors. The downside of the grace period is huge when the peer completely fails and never comes up, because that slows down the overall network convergence, which brings us to the final concept: nonstop routing (NSR).

NSR is an internal (vendor-specific) mechanism to extend the awareness of routing to the standby routing plane so that in case of failover, the newly active routing plane can take charge of the already established sessions.

Table 12-1 shows the compatibility and support matrix for ASR 1000 IOS XE software 2.2, and outlines the various states that are preserved during FP/ESP failover.

See the “Further Reading” section at the end of this chapter to find out where to look for complete route scale testing details.

Use Case: Achieving High Availability Using NSF/SSO

To command higher revenues and consistent profitability, service providers and enterprises are increasingly putting more mission-critical, time-sensitive services on their IP infrastructure. One of the key challenges to this is achieving and delivering high network availability with strict service level agreement (SLA) requirements. It is universally understood that availability of the network is directly linked with the overall total cost of ownership (TCO).

An enterprise has an ASR 1006 / ASR1000-ESP10 router used in the core of the network running OSPF as the routing protocol used to connect to multiple distribution hub routers, where distribution hub routers might not all be Cisco.

The goal is to reduce the route/prefix recomputation churn caused by RP switchover and reestablishment of OSPF peers.

Table 12-1 *Protocols and Their State Preservation via NSF/SSO*

Technology Focus	NSF	SSO
Routing protocols	Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First Version 2 (OSPFv2), OSPFv3, Intermediate System-to-Intermediate System (IS-IS), and Border Gateway Protocol Version 4 (BGPv4)	
IPv4 services	—	Address Resolution Protocol (ARP), Hot Standby Routing Protocol (HSRP), IPsec, Network Address Translation (NAT), IPv6 Neighbor Discovery Protocol (NDP), Unicast Reverse Path Forwarding (uRPF), Simple Network Management Protocol (SNMP), Gateway Load Balancing Protocol (GLBP), Virtual Router Redundancy Protocol (VRRP), Multicast (Internet Group Management Protocol [IGMP])
IPv6 services	—	IPv6 Multicast (Multicast Listener Discovery [MLD], Protocol Independent Multicast-Source Specific Multicast [PIM-SSM], MLD Access group)
L2/L3 protocols	—	Frame Relay, PPP, Multilink PPP (MLPPP), High-Level Data Link Control (HDLC), 802.1Q, bidirectional forwarding detection (BFD)
Multiprotocol Label Switching (MPLS)	—	MPLS Layer 3 VPN (L3 VPN), MPLS Label Distribution Protocol (LDP)
SBC	—	SBC Data Border Element (DBE)

To address the requirements, you need to implement Internet Engineering Task Force (IETF) NSF for OSPF because that is interoperable with all vendors that are NSF-aware (a term used for a neighboring router that understands the GR protocol extensions). In this case, when NSF-capable ASR 1000 switches over from active RP to standby RP, there will be no packet loss at all, and downstream neighbors will not restart adjacencies.

Figure 12-1 shows the ASR 1000 core router and its neighbors, which are all NSF-aware and can act as helpers during RP SSO.

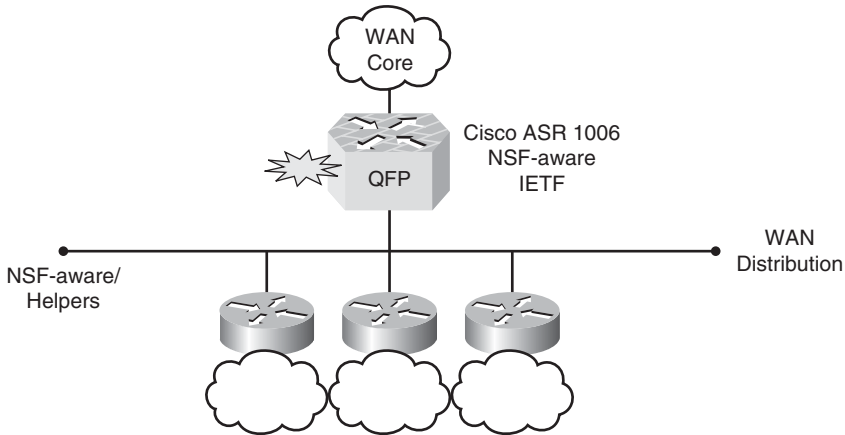


Figure 12-1 Logical view of many regional WAN aggregation routers coming into a consolidated WAN campus edge router.

To turn on IETF helper mode on all the distribution hub routers, including the Cisco ASR 1000, you need to execute the following configuration steps:

Step 1. Configure NSF within the given OSPF process ID:

```
ASR1006# configure terminal
ASR1006(config)# router ospf 100
ASR1006(config-router)# nsf ietf restart-interval 300
```

Note By default, both IETF and Cisco NSF helper modes are turned on.

Step 2. Check that the NSF is turned on, for sure, on the helper router:

```
Router-helper# show ip ospf 100

Routing Process "ospf 100" with ID 172.16.1.2
--output truncated--
IETF Non-Stop Forwarding enabled
  restart-interval limit: 300 sec
IETF NSF helper support enabled
Cisco NSF helper support enabled
Reference bandwidth unit is 100 mbps
Area BACKBONE(0)
ASR1006# sh ip ospf 100
Routing Process "ospf 1" with ID 10.1.1.1
--output truncated--
IETF Non-Stop Forwarding enabled
  restart-interval limit: 300 sec
IETF NSF helper support enabled
Cisco NSF helper support enabled
```

Step 3. Now you need to verify that both RPs are active (using the **show platform** command) and OSPF neighbor relationships are established (using the **show ip ospf neighbors** command):

```
! active ESP:
ASR1006# show platform software ip fp active cef summary
Forwarding Table Summary
Name          VRF id  Table id  Protocol  Prefixes  State
-----
Default       0       0         IPv4      10000     cpp:
                                0x10e265d8
                                (created)

! standby ESP:
ASR1006# show platform software ip fp standby cef summary
Forwarding Table Summary
Name          VRF id  Table id  Protocol  Prefixes  State
-----
Default       0       0         IPv4      10000     cpp: 0x10e265d8
                                (created)
```

You can also view the prefixes downloaded into both the active and standby Embedded Service Processor (ESP) before failing over the router.

The preceding output shows that about 10K routes are created and exist in both ESPs before the failover.

Step 4. Now you'll induce the RP SSO failover (using **redundancy force-switchover**) from the active RP enable mode CLI. The following output shows the effects from the newly active RP:

```
ASR1006# show ip ospf 100
--output truncated--
IETF Non-Stop Forwarding enabled
  restart-interval limit: 300 sec, last IETF NSF restart 00:00:10 ago
IETF NSF helper support enabled
Cisco NSF helper support enabled
```

Note The FIB remains detached from the Routing Information Base (RIB) until the routing protocol reconverges; therefore, both ESPs retain the pre-failover FIB copies until that time. Packet forwarding continues based on the last-known FIB and adjacency entries. The newly active RP sends GR link-state advertisements (LSAs, grace LSAs) to the NSF-aware/helper routers (again, distribution hub routers in this case). With ASR1000-RP1 and ASR1000-ESP10/ESP20, FIB download times are in the neighborhood of 5 k prefixes/sec. It is also worth noting that large number of peers will cause RIB/FIB calculation to put more stress on RP CPU, and similarly large number of prefixes (with or without large number of peers) will cause longer FIB download times. You can tune BGP timers and GR timers for scenarios where you have a large number of prefixes as default values, which might not be the most optimal.

Step 5. RP SSO will not result in any packet loss, because forwarding continues during this entire process. During this switchover process, you can execute the **show platform** command to verify that the former active RP is booting (“booting” state).

In case of ASR1000-ESP10 failover, some small packet loss will occur (packets that are being processed inside the QuantumFlow Processor [QFP]), although that would account for much less than 1-ms worth of transit traffic loss.

NSF/SSO allows RPs to fail over without any packet loss, and ESPs can fail over with extremely small packet loss. The Cisco ASR 1000 shows core benefits of a carrier-class router where failover times beat even the Automatic Protection Switching (APS) gold standard of 50 ms.

In today's networks, where SLAs are enforced and networks are participating in life- and mission-critical scenarios, a robust infrastructure with faster failover based on modern architectures is a must.

Packet Capture Using Encapsulated Remote SPAN

For various reasons, including compliance, enterprises are looking for ways to capture data for further analysis (using an intrusion detection/prevention system [IDS/IPS] or some other advanced analysis system). NetFlow proves handy for this purpose, where you can get detailed IP flow accounting information for the given network.

NetFlow, however useful, still does not provide full packet capture capability from Layer 2 to 7. This is where the Switch Port Analyzer (SPAN) function steps in, although as the name says, this is limited to switches only. SPAN or Remote SPAN (RSPAN), where monitored traffic can traverse a Layer 2 cloud or network, essentially creates an opportunity to capture and analyze traffic on two different switches that are part of a single Layer 2 domain (as opposed to a Layer 3 routing domain). Encapsulated Remote SPAN (ERSPAN), as the name says, brings generic routing encapsulation (GRE) for all captured traffic and allows it to be extended across Layer 3 domains. Until recently, ERSPAN has been available only on Catalyst 6500 and 7600 platforms.

The Cisco ASR 1000 originated with ERSPAN support and can operate in two ways:

- As source or destination for ERSPAN sessions
- As source and destination for ERSPAN sessions at the same time

Note, as well, that this implementation is interoperable with Catalyst 6500 and 7600, and so traffic captured on a port/interface attached to an ASR 1000 can be sent to a destination monitoring station over to a 6500/7600 across a Layer 3 domain as a GRE packet.

Use Case: Ethernet Frame Capture and Transport Across a Layer 3 Cloud

An enterprise has an ASR 1000 being used at one of the regional HQs in San Francisco, and needs to capture traffic from an interface on an on-demand basis and bring it to the centralized data center location in Austin, terminating it on a Catalyst 6500 switch in the core. The San Francisco and Austin locations are connected via a shared MPLS IP VPN cloud.

Note ERSPAN is a Cisco proprietary feature and is available only to Catalyst 6500, 7600, Nexus, and ASR 1000 platforms to date. The ASR 1000 supports ERSPAN source (monitoring) only on Fast Ethernet, Gigabit Ethernet, and port-channel interfaces. The ASR 1000, being a router, does not support regular SPAN or RSPAN functions. For source interface and source VLAN configuration, the default SPAN direction is “both.”

To meet the requirement needed for this enterprise, you need to implement ERSPAN on the ASR 1000 in the SF HQ location as a source session and terminate it at the Catalyst 6500 switch in the core.

Figure 12-2 shows the ERSPAN source (monitored) and destination (monitoring) ports on the ASR 1000 and Catalyst 6500, respectively.

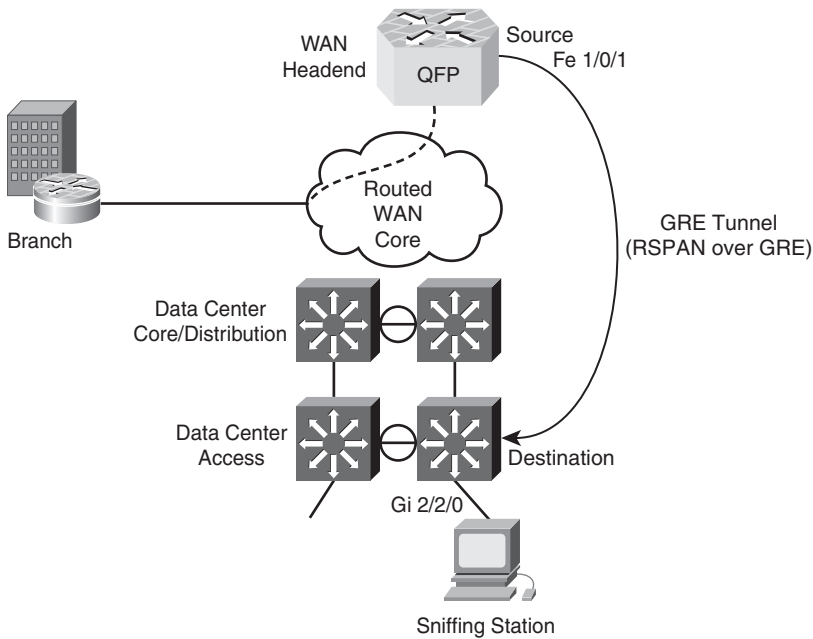


Figure 12-2 Ethernet frame capture at the WAN headend and transporting them to data center via a Layer 3 cloud.

Begin with the configuration on the ASR 1000. Here we'll configure source interface, direction of traffic, and ERSPAN session ID.

Step 1. Identify the ports/interfaces that need to be monitored, and the direction of traffic that needs to be captured, (for example, Rx) by entering the following commands:

```
ASR1006(config)# monitor session 1 type erspan-source
ASR1006(config-mon-erspan-src)# source interface Fe1/0/1 rx
ASR1006(config-mon-erspan-src)# destination
ASR1006(config-mon-erspan-src-dst)# erspan-id 100
ASR1006(config-mon-erspan-src-dst)# ip address 10.10.0.1
ASR1006(config-mon-erspan-src-dst)# ip ttl 32
ASR1006(config-mon-erspan-src-dst)# origin ip address 172.16.0.1
```

Note The ASR 1000 supports up to 1024 sessions that can be source, destination, or a combination per system. This provides tremendous flexibility in data capturing and monitoring to a routing platform.

- Step 2.** Configure the Catalyst 6500 to receive traffic from the source session on the ASR 1000 from Step 1:

```
Cat6500(config)# monitor session 2 type erspan-destination
Cat6500(config-mon-erspan-dst)# destination interface gigabitEthernet
2/2/0
Cat6500(config-mon-erspan-dst)# source
Cat6500(config-mon-erspan-dst-src)# erspan-id 100
Cat6500(config-mon-erspan-dst-src)# ip address 172.16.0.1
```

You can use the **show monitor session** command to verify the configuration:

```
ASR1006# show monitor session 1
Session 1
-----
Type                : ERSPAN Source Session
Status              : Admin Enabled
Source Ports        :
  RX Only           : Fe1/0/1
Destination IP Address : 10.10.0.1
Destination ERSPAN ID  : 100
Origin IP Address    : 172.16.0.1
IP TTL               : 32
```

- Step 3.** To be able to monitor the statistics of monitored traffic, you need to use **show platform hardware qfp active feature erspan state** command:

```
ASR1006# show platform hardware qfp active feature erspan state
ERSPAN State:
  Status      : Active
--output truncated--
System Statistics:
  DROP src session replica :      0 /      0
  DROP term session replica :      0 /      0
  DROP receive malformed   :      0 /      0
  DROP receive invalid ID  :      0 /      0
  DROP recycle queue full  :      0 /      0
  DROP no GPM memory       :      0 /      0
  DROP no channel memory   :      0 /      0
```

This will achieve the purpose of capturing received traffic on the ASR 1000 (FE1/0/1) to Catalyst 6500 GE2/2/0. This traffic will simply be captured, encapsulated in GRE by ASR 1000 natively by the QFP chipset and routed over to the Catalyst 6500. A sniffing station on the 6500 attached to GE2/2/0 will see the complete Ethernet frame (L2 to L7) up to jumbo size (assuming the routed WAN infrastructure can carry jumbo frames end to end).

The ASR 1000, being the first midrange routing platform to support ERSPAN, adds tremendous value to data capturing and data visibility end to end from a branch, or from HQ to data center, a common requirement in medium to large enterprise networks. ERSPAN packet replication is natively done by the QFP chipset, and therefore no external modules are required. ERSPAN, when combined with NetFlow, can result in detailed end-to-end network visibility.

Achieving Segmentation Using MPLS over GRE and MPLS VPNs over GRE Solutions

In today's world, an enterprise campus is home to many different and often competing users. Multitenant environments such as universities, airports, and some public-sector networks (including educational networks) fall under this category.

Such enterprises leverage their high-touch intelligent networking infrastructure to provide connectivity and network services for all stakeholders. For instance, different airlines could share one physical airport network and get billed for this connectivity. This setup accelerates the return on network infrastructure investment, and it optimizes network operations and operational expenses through virtualization. Regulatory compliance, mergers and acquisitions (M&A), and network infrastructure consolidation are among the many drivers. For the users of this single physical network, it results in seamless and instant-on delivery of services, which in turn results in increased revenue streams.

MPLS (or MPLS-based applications) has gained a lot of ground because of its capability to provide this virtualization within a large enterprise network and still provide the much-needed segmentation. The relevant technologies that you hear about are usually MPLS/LDP over GRE, and MPLS VPNs (2547) over GRE, in addition to a host of other MPLS-based technologies.

Use Case: Self-Managed MPLS and Enterprise Private WAN Segmentation

An enterprise is running a “self-managed” or “self-deployed MPLS” core to achieve this network segmentation. Deploying MPLS (or RFC 2547) over a mesh of GRE tunnels (enterprise provider edge [PE] to enterprise PE) allows the enterprise to extend their MPLS network over almost any IP network. Additional benefits include flexibility of edge router roles (provider [P] or PE), independence from the service provider (SP) cloud (which sees those packets as IP packets), and an easier add-on encryption capability, something you can call MPLS over GRE over IPsec. Several large enterprises today are running this environment in their production network.

Configurations of such deployments are fairly straightforward, where WAN edge routers (or customer edges [CE]) basically serve as enterprise Ps or PEs (also referred to as E-Ps or E-PEs), as documented in the text that follows.

Figure 12-3 shows the isolated self-deployed enterprise MPLS clouds that are connected together via an SP MPLS core using LDP over GRE.




Figure 12-3 Enterprise PEs (E-PE) are connected across the enterprise-owned/managed MPLS cloud.

A point-to-point GRE tunnel is set up between each WAN edge router pair if a full mesh is desired. From a control-plane perspective, the following protocols are to be run within the GRE tunnels:

- An IGP such as EIGRP or OSPF for MPLS device reachability. (This makes the E-PE, E-P, and route reflectors [RRs], if configured, reachable to each other.)
- LDP, to allow the formation of LSPs over which traffic is forwarded.
- MP-iBGP for VPN route and label distribution between the E-PE devices.

You will need to configure MPLS labeling, using the `mpls ip` command, on the tunnel interfaces rather than on the WAN edge router physical interfaces. You can verify this configuration with the `show platform software interface` command:

```
E-PE-SF(config)# interface Tunnel10
description GRE tunnel to E-P-NY
bandwidth 10000
ip address 172.16.10.5 255.255.255.0
ip mtu 1400
mpls ip
tunnel source Loopback10
tunnel destination 10.10.10.1
```

```
E-PE-SF# sh platform software interface fp active name Tunnel10
```

```
Name: Tunnel10, ID: 24, CPP ID: 25, Schedules: 0
```

```
--output truncated--
```

```
Flags: ipv4, mpls
```

```
ICMP Flags: unreachable, redirects, no-info-reply, no-mask-reply
```

```
ICMP6 Flags: unreachable, redirects
```

```
Dirty: unknown
```

```
AOM dependency sanity check: PASS
```

```
AOM Obj ID: 1081
```

Figure 12-4 shows the end-to-end protocol stacks for an MPLS/LDP over GRE scenario.

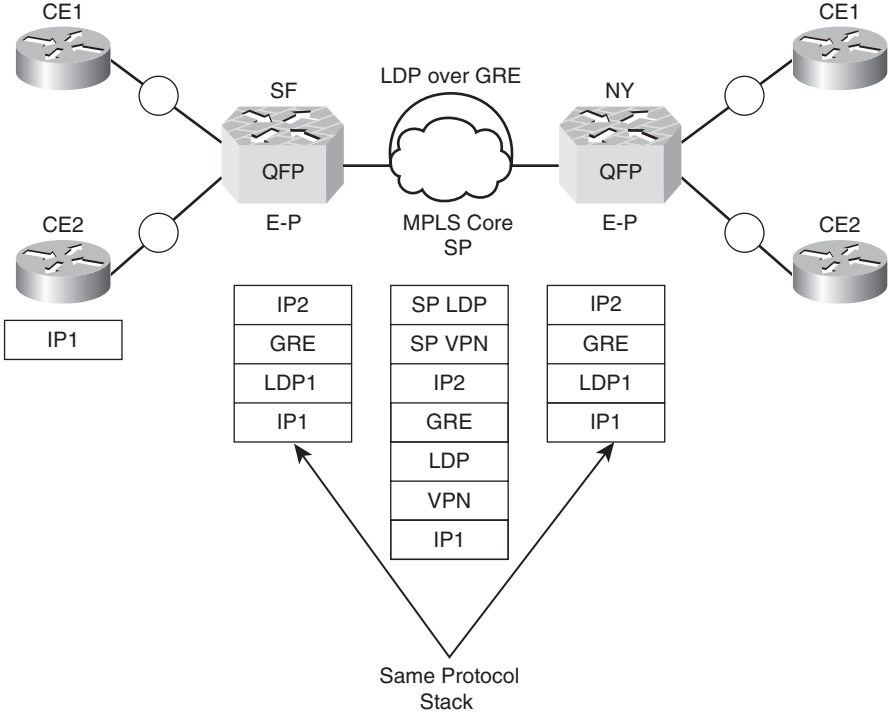


Figure 12-4 Protocol stacks for packets at both P and in the MPLS cloud.

This will effectively create an LSP from E-P-SF to E-P-NY, and the intermediary SP cloud does not have to be an MPLS-based service.

Note There is a subtle but critical difference between the enterprise WAN edge routers being configured as P or PE. In the former case, there will be an additional label and LSP, whereas in the latter case there will be no LSP and no additional label because of penultimate hop pop (PHP) behavior. The only use of the labeling in case of E-PE enterprise segmentation is to map Virtual Routing and Forwarding (VRF) instances (as in per customer or network) within the same GRE tunnel and still be able to multiplex and demultiplex them to correct customer VRF instances.

Cisco ASR 1000, starting from IOS XE 2.2, supports both MPLS/LDP over GRE and MPLS VPNs (2547) over GRE with and without encryption.

Figure 12-5 shows the end-to-end protocol stacks for an MPLS VPNs over GRE scenario, or something also known as 2547 VPNs over GRE.

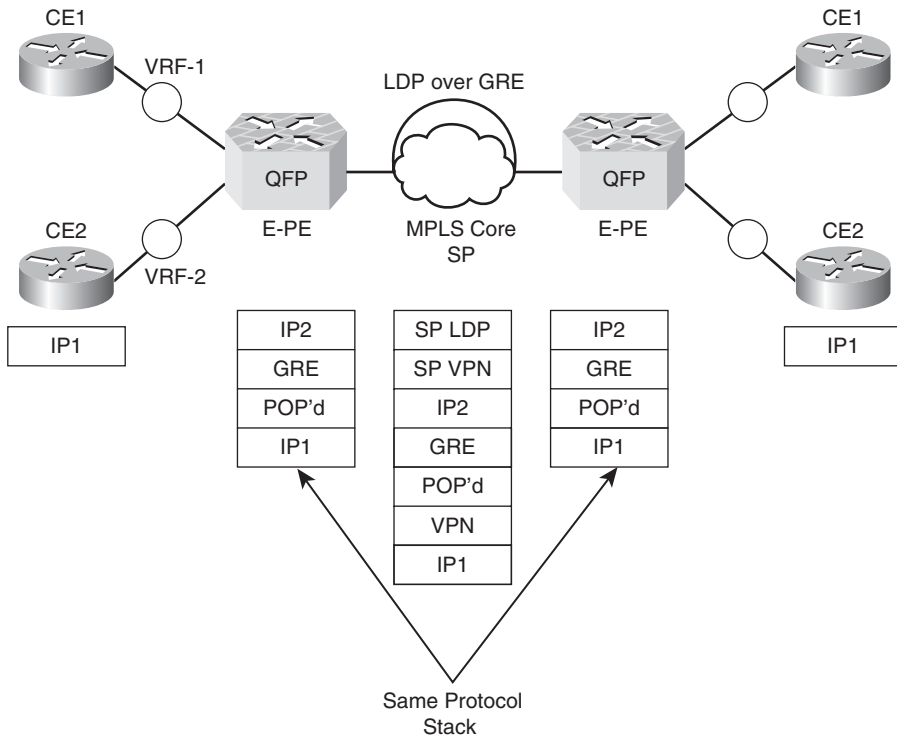


Figure 12-5 Protocol stacks at both PEs and in the MPLS cloud.

Full-mesh peer-to-peer (p2p) GRE tunnels can easily become an administrative hassle in a network with large number of WAN edge routers. In those cases, enterprises can also consider 2547 over Dynamic Multipoint VPN (DM VPN), or 2547 over mGRE over IPsec, to ease the burden of tunnel administration. These solutions will be supported on ASR 1000 in the future IOS XE versions.

The Cisco ASR 1000 provides the extreme flexibility necessary to meet the changing business environments that need virtualization in today's multitenant enterprise networks by supporting MPLS/2547 over GRE solutions at serial interface, Fast Ethernet, Gigabit Ethernet, or even 10 Gigabit Ethernet speeds natively or higher with the unique capability to perform all these encapsulations inside the single QFP chipset.

Scalable v4/VPNv4 Route Reflector

With the growing adoption of MPLS in the enterprises to achieve large-scale virtualization and segmentation, there is also a need for enterprises to have their own route reflector (RR) for VPNv4 routes, deployed separately or combined in a PE router. An RR simplifies the iBGP full-mesh restriction where all PEs don't have to mesh with all other PEs, rather just with the RR.

Use Case: Route Reflection

Figure 12-6 shows the RR used by the enterprise in the self-managed MPLS clouds.

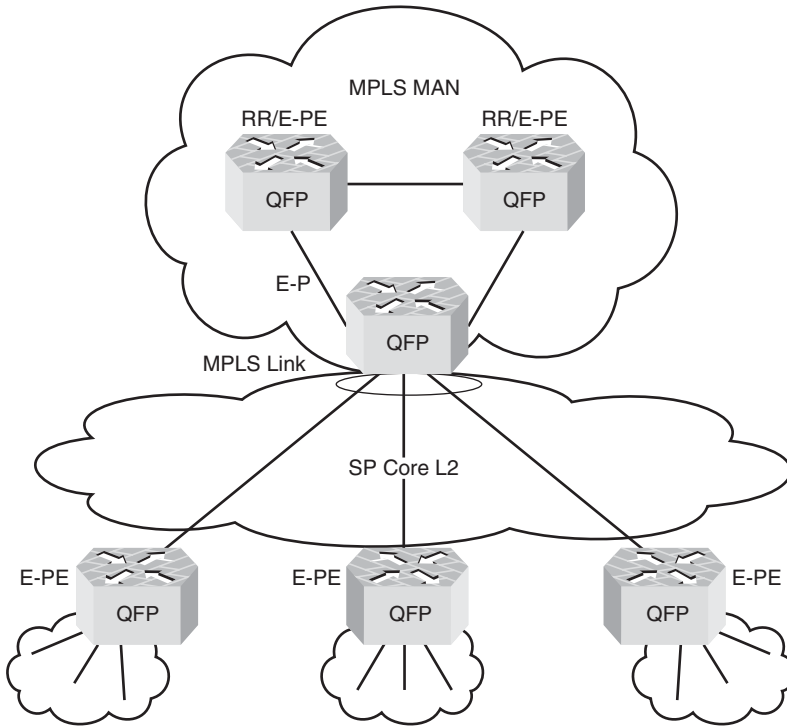


Figure 12-6 MAN using the same router for E-PE and VPNv4 RR roles.

To meet this requirement of avoiding the full mesh of iBGP, you need to configure the Cisco ASR 1000 as the RR for VPNv4 routes using the following steps:

- Step 1.** Configure RRs to peer with PEs to reflect VPNv4 routing information learned from other PEs:

```
ASR1004-RR(config)# router bgp 100
ASR1004-RR(config-router)# neighbor A-PE peer-group
ASR1004-RR(config)# neighbor A-PE remote-as 100
ASR1004-RR(config)# neighbor A-PE update-source Loopback100
ASR1004-RR(config)# neighbor PE loopback# peer-group A-PE
```

- Step 2.** Configure RRs for VPNv4 BGP peering between PEs and RRs:

```
ASR1004-RR(config-router)# address-family vpnv4
ASR1004-RR(config-router-af)# neighbor A-PE activate
ASR1004-RR(config-router-af)# neighbor A-PE route-reflector-client
```

```
ASR1004-RR(config-router-af)# neighbor A-PE send-community extended
ASR1004-RR(config-router-af)# neighbor PE loopback# peer-group A-PE
```

Step 3. Configure the PE for VPNv4 BGP peering between PEs and RRs (thus enabling PEs to exchange VPNv4 routing information with the RRs):

```
ASR1004-PE(config)# router bgp 100
ASR1004-PE(config-router)# no synchronization
ASR1004-PE(config-router)# bgp log-neighbor-changes
ASR1004-PE(config-router)# neighbor ASR1004-RR loopback ip# remote-as
100
ASR1004-PE(config-router)# neighbor ASR1004-RR loopback ip# update-
source Loopback0
ASR1004-PE(config-router)# address-family vpnv4
ASR1004-PE(config-router-af)# neighbor ASR1004-RR loopback ip# activate
ASR1004-PE(config-router-af)# neighbor 172.16.1.1 send-community
extended
```

Note It is assumed that LDP is enabled on the core routers, RR, and PE core-facing interfaces. When RRs are in use, and all the outgoing updates have the same policy, it makes sense to use peer groups on the RRs because this reduces the number of outgoing updates (per client) that an RR router has to generate.

Although this example uses the Cisco ASR 1004 as the VPNv4 RR, this is applicable to the IPv4 RR, too. The VPNv4 route scale is completely a function of the ASR1000-RP you have in the system. With the ASR1000-RP1 and ASR1000-RP2, the scale is up to 1M and 4M, respectively, for IPv4. For VPNv4 routes, ESP does not have to be in the data path, and therefore any ESP can be used. Currently for IPv4, FIB entries are still populated, hence limiting the RR scale. This will change in a future IOS XE version.

The Cisco ASR 1000, by virtue of the ASR1000-RP1 and ASR1000-RP2, provides the largest scale for Route Reflector deployments in the Cisco midrange routing portfolio. The ASR1000-RP2, with 16-GB DRAM, truly raises the bar, with 64-bit IOS XE that allows the routes to scale up to 20M, which essentially rivals even the largest core routers available today.

In general, the ASR1000-RP2 (16-GB DRAM) provides four times the route scale over RP1 (4-GB DRAM), three times the number of peers/sessions (with the given convergence time) and is at least twice as fast in terms of route convergence (for the given set of routes and peers).

Scalable and Flexible Internet Edge

When we talk about a router to be placed at the edge of the network facing the public Internet, a few things come to mind. An ideal router needs to be flexible and scalable with regard to features and variety of interfaces, without requiring service modules for every basic service, such as Network Based Application Recognition (NBAR), Flexible Packet Matching (FPM), firewalls, and IPsec. Other critical attributes include high availability, deep packet inspection, and near-line-rate quality of service (QoS).

High availability enables applications to remain available in case of software or hardware failure that causes a data- or control-plane problem. Deep packet inspection helps classify the data based on application header or payload; it also addresses zero-day attacks.

Use Case: Internet Gateway/Edge Router

An enterprise is looking for, in a smaller-compact factor, an Internet edge that can natively accelerate NAT, firewall, NetFlow, and access control lists (ACL), along with ISSU and RP SSO. This device should also be able to scale up to 10 Gbps if needed in the future.

To meet these requirements, you could use the ASR 1002 with ASR1000-ESP5, which provides 5-Gbps system bandwidth with four built-in Gigabit Ethernet ports ready to be used as fiber or copper and facing either the inside LAN or Internet (usually provisioned via an Ethernet link).

The ASR 1002 can also take the ASR1000-ESP10, which satisfies the requirements of 10 Gbps, essentially doubling the bandwidth from initial deployment.

Figure 12-7 shows the ASR 1002/ASR1000-ESP5 deployed at the Internet edge.

There are no configurations to be shared in this use case, but note the performance and scale numbers for the ASR 1000 series routers relevant to the previously mentioned features.

Table 12-2 shows the various features and their respective performance and scale relevant to Internet edge.

Note IOSD failover, IOS Firewall, and IPsec require their respective right to use (RTU) licenses. (At the time of this writing, however, these are only honor-based paper licenses and, as such, not enforced via the CLI.)

Tunnels per second (TPS), as mentioned in Table 12-2, is basically a function of RP compute cycles, because all Internet Key Exchange (IKE) packets are sent to the RP. In some environments, such as remote-access VPN aggregation/head-end where tunnel churn is expected, the ASR1000-RP2 is an option to get higher scalability.

The Cisco ASR 1000 not only meets the typical Internet gateway router requirements here, but also exceeds them from both control- and data-plane perspectives. The capability

to have two IOS daemons running at the same time, and providing IOSD-based SSO, is second to none!

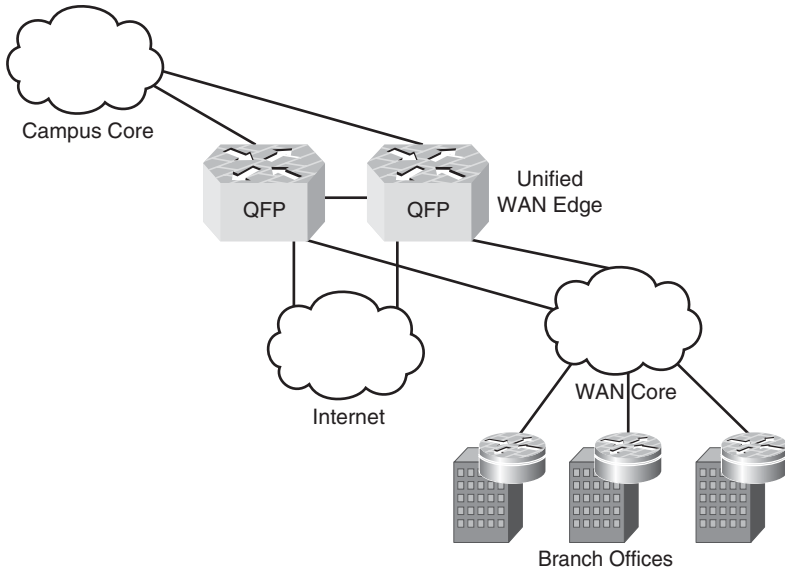


Figure 12-7 Single router used for both the WAN edge and Internet gateway router.

Table 12-2 Various ESPs and Their Scale and Performance for IOS Zone-Based Firewall, NetFlow, and IPsec

Feature	ASR1000-ESP5	ASR1000-ESP10	ASR1000-ESP20
IOS zone-based firewall (L4 inspection)	5 Gbps	10 Gbps	20 Gbps
NetFlow (v5, v8, v9)	500K flow cache entries	1M flow cache entries	2M flow cache entries
IPsec	1 Gbps at IMIX 4000 tunnels 90 tunnels/sec with ASR1000-RP1	2.5 Gbps at IMIX 4000 tunnels 90 tunnels/sec with ASR1000-RP1	5.2 Gbps at IMIX 4000 tunnels 90 tunnels/sec with ASR1000-RP1
Dual IOSD failover	< 50 ms	< 50 ms for ASR 1002-F/ASR 1002/ASR 1004 chassis	< 50 ms for ASR 1002-F/ASR 1002/ASR 1004 chassis

Scalable Data Center Interconnect

Today's businesses are seeing more and more consolidation for both file and application servers into a small number of data centers. Major drivers for this trend include cost savings, regulatory compliance, and ease of backup and administration.

At the heart of this, there is also a virtualization trend, where compute cycles are being isolated or abstracted from storage. This has created newer technologies for virtual machine high availability, and migration such as VMotion, clustering, or even geo-clustering of servers, which require extending Layer 2 VLANs across the WAN (data center interconnect).

Now, when looking at the data center connection and trying to tie it up with the application vendor requirements, almost all suggest using Layer 2 adjacent servers. To satisfy or emulate the requirement of L2 adjacencies across the WAN, various requirements emerge from these trends:

- **Loop prevention:** This refers to isolation of Spanning Tree Protocol (STP) to each data center itself, and not extended across the data center interconnect (DCI).
- **Redundancy:** This refers to the DCI solution itself not being prone to node or link failures. That, of course, requires redundancy.
- **Convergence times:** Apparently, there is no set standard for this requirement for DCI. It really depends on what applications are being run (for example, a requirement driven by VMotion stipulates no more than a couple of seconds for convergence).
- **Usage of multiple paths:** This is where technologies such as Virtual Switching Systems (VSS) and Multichassis EtherChannel (MEC) come into play. There is another similar solution known as virtual port channel (vPC), which essentially allows creating an EtherChannel where member links are across two different physical systems.

Note The Nexus 7000 supports vPC today. VSS/MEC are supported in the 6500 beginning with 12.2(33)SH1 code. Both technologies are slightly different but are beyond the scope of this book. TrustSec (802.1AE) is done using the port ASICs, and hence there is no degradation to the forwarding while encryption is added to the mix on Nexus 7K.

Interested readers should point to the links provided in the "Further Reading" section at the end of the chapter for more information.

Three types of transport are common for DCI:

- **Dark fiber:** Fiber that is not lit yet is called dark fiber. Not very many organizations have access to dark fiber, but the ones who have see it as the most preferred way of doing DCI. This is usually limited in distance.
- **IP:** This is a rather common medium and usually consists of some kind of private IP services that most SPs offer across geographies.
- **MPLS:** This is one of the more common ways to connect data centers.

The ASR 1000 supports almost all forms of Gigabit Ethernet coarse/dense wavelength-division multiplexing (CWDM/DWDM) optics, although the Catalyst 6500 with VSS/MEC has a solution that meets all the requirements in this arena, including multisite DC connectivity.

For IP and MPLS, the ASR 1000 offers (complementing the Cisco 6500 solution) Ethernet over MPLS and Ethernet over MPLS over GRE, starting from IOS XE 2.4 for dual-site DCI.

Figure 12-8 shows the MPLS transport and active/active EoMPLS pseudowires across DCI routers.

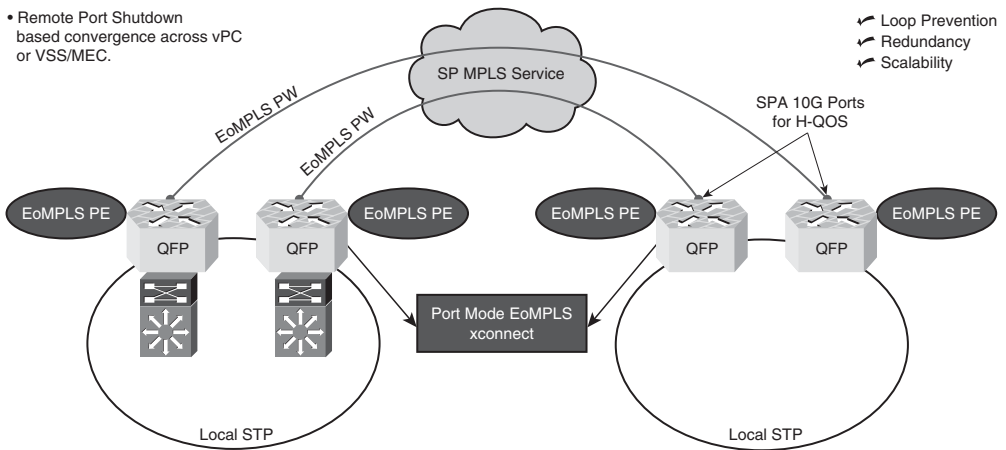


Figure 12-8 *Encrypting Ethernet frames at Layer 2 using TrustSec and avoiding the use Layer 3 encryption such as IPsec.*

Figure 12-9 shows the MPLS transport and active/active EoMPLS pseudowires across DCI routers. Here the DC core switches are Nexus 7Ks running TrustSec to encrypt packets at Layer 2 hop by hop.

The solution in Figure 12-9 shows a unique advantage where any traffic leaving the premise is required to be encrypted, as this provides a native way to encrypt all traffic. This requirement is common in government and state agencies.

Note The Nexus 7000 switch has always supported TrustSec. Interested readers should point to links in the “Further Reading” section at the end of the chapter for more information.

Figure 12-10 shows the IP transport and active/active EoMPLSoGRE tunnels across DCI routers.

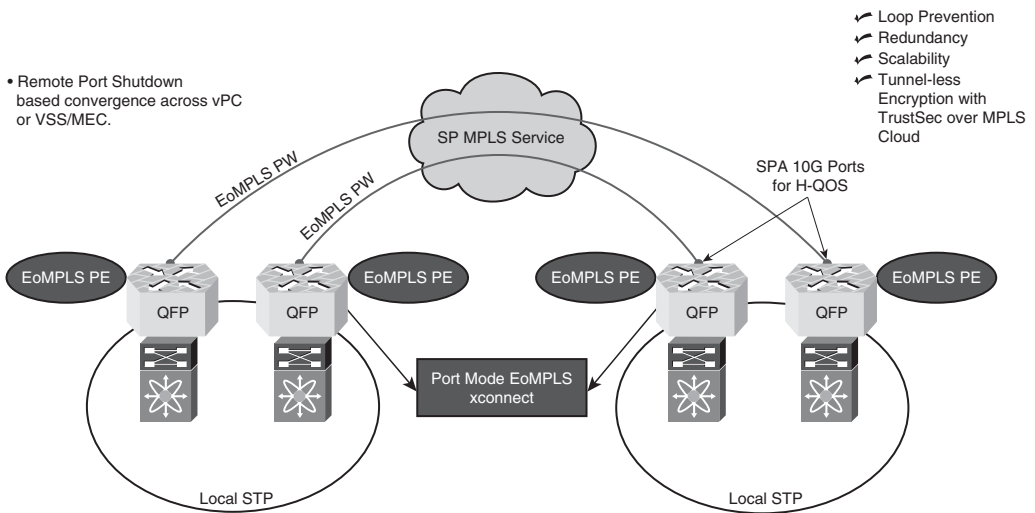


Figure 12-9 EoMPLS scenario where the transport cloud is MPLS.




Figure 12-10 EoMPLSoGRE scenario where the transport cloud is IP.

This can also be seen as a consolidation strategy, especially for green-field deployments, where ASR 1000 working as a DCI LAN extension router can also serve as a consolidated unified WAN services router. This brings down the TCO much lower and at the same time allows for faster qualification, where the ASR 1000 functions as a private WAN aggregation, and even perhaps the Internet edge can be collapsed at the consolidated WAN edge.

Figure 12-11 shows the unified WAN edge, which consolidates the DCI with multiple other functions.

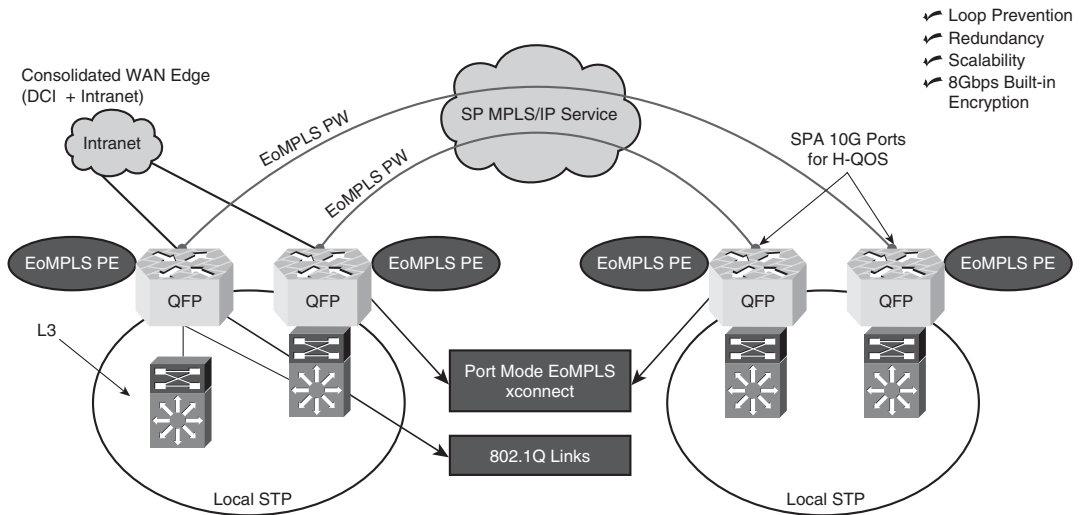


Figure 12-11 Example of DCI and WAN edge functional collapse in a single router.

Use Case: Encrypting Traffic over an EoMPLS Pseudowire at Layer 2 Using TrustSec

Assume that an organization wants to connect two different data centers and extend multiple VLANs across these two sites for connecting various clusters, geo-clusters. Assume that the transport method in the middle is MPLS and the organization is using Nexus 7Ks as data center core switches. The customer also wants to start with a clear-text (nonencrypted) DCI, but later on would also like to add encryption to it to deal with regulatory compliance. This must be met without much configuration overhead.

Now, to meet these requirements, you need to extend the Layer 2 connectivity across the sites. Because the transport medium is MPLS, you can start with clear-text requirement and deploy the ASR 1000, as shown in Figure 12-8, and then later move to the deployment illustrated in Figure 12-9 by turning on TrustSec (or 802.1AE) link-layer encryption on the Nexus transparently without changing anything on the ASR 1000!

The examples that follow examine the configuration for both the ASR 1000 (port mode EoMPLS) and TrustSec on the Nexus 7000.

Figure 12-12 shows the final topology with both Nexus 7K and ASR 1000 connected with vPC.

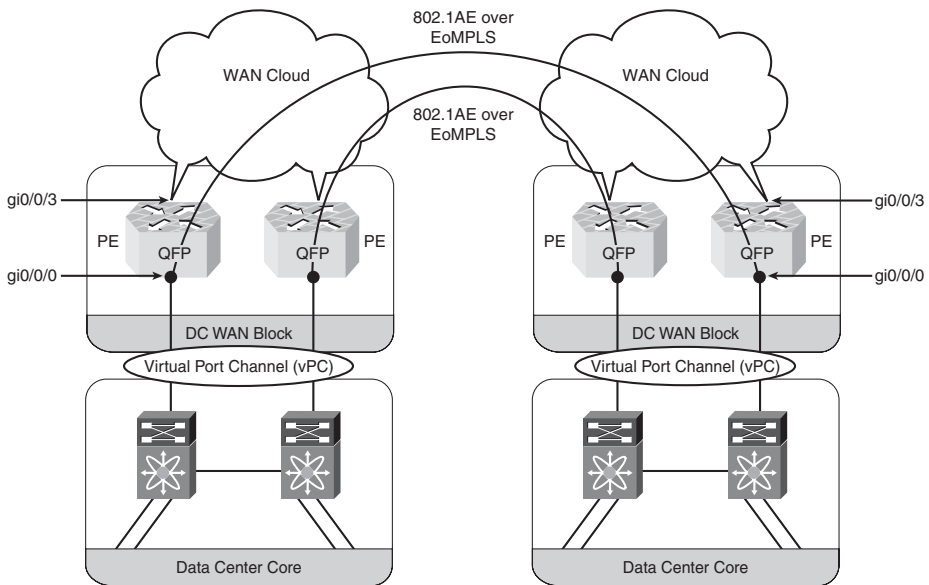


Figure 12-12 Nexus 7K and ASR 1000 connected via vPC.

Example 12-1 shows the ASR 1000 port mode EoMPLS configuration, which can be used to theoretically extend 4K VLANs.

Example 12-1 Port Mode EoMPLS Configuration for the ASR 1000 with Remote Port Shutdown Enabled by Default

```
ASR1000-1:
ASR1000-1(config)# interface Loopback0
ASR1000-1(config-if)# ip address 192.168.100.1 255.255.255.255
!
ASR1000-1(config)# interface GigabitEthernet0/0/0
ASR1000-1(config-if)# mtu 9216
ASR1000-1(config-if)# no ip address
ASR1000-1(config-if)# negotiation auto
ASR1000-1(config-if)# xconnect 192.168.100.2 100 encapsulation mpls
!
ASR1000-1(config)# interface GigabitEthernet0/0/1
ASR1000-1(config-if)# description to ASR-2
ASR1000-1(config-if)# mtu 9216
ASR1000-1(config-if)# ip address 10.1.2.1 255.255.255.0
ASR1000-1(config-if)# load-interval 30
```

```

ASR1000-1(config-if)# negotiation auto
ASR1000-1(config-if)# mpls label protocol ldp
ASR1000-1(config-if)# mpls ip
!
!
ASR1000-1(config)# interface GigabitEthernet0/0/3
ASR1000-1(config-if)# description to ASR-2
ASR1000-1(config-if)# mtu 9216
ASR1000-1(config-if)# ip address 10.1.1.1 255.255.255.0
ASR1000-1(config-if)# load-interval 30
ASR1000-1(config-if)# negotiation auto
ASR1000-1(config-if)# mpls label protocol ldp
ASR1000-1(config-if)# mpls ip
ASR1000-2:
ASR1000-2(config)# interface Loopback0
ASR1000-2(config-if)# ip address 192.168.100.2 255.255.255.255
!
ASR1000-2(config)# interface GigabitEthernet0/0/0
ASR1000-2(config-if)# mtu 9216
ASR1000-2(config-if)# no ip address
ASR1000-2(config-if)# negotiation auto
ASR1000-2(config-if)# xconnect 192.168.100.1 100 encapsulation mpls
!
ASR1000-2(config)# interface GigabitEthernet0/0/1
ASR1000-2(config-if)# description to ASR-1
ASR1000-2(config-if)# mtu 9216
ASR1000-2(config-if)# ip address 10.1.2.2 255.255.255.0
ASR1000-2(config-if)# load-interval 30
ASR1000-2(config-if)# mpls label protocol ldp
ASR1000-2(config-if)# mpls ip
!
!
ASR1000-2(config)# interface GigabitEthernet0/0/3
ASR1000-2(config-if)# description to ASR-1
ASR1000-2(config-if)# mtu 9216
ASR1000-2(config-if)# ip address 10.1.1.2 255.255.255.0
ASR1000-2(config-if)# mpls label protocol ldp
ASR1000-2(config-if)# mpls ip

```

Example 12-2 shows the Nexus 7000 configuration to use TrustSec for all traffic going outbound on to the EoMPLS pseudowires (over an MPLS cloud).

Example 12-2 *Nexus 7K TrustSec Configuration*

```

Nexus-7K-1# sh run cts
version 4.1(2)
feature dot1x
feature cts
cts device-id Nexus-7K-1 password 7 qxz12345

interface Ethernet1/12
    switchport
    switchport access vlan 666
    cts manual
        sap pmk abcdef123400000000000000000000000000000000000000000000000000000000000000000000000000
    mtu 9216
    no shutdown

interface Vlan666
    no shutdown
    ip address 155.5.5.1/24

```

```

Nexus-7K-2# sh run cts
version 4.1(2)
feature dot1x
feature cts
cts device-id Nexus-7K-2 password 7 qxz12345

interface Ethernet1/12
    switchport
    switchport access vlan 666
    cts manual
        sap pmk abcdef123400000000000000000000000000000000000000000000000000000000000000000000000000
    mtu 9216
    no shutdown

interface Vlan666
    no shutdown
    mtu 9216
    ip address 155.5.5.2/24

```

Example 12-3 shows that the TrustSec session is established.

Example 12-3 *Confirmation That TrustSec Is Negotiated and Is Up*

```

Operational Status (TrustSec 802.1AE SAP negotiation successful):
Nexus-7K-1# show cts interface e 1/12
CTS Information for Interface Ethernet1/12:
    CTS is enabled, mode:     CTS_MODE_MANUAL
    IFC state:                CTS_IFC_ST_CTS_OPEN_STATE
    Authentication Status:    CTS_AUTHC_SKIPPED_CONFIG

```

```

Peer Identity:
Peer is:                Not CTS Capable
802.1X role:           CTS_ROLE_UNKNOWN
Last Re-Authentication:
Authorization Status:   CTS_AUTHZ_SKIPPED_CONFIG
PEER SGT:              0
Peer SGT assignment:   Not Trusted
Global policy fallback access list:
SAP Status:            CTS_SAP_SUCCESS
Configured pairwise ciphers: GCM_ENCRYPT
Replay protection: Disabled
Replay protection mode: Strict
Selected cipher: GCM_ENCRYPT
Current receive SPI: sci:1b54c148d80000 an:2
Current transmit SPI: sci:225577968c0000 an:2
Operational Status (TrustSec 802.1AE SAP negotiation successful):
-----
Nexus-7K-2# show cts interface e 1/12
CTS Information for Interface Ethernet1/12:
CTS is enabled, mode:   CTS_MODE_MANUAL
IFC state:              CTS_IFC_ST_CTS_OPEN_STATE
Authentication Status: CTS_AUTHC_SKIPPED_CONFIG
Peer Identity:
Peer is:                Not CTS Capable
802.1X role:           CTS_ROLE_UNKNOWN
Last Re-Authentication:
Authorization Status:   CTS_AUTHZ_SKIPPED_CONFIG
PEER SGT:              0
Peer SGT assignment:   Not Trusted
Global policy fallback access list:
SAP Status:            CTS_SAP_SUCCESS
Configured pairwise ciphers: GCM_ENCRYPT
Replay protection: Disabled
Replay protection mode: Strict
Selected cipher: GCM_ENCRYPT
Current receive SPI: sci:225577968c0000 an:2
Current transmit SPI: sci:1b54c148d80000 an:2

```

As shown in Figure 12-12, the two ASR 1000s are connected via two active/active EoMPLS pseudowires. To deal with a failure scenario, the ASR 1000 uses a feature called Remote Pseudo Wire Shutdown. The behavior on the ASR 1000 is a bit different than on the Catalyst 6500/7600, where the feature does not depend on interworking with the Ethernet LMIs.

On the ASR 1000, this feature, upon pseudowire down state, shuts down the local laser on the port with “xconnect”, **gi0/0/0**, as shown in the use case. This is seen by the peer

Ethernet port as the interface going down, and it will go to down/down. This allows the downstream devices to stop sending traffic to the port and results in almost instant convergence. EoMPLS remote port shutdown provides faster failover times for both local/remote node or link failure scenarios.

This behavior is very helpful in the vPC scenario, because it will trigger the LACP (Link Aggregation Control Protocol) to converge instantly and will remove the member link from the virtual port channel.

Summary

This chapter discussed six use cases for Cisco ASR 1000 to provide the following variety of solutions:

- High availability using NSF/SSO in an enterprise
- Data capture using ERSPAN in a router
- MPLS over x solutions in a large enterprise that needs virtualization/segmentation at 10 Gbps or higher speeds
- VPNv4 RR in a self-deployed MPLS enterprise
- Highly available Internet gateway router
- DCI WAN router

The goal was really to go over a diverse set of technology problem statements and solutions that are common in an enterprise and to cover how the ASR 1000 addresses them.

Chapter Review Questions

1. Is NSF for IGP enabled by default?
2. What is the difference between an NSF-aware and NSF-capable router?
3. What is ERSPAN, and which Cisco platforms support ERSPAN today?
4. How does IOS, being a 32-bit OS, address 16-GB DRAM in ASR1000-RP2 to achieve such a high route scale?
5. Does the ASR 1000 require a feature license to turn on and use MPLS, BGP, NAT, GRE, or NetFlow?
6. What is DCI, and how does the pseudowire failover work for remote node/link failure?

Answers

1. No, it is not turned on by default. You need to turn it on by entering the `nsf` command within IGP configuration mode.
2. NSF-aware means that the device can participate in an NSF restart by virtue of understanding the GR LSA, but might not undertake the restart itself. NSF-capable

routers, on the other hand, can both understand GR LSA and can also undergo an NSF restart. Cisco ASR 1000 is an NSF-aware and -capable device.

3. ERSPAN stands for Encapsulated Remote SPAN, which essentially encapsulates the SPAN-ed traffic inside a GRE header so that it can be routed across a Layer 3 domain. This enables data capturing on one device on a given set of interfaces and direction, whereas monitoring station could be placed several L3 hops away on another device (such as Cisco ASR 1000). Cisco Catalyst 6500, 7600, Nexus, and ASR 1000 are the only platforms that support ERSPAN.
4. The RP IOS package for ASR1000-RP2 and most of the underlying software infrastructure has been extended to 64 bits, hence it can therefore address beyond 4 GB DRAM.
5. No, the Cisco ASR 1000 does not require any software RTU licenses for these basic features. Hence, they can be used as long as they are available in the given IOS image.
6. DCI stands for data center interconnect, which is a common way to extend Layer 2 or Layer 3 connectivity across the data centers. The ASR 1000 can be used at this time for p2p connectivity across two data centers for IP and MPLS transport types. EoMPLS can be used to extend the L2 connectivity and VLANs across the DCI WAN link. The ASR 1000 has a unique feature known as Remote Port Shutdown, which functions similar to GSR. So, to avoid traffic blackholing and to allow faster convergence, as soon as a pseudowire goes down, the router switches off the port laser to let the peer port (customer edge [CE]) know that the link has gone down, which immediately goes to down/down. This proves handy to achieve extremely fast convergence end to end. As soon as the pseudowire comes back up, it turns the laser on, signaling the CE port that it can resume traffic via the given EoMPLS PE. This feature is enabled by default and does not require any additional configuration.

Further Reading

Graceful OSPF Restart, document: <http://tools.ietf.org/html/rfc3623>

Configuring ERSPAN on Catalyst 6500 Switches, document: <http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/span.html#wp1063324>

Internet Gateway Router Design Using Cisco ASR 1000 Series Routers, document: <http://tinyurl.com/l6nbcP>

Cisco 6500 Virtual Switching Systems (VSS), document: <http://tinyurl.com/5zph8e>

Configuring vPC (Virtual Port Channel), document: <http://tinyurl.com/l37wqp>

Cisco Nexus 7000 Security Features, document: <http://tinyurl.com/n2nx99>

Data Center Interconnect, document: <http://tinyurl.com/rc1v2f>

“Route Reflector Scale,” report: <http://tinyurl.com/kmc89b>

Index

Symbols

7200VXR (QoS scaling guidelines), 216

A

active ESP, 220

active RP, 220

adoption, time to (WAN architectures), 24, 26

aggregating

 policies, 215

 SIP traffic, ESP Interconnect Scheduler, 213

application mobility, virtualization and, 20

arrival processing (transit packets), ASR 1000 series routers, 88–90

ASR 1000 GUI (Graphical User Interface)

 configuring, 142

 usage examples, 143, 146

 views of, 141–142

ASR 1000 series routers, 43–44

 availability, 46

 bootflash, ISSU, 117

 booting, 103

 ROMMON, 104–106

valid configurations, 103

 branch aggregation, 50

 BRAS, 53–54

 carrier-class routing, 57

BITS reference clocks, 60

chassis design/modularity, 57

integrated QoS, 59

 ISSU, 58

 LAN interfaces, 60

nonstop router management, 60

operating system modularity, 57

oversubscription, 59

plane separation, 58–59

 WAN interfaces, 60

 configuring, 107, 109

 control plane, 177–178

 CPE, 54

 CWDM/DWDM, 196

- DBE, 274
- embedded services, 56
- ERSPAN support, 184
 - Ethernet frame capture/transport across Layer 3 clouds, 184–186*
- ESP, 79, 103, 106
 - chassis manager, 87*
 - crypto engines, 82–83*
 - displaying insertion/uptime, 166–167*
 - DRAM, 79*
 - ESP subpackages, 135–136*
 - forwarding manager, 86*
 - initializing, 81*
 - interconnect ASIC, 79*
 - packet handling, 81*
 - QFP software, 86*
 - system board, 80*
 - TCAM, 80*
- file system structure, 109–110
- hardware components
 - chassis options, 61*
 - chassis slot naming/ numbering, 62*
 - ESP, 63*
 - RP, 62*
 - SIP, 63*
- hidden costs of, 49
- Internet gateway/edge routers, 193
- Internet gateways, 49–50
- IOS XE software, 98
 - releasing, 99*
- IP services case studies, 205
 - high-speed logging via NetFlow, 223–225*
 - scalable hierarchical QoS, 216–217, 219*
 - scalable IPv4/IPv6 multicast acceleration via QFP, 219–221*
 - scalable multigigabit NAT, 221–223*
 - scalable multigigabit NBAR/FPM, 225, 227–228*
 - scalable QoS via QFP Traffic Manager, 206–215*
- ISSU, 114, 117
 - consolidated packages on fully redundant 6RU, 117*
 - impact on subpackages, 114–115*
 - ISSU on 6RU system for IOSD, 118–123*
 - running dual IOSD on 2/4RU systems, 137–138*
 - subpackage ISSU on fully redundant 6RU, 124–131*
 - upgrading ESP subpackages, 135–136*
 - upgrading RP subpackages, 132–134*
 - upgrading SIP/SPA subpackages, 131*
- life span of, 49
- modular operation, 87
- multicast HA, 220–221
- multicast replication on ESP, 221
- operational savings, 49
- packet handling
 - arrival processing, 88–90*
 - egress processing, 91*
 - ingress processing, 88*
- partitioning, 68
 - data plane, 68*
 - I/O plane, 68*
 - routing plane, 68*

- PQ
 - conditional policing*, 214
 - unconditional policing*, 214
- QFP, 56, 106–107
 - displaying QFP PPE utilization information*, 167–168
- QoS, 48, 207
- QoS scaling guidelines, 216
- reliability, 46–47
- remote user aggregation, 50
- ROI, 49
- RP, 71, 103
 - bootflash*, 73
 - chassis manager*, 85
 - CPU tasks*, 85
 - displaying insertion/uptime*, 166–167
 - DRAM*, 73
 - forwarding manager*, 86
 - front panel*, 72
 - hardware-assisted control-plane protection*, 78
 - HDD*, 74
 - initializing*, 75
 - interconnect ASIC*, 74
 - interface manager*, 86
 - legacy protocol traffic*, 78
 - packet handling*, 75, 77
 - RP subpackages*, 132–134
- RR, 191–192
- SBC, 271
- SBE, 273, 277–278
- scalability, 48
- security headends, 50
- security services case studies, 231
 - Cisco Self-Defending Network schema*, 0, 231
 - DMVPN hub design*, 239–241
 - GETVPN*, 242, 244–248, 250
 - IKE initiation requests*, 235
 - integrated threat control solutions*, 251, 253–256
 - IOS firewalls*, 251, 253–256
 - IPsec HA*, 236
 - IPsec multicast encryption*, 236–237
 - IPsec packet flow (egress)*, 235
 - IPsec packet flow (ingress)*, 235
 - IPsec-based VPN solutions*, 232, 234
 - QoS scalable encryption*, 237, 239
- segmentation, MPLS over GRE, 187, 189
- SIP, 83–84, 103, 106
 - chassis manager*, 87
 - displaying insertion/uptime*, 166–167
 - displaying SPA status in SIP*, 163
 - interface manager*, 87
 - SIP subpackages*, 131
 - SPA drivers*, 87
- software components
 - ESP software*, 65
 - IOS XE*, 63
 - IOSD*, 67
 - Linux kernels*, 67
 - ROMMON*, 68
 - RP software*, 63, 65
 - SIP software*, 66
- SPA
 - displaying SPA status in SIP*, 163
 - displaying SPA-level statistics*, 164
 - SPA subpackages*, 131

subpackages

ISSU impact on subpackages,
114–115

*subpackage ISSU on fully
redundant 6RU, 124–131*

upgrading ESP subpackages,
135–136

upgrading RP subpackages,
132–134

*upgrading SIP/SPA
subpackages, 131*

system redundancy/modularity, 68

traffic encryption over EoMPLS
Pseudowire at Layer 2 via
TrustSec, 198, 200–203

traffic manager priority queues,
213–215

troubleshooting

debug commands, 150–153,
168–174

displaying drop statistics,
164–165

*displaying front-panel LED
status via show platform
hardware command, 163*

*displaying interface-level
feature binding, 165*

*displaying IPv4-related drops
for active QFP, 155–156*

displaying processors,
154–155

*displaying QFP memory
statistics for IRAM/DRAM/
SRAM usage, 156–157*

*displaying QFP memory
statistics on per-IOs
feature/internal-usage basis,
157–161*

*displaying QFP PPE utilization
information, 167–168*

displaying RP/ESP/SIP

insertion/uptime, 166–167

displaying SPA status in SIP, 163

*displaying SPA-level
statistics, 164*

memory utilization, 154–155

show commands, 150–153,
169–174

*tracking command output via
monitor command, 162*

*tracking control CPU usage
from Linux shell, 161–162*

*“Warning: Filesystem Is Not
clean” error messages,
174–175*

unified communications services
case studies, CUBE, 269, 271,
273–279

virtualization, 190

voice header compression via
cRTP, 267

VPN

SP L2 VPN, 51

SP L3 VPN, 51

WAAS, 262

WAN optimization, 262–265

WAN aggregation, 49

WAN optimization, 262

branch deployments, 264

*campus headend
deployments, 263*

IronPort appliances, 265

WSA, 265

WCCPv2, 259

troubleshooting, 265–266

web caching, 260–262

WebEx Node services module, 269,
280–283, 285

“ASR1000_PEM-3-PEMFAIL: The PEM in Slot 0 Is Switched Off or Encountering a Failure Condition” error message, 175

audit trail messages, NetFlow event logging, 256

automation (policy-based), encapsulation and, 20

availability

ASR 1000 series routers, 46

HA

NSF, 179, 181–183

SSO, 179, 181–183

multicast HA, ASR 1000 series routers, 220–221

B

bandwidth

excess bandwidth, examples of, 214

minimum bandwidth, examples of, 214

WAN architectures, bandwidth commoditization in, 22–23, 25

WebEx Node services module, 281

benchmarking routers, 40

data-plane performance/scale, 41

routing-plane performance/scale, 41

BITS (Building Integrated Timing Source) reference clocks, ASR 1000 series routers, 60

bootflash

ASR 1000 series routers, ISSU, 117

ASR 1000 series RP, 73

IOS XE software, 96

booting

ASR 1000 series routers, 103

ROMMON, 104–106

valid configurations, 103

IOS XE software, 96, 98

nonmodular boot procedure, 97

RP, “Warning: Filesystem Is Not Clean” error messages, 174–175

branch aggregation, ASR series routers, 50

branch deployments, WAN optimization via WAAS integration, 264

branch WAN aggregation, 2

connectivity options table, 4

feature requirements table, 5

secure WAN technologies table, 5

SLA requirements, 5

BRAS (broadband aggregation systems), ASR series 1000 routers, 53–54

buffering

egress SIP buffering, 211–212

GPM, 208

ingress SIP buffering, 207–208

multicast packets, 210

packet buffer DRAM (QFP Traffic Manager), 208–209

packet buffering (QFP Traffic Manager), 209–210

punt packets, 210

unicast packets, 210

business drivers, WAN architectures

bandwidth commoditization, 22–23, 25

carbon footprint reduction, 23, 26

infrastructure consolidation, 19, 25

regulatory compliance, 24, 26

reliability, 22, 25

security, 22, 25

segmentation/virtualization, 20–21, 25

service awareness/integration, 18, 25
 time to adoption, 24, 26
 time to understanding, 24, 26
 troubleshooting, 24, 26

C

campus headend deployments, WAN
 optimization via WAAS
 integration, 263

carbon footprint reduction

Enterprise Edge platforms, 31–32
 WAN architectures, 23, 26

carrier-class routing

ASR series 1000 routers, 57
 Enterprise Edge platforms, 29

chassis manager

ESP software, 65
ASR 1000 series routers, 87
 RP software, 64
ASR 1000 series routers, 85
 SIP software, 66
ASR 1000 series routers, 87

Cisco Self-Defending Network
 schema, 231

compliance (regulatory), WAN
 architectures, 24, 26

conditional policing, PQ (ASR 1000
 series routers), 214

configuring

ASR 1000 GUI, 142
 ASR 1000 series routers, 103,
 107, 109

consolidation (infrastructure), WAN
 architectures, 19, 25

control plane (ASR 1000 series
 routers), 59, 177–178

convergence times (DCI), 195

cooling systems, WAN architectures,
 23, 26

CoPP (control-plane policing)
 feature, 78

CPE (customer premises equipment),
 ASR series 1000 routers, 54

cRTP (compressed Real-Time
 Protocol), voice header
 compression, 267

crypto engines

ASR 1000 series router ESP, 82–83
 GETVPN, 245
 multicast encryption, 236
 QoS scalable encryption, 237, 239

CUBE (Cisco Unified Border Element
 for service provider), 269

business-to-business telepresence
 deployment scenario, 276–278

integrated CUBE, 271, 273

SP-to-managed enterprise and
 residential SIP trunking
 deployment scenario, 275–276

SP-to-SP peering deployment
 scenario, 274

troubleshooting, 279

CWDM/DWDM, ASR 1000 series
 routers, 196

D

dark fiber transport (DCI), 195

data plane, ASR 1000 series
 routers, 59

partitioning, 68

DBE (Data Border Elements), 272
 ASR 1000 series routers, 274

DCI (data center interconnects), 10
 convergence times, 195

dark fiber transport, 195

feature requirements table, 11–13

loop prevention, 195

redundancy, 195

scalability, 195–196

traffic encryption over EoMPLS

Pseudowire at Layer 2 via

TrustSec, 198, 200–203

DDoS (distributed denial of service), self-inflicted, 256

debug commands (ASR 1000 series routers), troubleshooting, 150–153, 168–174

DM VPN (Dynamic Multipoint VPN), 7

hub design, ASR 1000 series routers, 239–241

multipoint GRE tunnels, 239

NHRP, 239

DRAM

ASR 1000 series router ESP, 79

ASR 1000 series RP, 73

QFP memory statistics, displaying for, 156–157

drop messages, NetFlow event logging, 256

dual IOSD, running on 2/4RU systems, 137–138

dual IOSD failovers, ESP, 193

E

egress processing (transit packets), ASR 1000 series routers, 91

egress SIP buffering, 211–212

embedded services, ASR series 1000 routers, 56

encapsulation, virtualization and, 20

encryption

EoMPLS Pseudowire at Layer 2 using TrustSec, 198, 200–203

multicast encryption, ASR 1000 series routers, 236–237

scalable encryption, ASR 1000 series routers, 237, 239

Enterprise Edge platforms

carbon footprint reduction, 31–32

carrier-class routing, 29

feature velocity, 31

flexible system architectures, 30

industry standard compliance, 32

interface diversity/density, 31

QoS, 30

service integration, 29

system investment protection, 31

test plans, writing, 32

functional tests, 35

load testing methodology, 35

longevity testing

methodology, 35

negative testing

methodology, 35

performance tests, 35

positive testing

methodology, 35

scale tests, 35

stress testing methodology, 35

test case details, 36

test entry/exit criteria, 35

test resources, 34

test results reporting, 36

test schedules, 36

test scope/objective, 34

test setup/topology, 34

Enterprise Private WAN, segmentation, 187, 189

environmental concerns

Enterprise Edge platforms, 31–32

WAN architectures, 23, 26

EoMPLS Pseudowire, traffic encryption at Layer 2 using TrustSec, 198, 200–203

error messages

“ASR1000_PEM-3-PEMFAIL: The PEM in Slot 0 Is Switched Off or Encountering a Failure Condition” messages, 175

“Warning: Filesystem Is Not Clean” messages, 174–175

ERSPAN (Encapsulated Remote SPAN)

ASR 1000 series router support, 184–186

packet capturing, 184–186

ESP (embedded service processors), 220

active ESP, 220

ASR 1000 series routers, 63, 79, 103, 106

displaying RP insertion/uptime, 166–167

DRAM, 79

ESP chassis manager, 87

ESP crypto engines, 82–83

ESP forwarding manager, 86

ESP initialization, 81

ESP packet handling, 81

ESP subpackages, 135–136

interconnect ASIC, 79

QFP software, 86

system board, 80

TCAM, 80

IOS zone-based firewalls, 193

IOSD failovers, 193

IPsec, 193

ISO XE software, 94–95

LED color and description table, 79

multicast replication, 221

NetFlow, 193

software, ASR 1000 series routers, 65

ESP Interconnect Scheduler, aggregating SIP traffic, 213

Ethernet

frame capture/transport across Layer 3 clouds, 184–186

Metro Ethernet, 4

event logging (NetFlow)

audit trail messages, 256

drop messages, 256

F

feature velocity, Enterprise Edge platforms, 31

FIB (Forwarding Information Base), 183

firewalls

high-speed logging via NetFlow, 223–225

IOS zone-based firewalls, 251

ESP, 193

HA, 253

scalable multigigabit router firewalls, 254–256

zone pair scale, 253

WAAS, 261

WCCPv2, 261–262

flexibility, Internet edge routers, 193

flexible system architectures, Enterprise Edge platforms, 30

forwarding manager

ESP software, 66

ASR 1000 series routers, 86

RP software, 65

ASR 1000 series routers, 86

FPM, scalable multigigabit, 225, 227–228

front-panel LED, status of, displaying via `show platform hardware` command, 163

front panels, ASR 1000 series router RP, 72

functional tests, Enterprise Edge platforms, 35

G

gateways (Internet), ASR series routers, 49–50

GDOI (Group Domain of Interpretation), 242

GM registration, 244

GET VPN (Group Encrypted Transport VPN), 7

ASR 1000 series routers, 242, 244–246

- branch design deployment model*, 248
- DC design deployment model*, 248
- limitations in*, 248
- memory*, 247
- supported features*, 247
- troubleshooting*, 250

crypto engine, 245

GDOI, 242

- GM registration*, 244

GM, 242

- registering*, 244

key servers, 242

QFP, 245

GM (group members), GETVPN, 242

- registering in, 244

GPM (global packet memory), 208

GR (graceful restarts), 179

GRE (generic routing encapsulation)

- MPLS over GRE, segmentation, 187
- multipoint GRE tunnels, 239
- p2p GRE inside IPsec, 7
- p2p GRE tunnels, WAN edge routers, 190

GRE over IPsec, 236

H

HA (high availability)

- IOS firewalls, 253
- IOS XE software, 94
- IPsec, ASR 1000 series routers, 236
- multicast HA, ASR 1000 series routers, 220–221
- NSF, 179, 181–183
- SSO, 179, 181–183

hardware-assisted control-plane protection, ASR 1000 series RP, 78

HDD (hard disk drives)

- ASR 1000 series RP, 74
- RP, IOS XE software, 95–96

HQF (Hierarchical Queuing Framework), ASR 1000 series router QoS, 207

I-J

I/O (input/output) plane, ASR 1000 series routers, partitioning, 68

IKE (Internet Key Exchanges), initiation requests, ASR 1000 series routers, 235

industry standard compliance, Enterprise Edge platforms, 32

infrastructure consolidation, WAN architectures, 19, 25

ingress processing (transit packets), ASR 1000 series routers, 88

ingress SIP buffering, 207–208

initializing

ASR 1000 series router ESP, 81

ASR 1000 series RP, 75

injected packets

defining, 74

displaying, 77

installing WebEx Node services module, 283, 285

integrated QoS (Quality of Service), ASR 1000 series routers, 59

integrated threat control solutions, ASR 1000 series routers, IOS firewalls, 251, 253–256

interconnect ASIC (application-specific integrated circuits)

ASR 1000 series ESP, 79

ASR 1000 series RP, 74

interface diversity/density, Enterprise Edge platforms, 31

interface manager

RP software, 65

ASR 1000 series routers, 86

SIP software, 66

ASR 1000 series routers, 87

Internet edge role

feature requirements table, 9–10

router functionality table, 8–9

Internet edge routers, scalability/flexibility of, 193

Internet gateway/edge routers, 193

Internet gateways, ASR series routers, 49–50

IOS

CoPP feature, 78

troubleshooting, 169–174

IOS firewalls, 251

HA, 253

scalable multigigabit router firewalls, 254–256

zone pair scale, 253

IOS XE, ASR 1000 series routers, 63

IOS XE software, 93

ASR 1000 series routers, packaging, 98

benefits of, 94

bootflash, 96

booting, 96, 98

nonmodular boot procedure, 97

components of, 93

ESP, 94–95

feature support, 94

HA, 94

IOSD, 93

licensing, 100

Linux 2.6.x kernel, 94

middleware processes, 94

modularity, 94

packaging, 96–98

subpackages of, 94–95

QFP software, 94

redundancy, 98

releases, overall release plan, 99

RP, 94–95

HDD file system structure, 95–96

security, 94

SIP, 94–95

UMI, 94

IOS zone-based firewalls, ESP, 193

IOSD (IOS daemon), 93

- ASR 1000 series routers, 67
- dual IOSD, running on 2/4RU systems, 137–138

IOSD failovers, ESP, 193

IP, DCI transport, 195

IP routing

- DCI
 - scalability, 195–196*
 - traffic encryption over EoMPLS Pseudowire at Layer 2 via TrustSec, 198, 200–203*
- Internet edge routers, scalability/flexibility of, 193
- NetFlow, 184
- packet capturing, ERSPAN, 184–186
- RR, 191–192
- segmentation, MPLS over GRE, 187, 189

IP services, 205

- high-speed logging via NetFlow, 223–225
- IPv4/IPv6 scalable multicast acceleration via QFP, 219
 - multicast HA, 220–221*
 - multicast replication on ESP, 221*
- scalable hierarchical QoS, 216–217, 219
- scalable multigigabit NAT, 221–223
- scalable multigigabit NBAR/FPM, 225, 227–228
- scalable QoS via QFP Traffic Manager, 206–207
 - aggregating SIP traffic, 213*
 - egress SIP buffering, 211–212*
 - ingress SIP buffering, 207–208*

packet buffering, 209–210

priority queues, 213–215

IPsec, 7

- ESP, 193
- GRE over IPsec, 236
- HA, ASR 1000 series routers, 236
- multicast encryption, ASR 1000 series routers, 236–237
- p2p GRE inside IPsec, 7
- packet flow (egress), ASR 1000 series routers, 235
- packet flow (ingress), ASR 1000 series routers, 235
- VPN, ASR 1000 series routers, 232, 234

IPv4, scalable multicast acceleration via QFP, 219

- multicast HA, 220–221
- multicast replication on ESP, 221

IPv6, scalable multicast acceleration via QFP, 219

- multicast HA, 220–221
- multicast replication on ESP, 221

IRAM, QFP memory statistics, displaying for, 156–157

IronPort appliances, WAN optimization, 265

ISSU (in-service software upgrades)

- ASR 1000 series routers, 58, 114, 117
 - consolidated packages on fully redundant 6RU, 117*
 - ISSU impact on subpackages, 114–115*
 - ISSU on 6RU system for IOSD, 118–123*
 - running dual IOSD on 2/4RU systems, 137–138*
 - subpackage ISSU on fully redundant 6RU, 124–131*

upgrading ESP subpackages,
135–136

upgrading RP subpackages,
132–134

upgrading SIP/SPA subpackages, 131

benefits of

business benefits, 114

operational benefits, 113

issu acceptversion command, 118, 122

issu commitversion command, 118,
133, 135

issu loadversion command, 118–119,
131, 135

issu runversion command, 118, 121

issu set rollback-timer command, 118

K

key servers, GETVPN, 242

L

LAN (local area networks), ASR 1000
series routers, 60

large branch WAN (wide area
networks), office deployment
requirements table, 13–14

Layer 3 clouds, Ethernet frame
capture/transport, 184–186

LED, displaying status of via show
platform hardware command, 163

licensing IOS XE software, 100

Linux

ASR 1000 series routers, 67

control CPU usage, tracking from
Linux shell, 161–162

Linux 2.6.x kernel, IOS XE
software, 94

LLQ (low-latency queuing), ASR
1000 series routers, 83

load testing methodology, Enterprise
Edge platforms, 35

logs

high-speed logging via NetFlow,
223–225

NetFlow event logging, 256

longevity testing methodology,
Enterprise Edge platforms, 35

loop prevention, DCI, 195

M

management plane, ASR 1000 series
routers, 59

MEC (Multichassis EtherChannel),
195

memory, ASR 1000 series routers,
154–155

GETVPN, 247

Metro Ethernet, 4

MFIB, 219

middleware, IOS XE software, 94

mobility (applications), virtualization
and, 20

modular control plane (ASR 1000
series routers), 177–178

modularity, IOS XE software, 94

momentary packet loss, 220

monitor command, command output,
tracking via, 162

MPLS

DCI transport, 195

MPLS over GRE, segmentation, 187

MRIB, 219

multicast encryption, ASR 1000
series routers, 236–237

multicast HA (high availability), ASR 1000 series routers, 220–221

multicast packets, buffering, 210

multipoint GRE tunnels, 239

N

NAT (Network Address Translation)

high-speed logging via NetFlow, 223–225

scalable in-built multigigabit NAT, 221–223

NBAR (Network Based Application Recognition)

scalable multigigabit NBAR, 225, 227–228

WCCPv2, 261

NDR (non drop rates), 35

negative testing methodology, Enterprise Edge platforms, 35

NetFlow, 184

ESP, 193

event logging, 256

firewalls, high-speed logging, 223–225

NAT, high-speed logging, 223–225

NHRP (Next-Hop Resolution Protocol), 239

nonmodular boot procedure, IOS XE software, 97

nonstop router management, ASR 1000 series routers, 60

NSF (nonstop forwarding), 179

HA, 179, 181–183

NSR (nonstop routing), 179

O

optimized WAN aggregation, 2

oversubscription, ASR 1000 series routers, 59

P

p2p GRE inside IPsec, 7

p2p GRE tunnels, WAN edge routers, 190

packet buffer DRAM (QFP Traffic Manager), 208–209

packet buffering, QFP Traffic Manager, 209

multicast packets, 210

punt packets, 210

unicast packets, 210

packet capturing, ERSPAN, 184

Ethernet frame capture/transport across Layer 3 clouds, 184–186

packet handling

ASR 1000 series router ESP, 81

ASR 1000 series RP, 75, 77

transit packets

arrival processing, 88–90

egress processing, 91

ingress processing, 88

packets, momentary packet loss, 220

partitioning ASR 1000 series routers, 68

PBR (Policy Based Routing), WCCPv2, 261

performance

data-plane performance (benchmarking routers), 41

- routing-plane performance
 - (benchmarking routers), 41
- tests, Enterprise Edge platforms, 35
- policy aggregation, 215**
- policy-based automation, virtualization and, 20**
- positive testing methodology, Enterprise Edge platforms, 35**
- power usage**
 - Enterprise Edge platforms, 31–32
 - WAN architectures, 23, 26
- PQ (priority queuing), ASR 1000 series routers, 214**
- private WAN aggregation, 2**
 - connectivity options table, 4
 - feature requirements table, 5
 - secure WAN technologies table, 5
 - SLA requirements, 5
- Pseudowire, traffic encryption over EoMPLS Pseudowire at Layer 2 using TrustSec, 198, 200–203**
- punt packets, buffering, 210**
- punted traffic**
 - defining, 74
 - displaying, 75, 77

Q

- QFP (QuantumFlow Processor), 78**
 - ASR 1000 series routers, 56, 106–107
 - ASR 1000 series router ESP, 86
 - GETVPN, 245
 - IOS XE software, 94
 - IPv4-related drops, displaying for, 155–156
 - IPv4/IPv6 scalable multicast acceleration, 219

- multicast HA, 220–221*
- multicast replication on ESP, 221*

- memory statistics
 - displaying for IRAM/DRAM/SRAM, 156–157*
 - displaying on per-IOS feature/internal-usage basis, 157–161*
- QFP utilization information, displaying, 167–168
- Traffic Manager
 - packet buffer DRAM, 208–209*
 - packet buffering, 209–210*
 - scalable QoS via, 206–215*

QoS (Quality of Service)

- 7200VXR, QoS scaling guidelines, 216
- ASR 1000 series routers, 48, 207
 - QoS scaling guidelines, 216*
- Enterprise Edge platforms, 30
- integrated QoS, ASR 1000 series routers, 59
- QoS preclassify, 237
- scalability
 - hierarchical QoS, 216–217, 219*
 - QFP Traffic Manager, 206–215*
- scalable encryption, ASR 1000 series routers, 237, 239

R

redundancy

- ASR 1000 series routers, 68
- DCI, 195
- IOS XE software, 98
- RP subpackages, ASR 1000 series routers, 134

reliability

- ASR 1000 series routers, 46–47, 49
- WAN architectures, 22, 25

remote user aggregation, ASR series routers, 50**remote-access VPN (virtual private networks), 7****replay protection via TBAR (time-based antireplay), 245****ROI (return on investment), ASR 1000 series routers, 49****ROMMON, ASR 1000 series routers, 68, 104–106****routers**

- ASR 1000 series routers, 43–44
 - availability, 46*
 - bootflash, 117*
 - booting, 103–106*
 - branch aggregation, 50*
 - BRAS, 53–54*
 - carrier-class routing, 57–60*
 - configuring, 107, 109*
 - control plane, 177–178*
 - CPE, 54*
 - CWDM/DWDM, 196*
 - DBE, 274*
 - embedded services, 56*
 - ERSPAN support, 184*
 - ESP, 79, 81–83, 86–87, 103, 106, 135–136, 166–167*
 - Ethernet frame capture/transport across Layer 3 clouds, 184–186*
 - file system structure, 109–110*
 - hardware components, 61–63*
 - hidden costs of, 49*
 - IKE initiation requests, 235*
 - Internet gateway/edge routers, 193*

*Internet gateways, 49–50**IOS XE software, 98–99**IP services case studies, 205–217, 219–225, 227–228**ISSU, 114–115, 117–138**life span of, 49**modular operation, 87**multicast HA, 220–221**multicast replication on ESP, 221**operational savings, 49**packet handling, 88–91**partitioning, 68**QFP, 56, 106–107, 167–168**QoS, 48, 207**QoS scaling guidelines, 216**reliability, 46–47**remote user aggregation, 50**ROI, 49**RP, 71–73, 75, 77–78, 85–86, 103, 132–134, 166–167**RR, 191–192**SBC, 271**SBE, 273, 277–278**scalability, 48**security headends, 50**security services case studies, 231–232, 234–237, 239–242, 244–248, 250–251, 253–256**segmentation, 187, 189**SIP, 83–84, 87, 103, 106, 131, 163, 166–167**software components, 63, 65–68**SP L2 VPN, 51**SP L3 VPN, 51**SPA, 131, 163–164**subpackages, 114–115, 124–136*

- system redundancy/modularity*, 68
- traffic encryption over EoMPLS Pseudowire at Layer 2 via TrustSec*, 198, 200–203
- traffic manager priority queues*, 213–215
- troubleshooting*, 150–175
- unified communications services case studies*, 269, 271, 273–279
- virtualization*, 190
- voice header compression via cRTP*, 267
- WAAS, 262–265
- WAN aggregation, 49
- WAN optimization, 262–265
- WCCPv2, 259–262, 265–266
- WebEx Node services module, 269, 280–283, 285
- benchmarking, 40
 - data-plane performance/scale*, 41
 - routing-plane performance/scale*, 41
- characteristics of, 39–40
- choosing, 39–40
- DCI feature requirements table, 12–13
- Internet edge role
 - feature requirements table*, 9–10
 - router functionality table*, 8–9
- Internet edge routers, scalability/flexibility of, 193
- Internet gateway/edge routers, 193
- nonstop router management, ASR 1000 series routers, 60
- WAN edge routers, p2p GRE tunnels, 190

- routing plane (ASR 1000 series routers), partitioning, 68
- RP (Route Processors), 220
 - active RP, 220
 - ASR 1000 series routers, 62, 71, 103
 - bootflash*, 73
 - displaying RP insertion/uptime*, 166–167
 - DRAM, 73
 - front panel*, 72
 - hardware-assisted control-plane protection*, 78
 - HDD, 74
 - initializing RP*, 75
 - interconnect ASIC*, 74
 - legacy protocol traffic*, 78
 - packet handling*, 75, 77
 - RP chassis manager*, 85
 - RP CPU tasks*, 85
 - RP forwarding manager*, 86
 - RP interface manager*, 86
 - RP subpackages*, 132–134
 - booting, “Warning: Filesystem Is Not Clean” error messages, 174–175
 - IOS XE software, 94–95
 - HDD file system structure*, 95–96
 - LED color and description table, 72–73
- RP software, ASR 1000 series routers, 63, 65
- RR (route reflection), 191–192

S

- SBC, ASR 1000 series routers, 271
- SBE (Signaling Border Elements), 272
 - ASR 1000 series routers, 273, 277–278

scalability

- ASR 1000 series routers, 48
- DCI, 195–196
- Internet edge routers, 193
- IPv4/IPv6 scalable multicast
 - acceleration via QFP, 219–221
- multigigabit FPM, 225, 227–228
- multigigabit NAT, 221–223
- multigigabit NBAR, 225, 227–228
- QoS
 - hierarchical QoS*, 216–217, 219
 - QFP Traffic Manager*, 206–215
- scalable control plane (ASR 1000 series routers), 177–178
- scalable encryption (ASR 1000 series routers), 237, 239
- scale tests (Enterprise Edge platforms), 35
- schedulers, ESP Interconnect Scheduler, aggregating SIP traffic, 213
- secure WAN aggregation, 2
- security
 - IOS XE software, 94
 - WAN architectures, 22, 25
 - WebEx Node services module, 281
- security headends, ASR series routers, 50
- security services, 231
 - Cisco Self-Defending Network schema, 231
 - crypto engines, multicast encryption, 236
 - DMVPN
 - ASR 1000 series routers*, 239–241
 - hub design*, 239–241

- GETVPN, ASR 1000 series routers, 242, 244–248, 250
- integrated threat control solutions, ASR 1000 series routers, 251, 253–256
- IOS firewalls, 251
 - HA*, 253
 - scalable multigigabit router firewalls*, 254–256
 - zone pair scale*, 253
- IPsec
 - GRE over IPsec*, 236
 - HA*, 236
 - multicast encryption*, 236–237
 - packet flow (egress)*, 235
 - packet flow (ingress)*, 235
 - VPN*, 232, 234
- QoS, scalable encryption, 237, 239
- SPD, 235
- segmentation
 - Enterprise Private WAN, 187, 189
 - MPLS over GRE, 187
 - WAN architectures, 20–21, 25
- self-inflicted DDoS (distributed denial of service), 256
- service awareness, WAN architectures, 18, 25
- service integration
 - Enterprise Edge platforms, 29
 - WAN architectures, 18, 25
- show commands, troubleshooting
 - ASR 1000 series routers, 150–153, 169–174
- show inventory command, 64
- SIP (SPA interface processors), 220
 - aggregating traffic, ESP Interconnect Scheduler, 213

- ASR 1000 series routers, 63, 83–84, 103, 106
 - displaying SIP*
 - insertion/uptime*, 166–167
 - SIP chassis manager*, 87
 - SIP interface manager*, 87
 - SIP subpackages*, 131
 - SPA drivers*, 87
- egress SIP buffering, 211–212
- ingress SIP buffering, 207–208
- SPA, displaying status in SIP, 163
- IOS XE software, 94–95
- SIP software, ASR 1000 series routers, 66
- SLA (Service Level Agreements), branch/private WAN aggregation, SLA requirements for, 5
- SP L2 VPN (Service Provider Layer 2 VPN), ASR series 1000 routers, 51
- SP L3 VPN (Service Provider Layer 3 VPN), ASR series 1000 routers, 51
- SPA
 - ASR 1000 series routers, SPA subpackages, 131
 - SIP, displaying SPA status in, 163
 - SPA-level statistics, displaying, 164
- SPA drivers
 - SIP, ASR 1000 series routers, 87
 - SIP software, 66
- SPAN (Switch Port Analyzers), ERSPAN
 - ASR 1000 series router packet capturing, 184–186
 - ASR 1000 series router support, 184
 - packet capturing, 184–186
- SPD (security policy databases), 235
- SRAM, QFP memory statistics, displaying for, 156–157

- SSO (stateful switchovers), 179
 - HA, 179, 181–183
- stress testing methodology, Enterprise Edge platforms, 35
- switches, DCI feature requirements table, 12–13
- system investment protection, Enterprise Edge platforms, 31
- system redundancy, ASR 1000 series routers, 68

T

- TBAR (time-based antireplay), GETVPN, 245
- TCAM (tertiary content-addressable memory), ASR 1000 series ESP, 80
- test plans, writing for Enterprise Edge platforms, 32
 - functional tests, 35
 - load testing methodology, 35
 - longevity testing methodology, 35
 - negative testing methodology, 35
 - performance tests, 35
 - positive testing methodology, 35
 - scale tests, 35
 - stress testing methodology, 35
- test case details, 36
- test entry/exit criteria, 35
- test resources, 34
- test results reporting, 36
- test schedules, 36
- test scope/objective, 34
- test setup/topology, 34
- Traffic Manager (QFP)
 - packet buffer DRAM, 208–209
 - packet buffering, 209

- multicast packets*, 210
- punt packets*, 210
- unicast packets*, 210
- transit packets (ASR 1000 series routers)**
 - arrival processing, 88–90
 - egress processing, 91
 - ingress processing, 88
- troubleshooting**
 - ASR 1000 series routers
 - debug commands*, 150–153, 168–174
 - displaying drop statistics*, 164–165
 - displaying front-panel LED status via show platform hardware command*, 163
 - displaying interface-level binding*, 165
 - displaying IPv4-related drops for active QFP*, 155–156
 - displaying processors*, 154–155
 - displaying QFP memory statistics for IRAM/DRAM/SRAM usage*, 156–157
 - displaying QFP memory statistics on per-IOS feature/internal-usage basis*, 157–161
 - displaying QFP PPE utilization information*, 167–168
 - displaying RP/ESP/SIP insertion/uptime*, 166–167
 - displaying SPA status in SIP*, 163
 - displaying SPA-level statistics*, 164
 - memory utilization*, 154–155
 - show commands*, 150–153, 169–174

- tracking command output via monitor command*, 162
- tracking control CPU usage from Linux shell*, 161–162
- “Warning: Filesystem Is Not Clean” error messages, 174–175

- CUBE, 279
- GETVPN, ASR 1000 series routers, 250
- IOS, 169–174
- methodology of, 149
- WAN architectures, 24, 26
- WCCPv2, 265–266
- TrustSec, traffic encryption over EoMPLS Pseudowire at Layer 2**, 198, 200–203

U

- UMI (unified management interfaces), IOS XE software, 94
- unconditional policing, PQ (ASR 1000 series routers)**, 214
- understanding, time to (WAN architectures)**, 24, 26
- unicast packets, buffering**, 210
- unified communications services**
 - CUBE, 269
 - business-to-business telepresence deployment scenario*, 276–278
 - integrated CUBE*, 271, 273
 - SP-to-managed enterprise and residential SIP trunking deployment scenario*, 275–276
 - SP-to-SP peering deployment scenario*, 274
 - troubleshooting*, 279

WebEx Node services module,
 269, 280
bandwidth, 281
deploying, 282
installing, 283, 285
security, 281
VoIP/video solution mode, 281
web presentation mode, 281

upgrades

ESP subpackages, ASR 1000 series routers, 135–136

ISSU

ASR 1000 series routers, 58, 114–115, 117–136
benefits of, 113–114
issu acceptversion command, 118, 122
issu commitversion command, 118, 133, 135
issu loadversion command, 118–119, 131, 135
issu runversion command, 118, 121
issu set rollback-timer command, 118

RP subpackages, ASR 1000 series routers, 132–134

SIP subpackages, ASR 1000 series routers, 131

SPA subpackages, ASR 1000 series routers, 131

users, remote user aggregation, ASR 1000 series routers, 50

V

virtualization

application mobility and, 20
 ASR 1000 series routers, 190
 encapsulation and, 20

WAN architectures, 20–21, 25
 voice header compression, cRTP, 267

VoIP/video solution mode (WebEx Node services module), 281

vPC (virtual port channels), 195

VPN (virtual private networks)

DM VPN, 7

ASR 1000 series routers, 239–241

hub design, 239–241

multipoint GRE tunnels, 239

NHRP, 239

GET VPN, 7

ASR 1000 series routers, 242, 244–248, 250

crypto engine, 245

GDOI, 242, 244

GM, 242, 244

key servers, 242

QFP, 245

IPsec, ASR 1000 series routers, 232, 234

remote-access VPN, 7

SP L2 VPN, ASR series 1000 routers, 51

SP L3 VPN, ASR series 1000 routers, 51

VPNv4 routes, RR, 191–192

VSS (Virtual Switching Systems), 195

W-Z

WAAS (Wide Area Application Services), 259

ASR 1000 series routers, 262

WAN optimization, 262–265

firewalls, 261

WCCPv2, 259

WAN (wide area networks)

- ASR 1000 series routers, 60
 - aggregation*, 49
- branch WAN aggregation, 2
 - connectivity options table*, 4
 - feature requirements table*, 5
 - secure WAN technologies table*, 5
 - SLA requirements*, 5
- DCI, 10
 - feature requirements table*, 11–13
- edge routers, p2p GRE tunnels, 190
- Enterprise Private WAN,
 - segmentation, 187, 189
- Internet edge role
 - feature requirements table*, 9–10
 - router functionality table*, 8–9
- large branch WAN, office
 - deployment requirements
 - table, 13–14
- private WAN aggregation, 2
 - connectivity options table*, 4
 - feature requirements table*, 5
 - secure WAN technologies table*, 5
 - SLA requirements*, 5

WAN architectures

- business drivers
 - bandwidth commoditization*, 22–23, 25
 - carbon footprint reduction*, 23, 26
 - infrastructure consolidation*, 19, 25
 - regulatory compliance*, 24, 26
 - reliability*, 22, 25
 - security*, 22, 25
 - segmentation/virtualization*, 20–21, 25

- service awareness/integration*, 18, 25
- time to adoption*, 24, 26
- time to understanding*, 24, 26
- troubleshooting*, 24, 26
- evolution of, 17

WAN optimization, 262

- branch deployments, 264
- campus headend deployments, 263
- IronPort appliances, 265
- WSA, 265

“Warning: Filesystem Is Not Clean” error messages, 174–175**WCCPv2 (Web Cache Control Protocol version 2), 259**

- firewalls, 261–262
- NBAR, 261
- PBR, 261
- troubleshooting, 265–266
- web caching, 260–262

web caching, WCCPv2, 260–262**web interfaces, ASR 1000 GUI**

- configuring, 142
- usage examples, 143, 146
- views of, 141–142

web presentation mode (WebEx Node services module), 281**WebEx Node services module, 269, 280**

- bandwidth, 281
- deploying, 282
- installing, 283, 285
- security, 281
- VoIP/video solution mode, 281
- web presentation mode, 281

WSA (Web Security Appliances), WAN optimization, 265