CISCO VALIDATED DESIGN

Intelligent WAN Technology Design Guide

February 2016

Table of Contents

Introduction 1

- Technology Use Cases 1
- Cisco Intelligent WAN Overview 2
- Design Overview 4

Deploying the Cisco Intelligent WAN 17

- Overall IWAN Architecture Design Goals 17

Deploying the Transport Independent Design 19

- Design Overview 19
- Deployment Details 26
- Configuring an IOS Certificate Authority 26
- Configuring DMVPN Hub Router 30
- Configuring the Firewall and DMZ Switch 53
- Configuring Remote-Site DMVPN Router 62
- Adding Second DMVPN for a Single-Router Remote Site 89
- Adding LTE fallback DMVPN for a single-router remote site 104
- Modifying the First Router for Dual Router Design 120
- Configuring Remote-Site DMVPN Router (Router 2) 125

Deploying an IWAN Remote-Site Distribution Layer 156

- Connecting Remote-site Router to Distribution Layer 156
- Connecting Remote-Site Router to Distribution Layer (Router 2) 163

Deploying IWAN Performance Routing 169

Configuring Hub Master Controller 172

Configuring PfR for Hub Location 176

Configuring PfR for Remote Site Locations 189

Configuring Hub Master Controller High Availability 201

Configuring Hub Border Router Scalability 208

Deploying a Second Data Center Location 221

Configuring Transit Border Routers 223

Configuring Transit Master Controller 235

Configuring PfR for Transit Location 238

Deploying IWAN Quality of Service 245

Configuring QoS for DMVPN Routers 245

Applying DMVPN QoS Policy to DMVPN Hub Routers 252

Applying QoS Configurations to Remote Site Routers 257

Deploying IWAN Monitoring 263

Configuring Flexible NetFlow for IWAN Monitoring 263

Appendix A: Product List 272

Appendix B: Technical Feature Supplement 276

Front Door VRF for DMVPN 276

Appendix C: Common Procedures 279

Appendix D: Changes 282

Appendix E: Device Configuration Files 283

Introduction

The Cisco Intelligent WAN (IWAN) solution provides design and implementation guidance for organizations looking to deploy wide area network (WAN) transport with a transport-independent design, intelligent path control, application optimization, and secure encrypted communications between branch locations while reducing the operating cost of the WAN. IWAN takes full advantage of cost-effective transport services in order to increase bandwidth capacity without compromising performance, reliability, or security of collaboration or cloud-based applications.

TECHNOLOGY USE CASES

Organizations require the WAN to provide sufficient performance and reliability for the remote-site users to be effective in supporting the business. Although most of the applications and services that the remote-site worker uses are centrally located, the WAN design must provide the workforce with a common resource-access experience, regardless of location.

Carrier-based multiprotocol label switching (MPLS) service is not always available or cost-effective for an organization to use exclusively for remote-site WAN connectivity. There are multiple WAN transport offerings that can be used simultaneously to create a robust, secure, and cost-effective WAN, including MPLS virtual private networks (VPNs), Internet, Cellular (3G/LTE), and Carrier Ethernet. Internet-based IP VPNs offer attractive bandwidth pricing and can augment premium MPLS offerings or replace MPLS in some scenarios. A flexible network architecture should include all common WAN transport offerings as options without significantly increasing the complexity of the overall design.

While Internet IP VPN networks present an attractive option for effective WAN connectivity, anytime an organization sends data across a public network there is risk that the data will be compromised. Loss or corruption of data can result in a regulatory violation and can present a negative public image, either of which can have significant financial impact on an organization. Secure data transport over public networks like the Internet requires adequate encryption to protect business information.

Use Case: Secure Site-to-Site WAN Communications

This guide helps organizations connect remote sites over private (MPLS VPN) and public (Internet) IP networks, efficiently and securely.

This design guide enables the following network capabilities:

- Secure, encrypted communications solutions for up to 2000 locations within a single domain using a dynamic multipoint VPN (DMVPN) IPsec tunnel overlay configuration
- A multi-homed solution that can have two or more connectivity options for resiliency and efficient use of all WAN bandwidth, using single or dual routers in remote locations
- Support for IP Multicast and replication performed on core, hub-site routers
- Compatibility with public Internet networks where network address translation (NAT) is implemented

CISCO INTELLIGENT WAN OVERVIEW

With the advent of globalization, WANs have become a major artery for communication between remote offices and customers in any corner of the world. Additionally, with data center consolidation, applications are moving to centralized data centers and clouds. WANs now play an even more critical role, because business survival is dependent on the availability and performance of the network.

Until now, the only way to get reliable connectivity with predictable performance was to take advantage of a private WAN using MPLS or leased line service. However, carrier-based MPLS and leased line services can be expensive and are not always cost-effective for an organization to use for WAN transport in order to support growing bandwidth requirements for remote-site connectivity. Organizations are looking for ways to lower operating budget while adequately providing the network transport for a remote site.

As bandwidth demands have increased, the Internet has become a much more stable platform, and the price-to-performance gains are very attractive. However, businesses are primarily deploying “Internet as WAN” in their smaller sites or as a backup path because of the risks. Now this cost-effective, performance-enhancing opportunity can be realized at all your branch offices with Cisco IWAN.

Cisco IWAN enables organizations to deliver an uncompromised experience over any connection. With Cisco IWAN IT organizations can provide more bandwidth to their branch office connections by using less expensive WAN transport options without affecting performance, security, or reliability. With the IWAN solution, traffic is dynamically routed based on application service-level agreement (SLA), endpoint type, and network conditions in order to deliver the best quality experience. The realized savings from IWAN not only pays for the infrastructure upgrades, but also frees resources for business innovation.

Figure 1 *Cisco IWAN solution components*

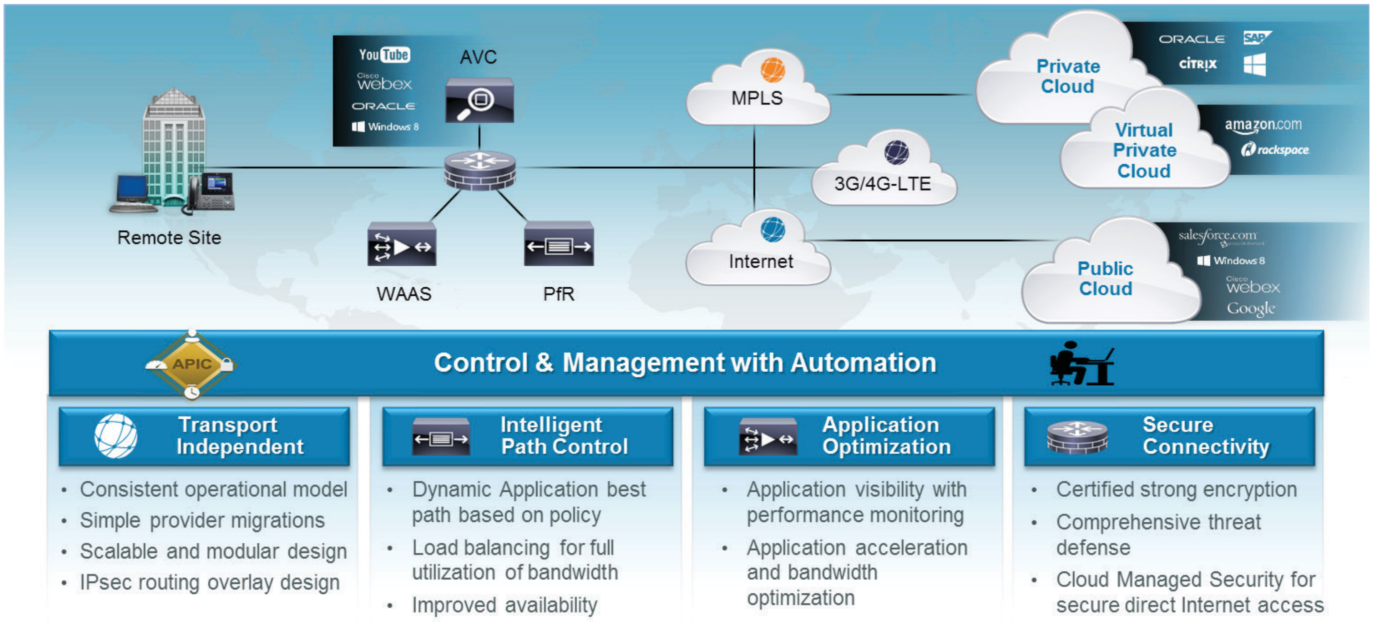
YouTube
Cisco webex
ORACLE
Windows 8
AVC
MPLS
Private Cloud
ORACLE
SAP
CITRIX
Virtual Private Cloud
amazon.com
rackspace
Remote Site
3G/4G-LTE
Internet
WAAS
PfR
Public Cloud
salesforce.com
Windows 8
Cisco webex
Google
APIC
Control & Management with Automation
Transport Independent
Consistent operational model
Simple provider migrations
Scalable and modular design
IPsec routing overlay design
Intelligent Path Control
Dynamic Application best path based on policy
Load balancing for full utilization of bandwidth
Improved availability
Application Optimization
Application visibility with performance monitoring
Application acceleration and bandwidth optimization
Secure Connectivity
Certified strong encryption
Comprehensive threat defense
Cloud Managed Security for secure direct Internet access

Transport Independence

Using DMVPN, IWAN provides capabilities for easy multi-homing over any carrier service offering, including MPLS, broadband, and cellular 3G/4G/LTE. More importantly, the design simplifies the routing design with a single routing control plane and minimal peering to providers, making it easy for organizations to mix and match and change providers and transport options. Two or more WAN transport providers are recommended in order to increase network availability up to 99.999%. Additionally, the Cisco DMVPN solution provides an industry-proven and U.S. government FIPS 140-2 certified IPsec solution for data privacy and integrity protection, as well as dynamic site-to-site DMVPN tunnels. These tunnels are encrypted using IPSec and two nodes can authenticate each other using pre-shared keys or using a public key infrastructure with a certificate authority in the demilitarized zone (DMZ) in order to enroll and authorize the use of keys between routers.

Intelligent Path Control

Cisco Performance Routing (PfR) improves application delivery and WAN efficiency. PfR dynamically controls data packet forwarding decisions by looking at application type, performance, policies, and path status. PfR monitors the network performance—jitter, packet loss, and delay—and makes decisions to forward critical applications over the best-performing path based on the defined application policy. PfR can intelligently load-balance traffic efficiently by using all available WAN bandwidth. IWAN intelligent path control is the key to providing a business-class WAN over Internet transport.

Application Optimization

Cisco Application Visibility and Control (AVC) and Cisco Wide Area Application Services (WAAS) provide application performance visibility and optimization over the WAN. With applications becoming increasingly opaque due to the increased reuse of well-known ports such as HTTP (port 80), static port classification of applications is no longer sufficient. Cisco AVC provides application awareness with deep packet inspection of traffic in order to identify and monitor applications' performance. Cisco AVC allows IT to determine what traffic is running across the network, tune the network for business-critical services, and resolve network problems. With increased visibility into the applications on the network, better QoS and PfR policies can be enabled to help ensure that critical applications are properly prioritized across the network. Cisco WAAS provides application-specific acceleration capabilities that improve response times while reducing WAN bandwidth requirements.

Secure Connectivity

Secure connectivity protects the corporate communications and offloads user traffic directly to the Internet. Strong IPsec encryption, zone-based firewalls, and strict access controls are used to protect the WAN over the public Internet. Routing remote-site users directly to the Internet improves public cloud application performance while reducing traffic over the WAN. Cisco Cloud Web Security service provides a cloud-based web proxy to centrally manage and secure user traffic accessing the Internet.

DESIGN OVERVIEW

The *Cisco Intelligent WAN Design Guide* provides a design that enables highly available, secure, and optimized connectivity for multiple remote-site local area networks (LANs).

Transport-Independent WAN Design

A transport-independent design simplifies the WAN deployment by using a GRE/IPsec VPN overlay over all WAN transport options including MPLS, Internet, and Cellular (3G/4G). A single VPN overlay reduces routing and security complexity, and provides flexibility in choosing providers and transport options. Cisco DMVPN provides the IWAN IPsec overlay.

DMVPN makes use of multipoint generic routing encapsulation (mGRE) tunnels to interconnect the hubs and all of the spoke routers. These mGRE tunnel networks are also sometimes referred to as DMVPN clouds in this context. This technology combination supports unicast, multicast, and broadcast IP, including the ability to run routing protocols within the tunnels.

Internet as WAN Transport

The Internet is essentially a large-scale public IP WAN composed of multiple interconnected service providers. The Internet can provide reliable high-performance connectivity between various locations, although it lacks any explicit guarantees for these connections. Despite its “best effort” nature, the Internet is a sensible choice for augmenting premium MPLS VPN transports or as a primary WAN transport in some cases. The IWAN architecture leverages two or more providers for resiliency and application availability. Provider path diversity provides the foundation for PfR to route around fluctuations in the providers’ performance.

Internet connections are typically included in discussions relevant to the Internet edge, specifically for the primary site. Remote-site routers also commonly have Internet connections but do not provide the same breadth of services using the Internet. For security and other reasons, Internet access at remote sites is often routed through the primary site.

This design guide uses both MPLS and the Internet for VPN site-to-site connections which is commonly referred to as a *hybrid WAN deployment*.

Dynamic Multipoint VPN

DMVPN is a solution for building scalable site-to-site VPNs that support a variety of applications. DMVPN is widely used for encrypted site-to-site connectivity over public or private IP networks and can be implemented on all WAN routers used in this design guide.

DMVPN was selected for the secure overlay WAN solution because DMVPN supports full mesh connectivity over any carriers transport with a simple hub-and-spoke base configuration and dynamic on-demand spoke-to-spoke tunnels. DMVPN also supports spoke routers that have dynamically assigned IP addresses.

Ethernet

The WAN transports mentioned previously use Ethernet as a standard media type. Ethernet is becoming a dominant carrier handoff in many markets and it is relevant to include Ethernet as the primary media in the tested architectures. Much of the discussion in this guide can also be applied to non-Ethernet media (such as T1/E1, DS-3, OC-3, and so on), but they are not explicitly discussed.

WAN-Aggregation Designs

This guide describes two IWAN design models.

The first design model is the IWAN hybrid, which uses MPLS paired with Internet VPN as WAN transports. In this design model, the MPLS WAN provides bandwidth for the critical classes of services needed for key applications and can provide SLA guarantees for these applications. The second design model is the IWAN dual Internet, which uses a pair of Internet service providers to further reduce cost while maintaining a high level of resiliency for the WAN.

A third design model, the IWAN Dual MPLS, is not covered in this guide. The Dual MPLS configuration follows the same concept as dual Internet, with the exception of the WAN underlay, which uses MPLS instead of Internet for the WAN transports. The CVD principles can be applied to other transport options, such as Carrier Ethernet, as well.

Figure 2 *Cisco IWAN design models*

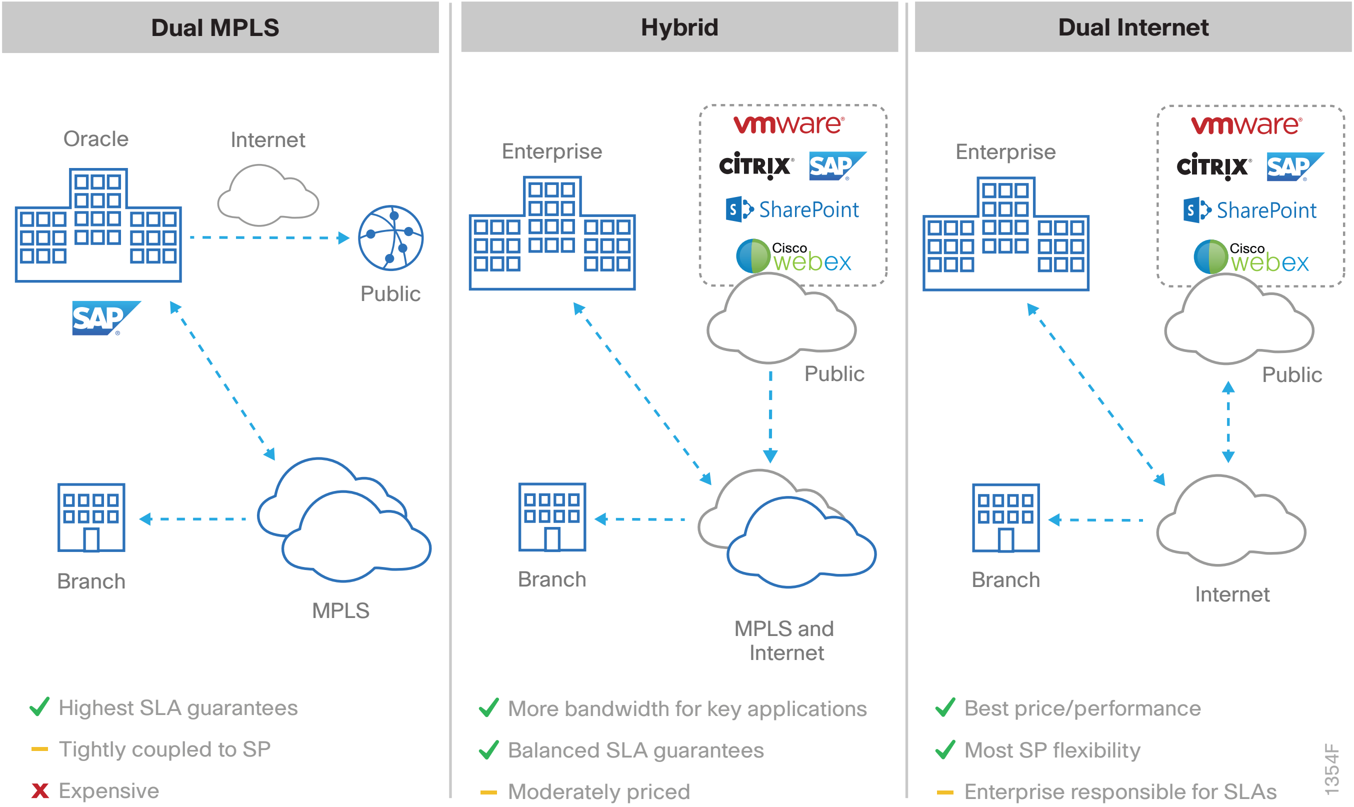

The IWAN WAN-aggregation (hub) designs for both design models include two WAN edge routers.

When WAN aggregation routers are referred to in the context of the connection to a carrier or service provider, they are typically known as *customer edge* (CE) routers. WAN aggregation routers that terminate VPN traffic are referred to as *VPN hub routers*. In the context of IWAN, an MPLS CE router is also used as a VPN hub router. Regardless of the design model, the WAN aggregation routers always connect into a pair of distribution layer switches.

Each design model is shown with LAN connections into either a collapsed core/distribution layer or a dedicated WAN distribution layer. From the WAN-aggregation perspective, there are no functional differences between these two methods.

In all of the WAN-aggregation designs, tasks such as IP route summarization are performed at the distribution layer. There are other various devices supporting WAN edge services, and these devices should also connect into the distribution layer.

The characteristics of each design are discussed in the following sections.

IWAN Hybrid Design Model

The IWAN hybrid design model:

- Has a single MPLS VPN carrier.
- Uses a single Internet carrier.
- Uses front-door virtual routing and forwarding (FVRF) on both MPLS and Internet links, with static default routing within the FVRF.
- The FVRFs provide control plane separation from and between the providers and an additional security layer between inside and outside networks.

Figure 3 WAN aggregation: IWAN hybrid design model

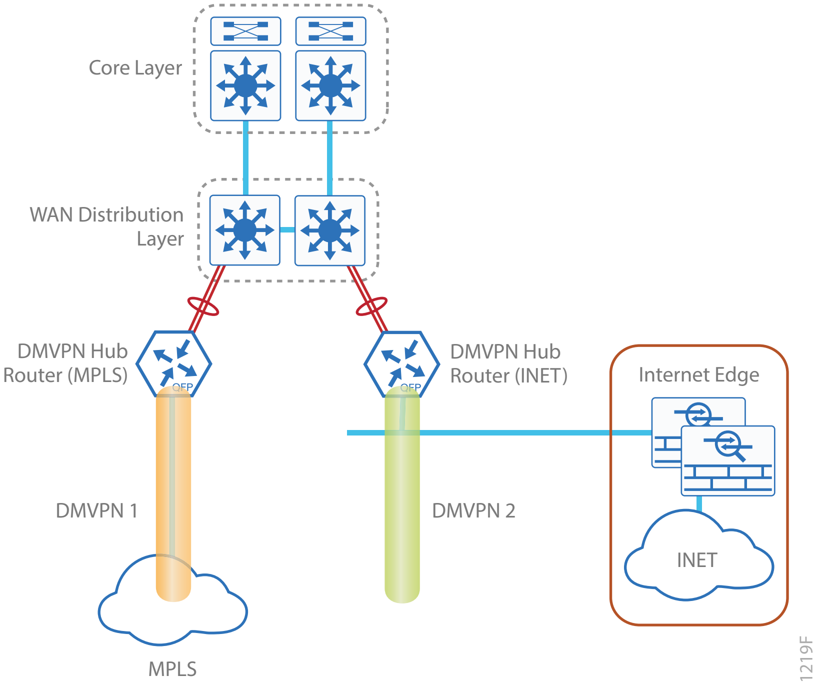

IWAN Dual Internet Design Model

The IWAN dual Internet design model:

- Uses two Internet carriers.
- Uses FVRF on both Internet links, with static default routing within the FVRF.

Figure 4 WAN aggregation: IWAN dual Internet design model

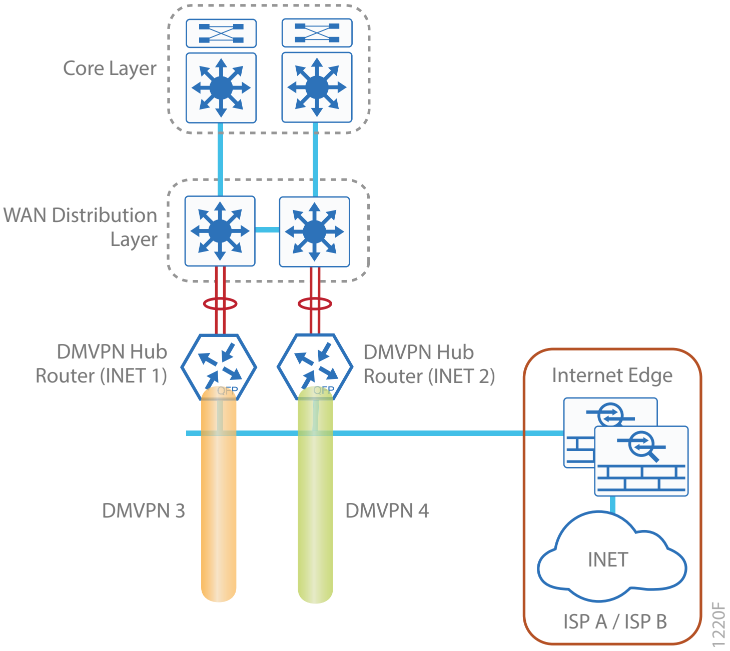

In both the IWAN hybrid and IWAN dual Internet design models, the DMVPN hub routers connect to the Internet indirectly through a firewall DMZ interface contained within the Internet edge. For details about the connection to the Internet, see the Firewall and IPS Technology Design Guide. The VPN hub routers are connected into the firewall DMZ interface, rather than connected directly with Internet service-provider routers. A firewall connection is typically not used when the VPN hub router connects to a MPLS carrier.

WAN Remote-Site Designs

This guide documents multiple WAN remote-site designs, and they are based on various combinations of WAN transports mapped to the site specific requirements for service levels and redundancy.

Figure 5 *WAN remote-site design options*

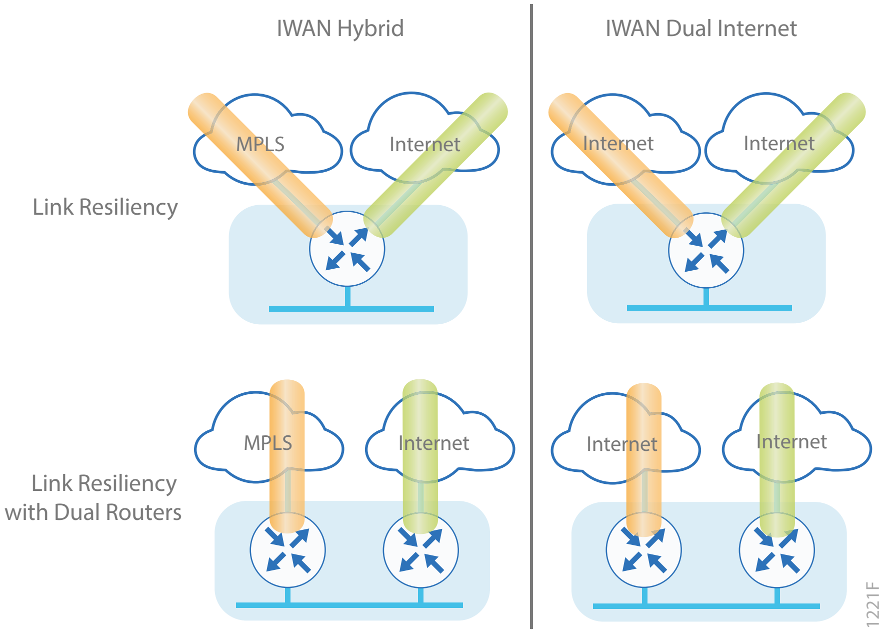

The remote-site designs include single or dual WAN edge routers. The remote-site routers are DMVPN spokes to the primary site DMVPN hubs.

Most remote sites are designed with a single router WAN edge; however, certain remote-site types require a dual router WAN design. Dual router candidate sites include regional office or remote campus locations with large user populations or sites with business critical needs that justify additional redundancy to remove single points of failure. Any individual transport can only have a single connection into a spoke site, whether or not single or dual spoke routers are used at that site.

The overall WAN design methodology is based on a primary WAN-aggregation site design that can accommodate all of the remote-site types that map to the various link combinations listed in the following table.

Table 1 *WAN remote-site transport options*

WAN remote- site routers	WAN transports	Primary transport	Secondary transport
Single	Dual	MPLS VPN	Internet
Dual	Dual	MPLS VPN	Internet
Single	Dual	Internet	Internet
Dual	Dual	Internet	Internet

This design guide also includes information for adding an LTE fallback DMVPN for a single-router remote site.

Table 2 WAN remote-site transport options with LTE fallback

WAN remote- site routers	WAN transports	Primary transport	Secondary transport	Tertiary transport
Single	Dual w/ fallback	MPLS VPN	Internet	4G LTE
Single	Dual w/ fallback	Internet	Internet	4G LTE

The modular nature of the IWAN network design enables you to create design elements that can be replicated throughout the network.

The WAN-aggregation designs and all of the WAN remote-site designs are standard building blocks in the overall design. Replication of the individual building blocks provides an easy way to scale the network and allows for a consistent deployment method.

WAN/LAN Interconnection

The primary role of the WAN is to interconnect primary site and remote-site LANs. The LAN discussion within this guide is limited to how the WAN-aggregation site LAN connects to the WAN-aggregation devices and how the remote-site LANs connect to the remote-site WAN devices. Specific details regarding the LAN components of the design are covered in the Campus LAN Layer 2 Access with Simplified Distribution Deployment Guide.

At remote sites, the LAN topology depends on the number of connected users and physical geography of the site. Large sites may require the use of a distribution layer to support multiple access layer switches. Other sites may only require an access layer switch directly connected to the WAN remote-site routers. The variants that are tested and documented in this guide are shown in the following table.

Table 3 WAN remote-site LAN options

WAN remote-site routers	WAN transports	LAN topology
Single	Dual	Access only Distribution/Access
Dual	Dual	Access only Distribution/Access

WAN Remotes Sites–LAN Topology

For consistency and modularity, all WAN remote sites use the same virtual LAN (VLAN) assignment scheme, which is shown in the following table. This design guide uses a convention that is relevant to any location that has a single access switch and this model can also be easily scaled to additional access closets through the addition of a distribution layer.

Table 4 WAN remote-sites: VLAN assignment

VLAN	Usage	Layer 2 access	Layer 3 distribution/access
VLAN 64	Data 1	Yes	–
VLAN 69	Voice 1	Yes	–
VLAN 99	Transit	Yes (dual router only)	Yes (dual router only)
VLAN 50	Router Link (1)	–	Yes
VLAN 54	Router Link (2)	–	Yes (dual router only)

Layer 2 Access

WAN remote sites that do not require additional distribution layer routing devices are considered to be flat or from a LAN perspective they are considered un-routed Layer 2 sites. All Layer 3 services are provided by the attached WAN routers. The access switches, through the use of multiple VLANs, can support services such as data and voice. The design shown in the following figure illustrates the standardized VLAN assignment scheme. The benefits of this design are clear: all of the access switches can be configured identically, regardless of the number of sites in this configuration.

Access switches and their configuration are not included in this guide. The Campus LAN Layer 2 Access with Simplified Distribution Deployment Guide provides configuration details on the various access switching platforms.

IP subnets are assigned on a per-VLAN basis. This design only allocates subnets with a 255.255.255.0 netmask for the access layer, even if less than 254 IP addresses are required. (This model can be adjusted as necessary to other IP address schemes.) The connection between the router and the access switch must be configured for 802.1Q VLAN trunking with sub-interfaces on the router that map to the respective VLANs on the switch. The various router sub-interfaces act as the IP default gateways for each of the IP subnet and VLAN combinations.

Figure 6 *WAN remote site with flat layer 2 LAN (single router)*

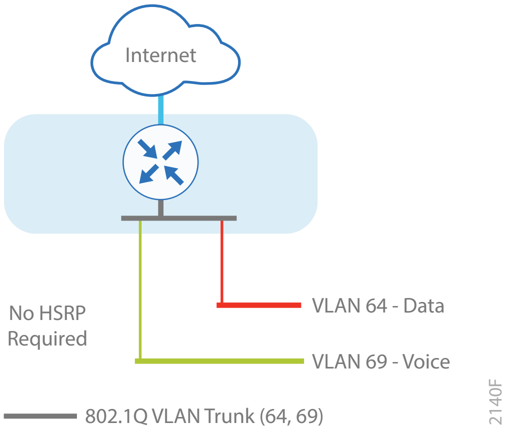

A similar LAN design can be extended to a dual-router edge as shown in the following figure. This design change introduces some additional complexity. The first requirement is to run a routing protocol. You need to configure enhanced interior gateway routing protocol (EIGRP) between the routers.

Because there are now two routers per subnet, a first-hop redundancy protocol (FHRP) must be implemented. For this design, Cisco selected hot standby router protocol (HSRP) as the FHRP. HSRP is designed to allow for transparent failover of the first-hop IP router. HSRP provides high network availability by providing first-hop routing redundancy for IP hosts configured with a default gateway IP address. HSRP is used in a group of routers for selecting an active router and a standby router. When there are multiple routers on a LAN, the active router forwards the packets; the standby router is the router that takes over when the active router fails or when preset conditions are met.

Figure 7 *WAN remote site with flat layer 2 LAN (dual router)*

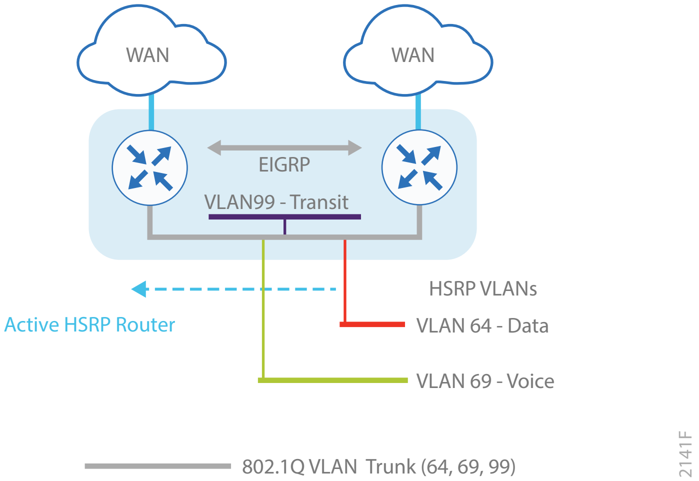

Enhanced object tracking (EOT) provides a consistent methodology for various router and switching features to conditionally modify their operation based on information objects available within other processes. The objects that can be tracked include interface line protocol, IP route reachability, and IP SLA reachability, as well as several others.

To improve convergence times after a primary WAN failure, HSRP has the capability to monitor the line-protocol status of the DMVPN tunnel interface. This capability allows for a router to give up its HSRP Active role if its DMVPN hub becomes unresponsive, and that provides additional network resiliency.

HSRP is configured to be active on the router with the highest priority WAN transport. EOT of the primary DMVPN tunnel is implemented in conjunction with HSRP so that in the case of WAN transport failure, the standby HSRP router associated with the lower priority (alternate) WAN transport becomes the active HSRP router.

The dual router designs also warrant an additional component that is required for proper routing in certain scenarios. In these cases, a traffic flow from a remote-site host might be sent to a destination reachable via the alternate WAN transport (for example, a dual DMVPN remote site communicating with a DMVPN2-only remote site). The primary WAN transport router then forwards the traffic out the same data interface to send it to the alternate WAN transport router, which then forwards the traffic to the proper destination. This is referred to as *hairpinning*.

The appropriate method to avoid sending the traffic out the same interface is to introduce an additional link between the routers and designate the link as a transit network (Vlan 99). There are no hosts connected to the transit network, and it is only used for router-router communication. The routing protocol runs between router sub-interfaces assigned to the transit network. No additional router interfaces are required with this design modification because the 802.1Q VLAN trunk configuration can easily accommodate an additional sub-interface.

Distribution and Access Layer

Large remote sites may require a LAN environment similar to that of a small campus LAN that includes a distribution layer and access layer. This topology works well with either a single or dual router WAN edge. To implement this design, the routers should connect via EtherChannel links to the distribution switch. These EtherChannel links are configured as 802.1Q VLAN trunks, to support both a routed point-to-point link to allow EIGRP routing with the distribution switch, and in the dual router design, to provide a transit network for direct communication between the WAN routers.

Figure 8 *IWAN single router remote-site: Connection to distribution layer*

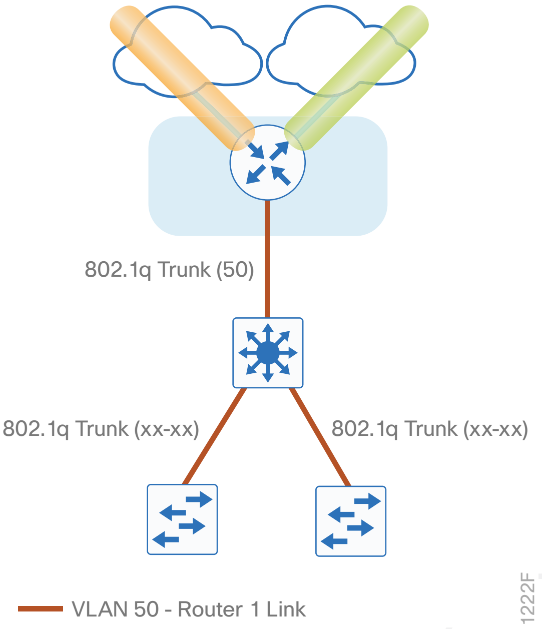

Figure 9 *IWAN dual router remote-site: Connection to distribution layer*

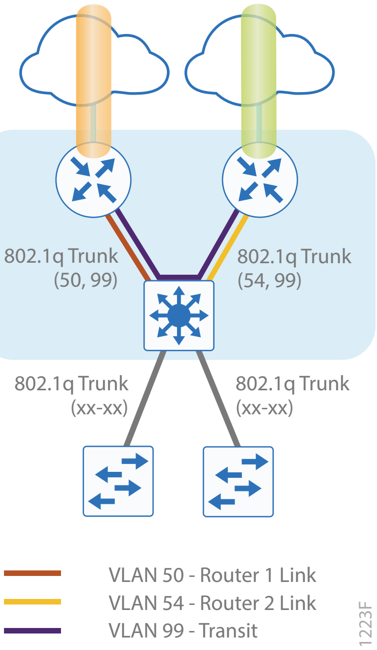

The distribution switch handles all access layer routing, with VLANs trunked to access switches. No HSRP is required when the design includes a distribution layer. A full distribution and access layer design is shown in the following figure.

Figure 10 *IWAN dual router remote-site: Distribution and access layer*

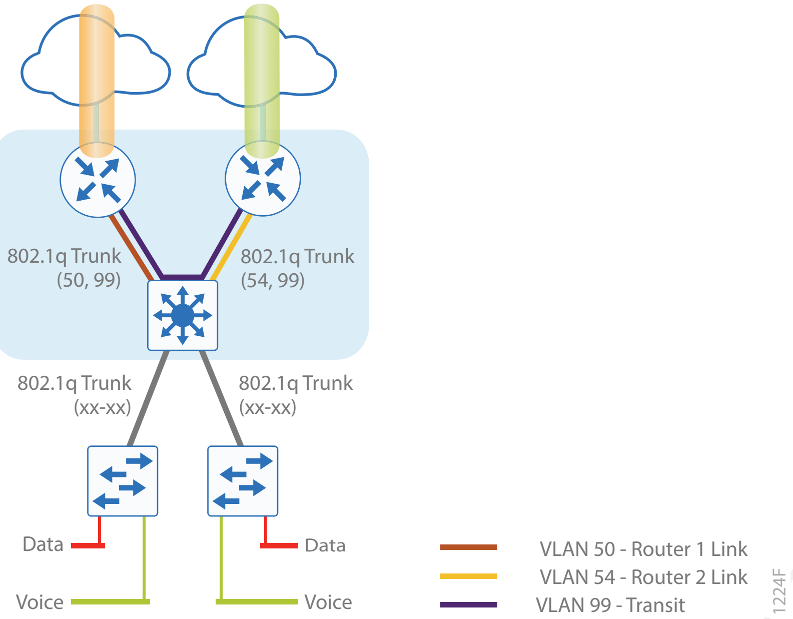

IP Multicast

IP Multicast allows a single IP data stream to be replicated by the infrastructure (routers and switches) and sent from a single source to multiple receivers. IP Multicast is much more efficient than multiple individual unicast streams or a broadcast stream that would propagate everywhere. IP telephony music on hold (MOH) and IP video broadcast streaming are two examples of IP Multicast applications.

To receive a particular IP Multicast data stream, end hosts must join a multicast group by sending an Internet group management protocol message to their local multicast router. In a traditional IP Multicast design, the local router consults another router in the network acting as a rendezvous point (RP). An RP maps the receivers to active sources so the end hosts can join their streams.

The RP is a control-plane operation that should be placed in the core of the network or close to the IP Multicast sources on a pair of Layer 3 switches or routers. IP Multicast routing begins at the distribution layer if the access layer is Layer 2 and provides connectivity to the IP Multicast RP. In designs without a core layer, the distribution layer performs the RP function. The RP must reside inside the DMVPN hub site because IP Multicast is supported only between the spoke and hub routers.

This design is fully enabled for a single global scope deployment of IP Multicast. The design uses an Anycast RP implementation strategy. This strategy provides load sharing and redundancy in protocol-independent multicast sparse mode (PIM SM) networks. Two RPs share the load for source registration and the ability to act as hot backup routers for each other.

The benefit of this strategy from the WAN perspective is that all IP routing devices within the WAN use an identical configuration referencing the Anycast RPs. IP PIM-SM is enabled on all interfaces including loopbacks, VLANs and sub-interfaces.

Quality of Service

Most users perceive the network as just a transport utility mechanism to shift data from point A to point B as fast as it can. Many sum this up as just “speeds and feeds.” While it is true that IP networks forward traffic on a best-effort basis by default, this type of routing only works well for applications that adapt gracefully to variations in latency, jitter, and loss. However networks are multiservice by design and support real-time voice and video as well as data traffic. The difference is that real-time applications require packets to be delivered within the specified delay, jitter, and loss parameters.

In reality, the network affects all traffic flows and must be aware of end-user requirements and services being offered. Even with unlimited bandwidth, time-sensitive applications are affected by jitter, delay, and packet loss. Quality of service (QoS) enables a multitude of user services and applications to coexist on the same network.

Within the architecture, there are connectivity options that provide advanced classification, prioritizing, queuing, and congestion-avoidance as part of the integrated QoS in order to help ensure optimal use of network resources. This functionality allows for the differentiation of applications, ensuring that each has the appropriate share of the network resources to protect the user experience and ensure the consistent operations of business critical applications.

QoS is an essential function of the network infrastructure devices used throughout this architecture. QoS enables a multitude of user services and applications, including real-time voice, high-quality video, and delay-sensitive data to coexist on the same network. In order for the network to provide predictable, measurable, and sometimes guaranteed services, it must manage bandwidth, delay, jitter, and loss parameters.

There are twelve common service classes that are grouped together based on interface speed, available queues, and device capabilities. The treatment of the twelve classes can be adjusted according to the policies of your organization. Cisco recommends marking your traffic in a granular manner to make it easier to make the appropriate queuing decisions at different places in the network. The goal of this design is to allow you to enable voice, video, critical data applications, bulk data applications and management traffic on the network, either during the initial deployment or later, with minimal system impact and engineering effort.

The twelve mappings in the following table are applied throughout this design by using an eight-class model in the enterprise and a six-class model in the service provider network.

Table 5 QoS service class mappings

Service class	Per-hop-behavior (PHB)	Differentiated services code point (DSCP)	Application examples
Network control	CS6	48	EIGRP, OSPF, BGP, HSRP, IKE
VoIP telephony	EF	46	Cisco IP Phones (G.711, G.729)
Call signaling	CS3	24	SCCP, SIP, H.323
Multimedia conferencing	AF4	34, 36, 38	Cisco TelePresence, Jabber, UC Video, WebEx
Real-time interactive	CS4	32	Cisco TelePresence (previous)
Multimedia streaming	AF3	26, 28, 30	Cisco Digital Media System (VoDs)
Broadcast video	CS5	40	Cisco IP Video Surveillance/Cisco Enterprise TV
Transactional data	AF2	18, 20, 22	ERP Apps, CRM Apps, Database Apps
Operation, administration, and maintenance (OAM)	CS2	16	SNMP, SSH, Syslog
Bulk data	AF1	10, 12, 14	E-mail, FTP, Backup Apps, Content Distribution
Default “best effort”	DF	0	Default class
Scavenger	CS1	8	YouTube, iTunes, BitTorrent (unnecessary for business operations)

Per-Tunnel QoS for DMVPN

The Per-Tunnel QoS for DMVPN feature allows the configuration of a QoS policy on a DMVPN hub to be applied dynamically on a per-tunnel (spoke) basis, as the spokes register with the hub. This feature allows you to apply a QoS policy on a tunnel instance (per-endpoint or per-spoke basis) in the egress direction for DMVPN hub-to-spoke tunnels. The QoS policy on a tunnel instance allows you to shape the tunnel traffic to individual spokes (parent policy) and to differentiate between traffic classes within the tunnel for appropriate treatment (child policy).

With simplified configurations, the hub site is prevented from sending more traffic than any single remote-site can handle. This ensures high bandwidth hub sites do not overrun remote-sites with lower bandwidth allocations.

Intelligent Path Control

Intelligent path control improves application delivery and WAN efficiency using PfR. PfR uses policies to dynamically control data packet forwarding by looking at application type, performance, and path status. PfR continuously monitors the network performance for jitter, packet loss and delay, and then it makes decisions to forward critical applications over the best performing path based on the application policy. PfR can evenly distribute traffic to maintain equivalent link utilization levels by using an advanced load balancing technique, even over links with differing bandwidth capacities.

Cisco PfR consists of border routers (BRs) that connect to the DMVPN overlay networks for each carrier network and a master controller (MC) application process that enforces policy. The BR collects traffic and path information and sends it to the MC at each site. The MC and BR can be configured on separate routers or the same router as shown in the figures below.

Figure 11 Cisco Performance Routing: Hub location

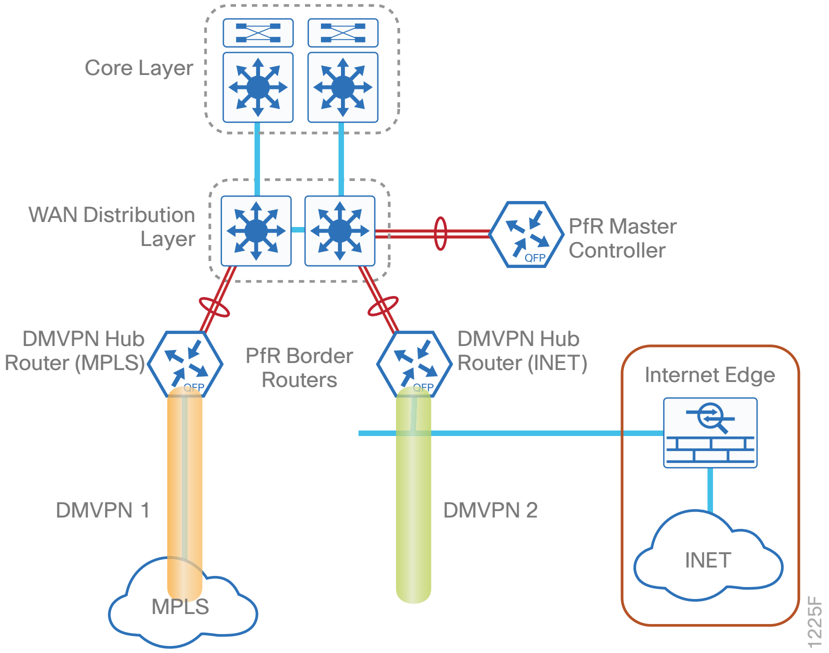

Figure 12 Cisco Performance Routing: Remote site options

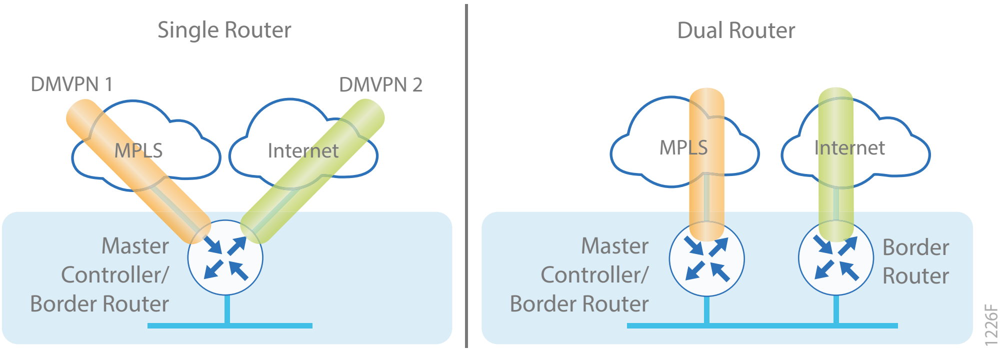

IWAN intelligent path control is the key to providing a business-class WAN over an Internet transport.

Deploying the Cisco Intelligent WAN

OVERALL IWAN ARCHITECTURE DESIGN GOALS

Overlay Transport (DMVPN)

All remote-site traffic must be encrypted when transported over public IP networks such as the Internet. This design also encrypts traffic over private IP networks such as MPLS and 4G LTE. It is recommended that you enable encryption on DMVPN over all paths in order to ensure consistency in data privacy and operations.

The use of encryption should not limit the performance or availability of a remote-site application and should be transparent to end users.

IP Routing (EIGRP)

The design has the following IP routing goals:

- Provide optimal routing connectivity from primary WAN-aggregation sites to all remote locations
- Isolate WAN routing topology changes from other portions of the network
- Ensure active/standby symmetric routing when multiple paths exist, for ease of troubleshooting and to prevent oversubscription of IP telephony call admission control (CAC) limits
- Provide a solid underlying IP routed topology in order to support the Intelligent Path Control provided by Cisco Performance Routing.
- Provide site-site remote routing via the primary WAN-aggregation site (hub-and-spoke model)
- Permit dynamic optimal direct site-to-site remote routing provided by NHRP (spoke-to-spoke model)
- Support IP Multicast sourced from the primary WAN-aggregation site

At the WAN remote sites, there is no local Internet access for web browsing or cloud services. This model is referred to as a *centralized Internet model*. It is worth noting that sites with Internet/DMVPN could potentially provide local Internet capability; however, for this design, only encrypted traffic to other DMVPN sites is permitted to use the Internet link. In the centralized Internet model, a default route is advertised to the WAN remote sites in addition to the internal routes from the data center and campus.

The use of local Internet access is covered separately from this guide.

The network must tolerate single failure conditions including the failure of any single WAN transport link or any single network device at the primary WAN-aggregation site.

Quality of Service

The network must ensure that business applications perform across the WAN during times of network congestion. Traffic must be classified and queued and the WAN connection must be shaped to operate within the capabilities of the connection. When the WAN design uses a service provider offering with QoS, the WAN edge QoS classification and treatment must align to the service provider in order to ensure consistent end-to-end QoS treatment of traffic.

Path Optimization (Performance Routing)

The network must protect business critical applications from fluctuating WAN performance by using the best-performing path based on the application policy. The design must also intelligently load-balance traffic in order to reduce an organization's overall communications expenses by allowing them to use a less expensive Internet transport without negatively affecting their mission critical traffic.

Remote sites classified as single-router, dual-link must be able to tolerate the loss of either WAN transport. Remote sites classified as dual-router, dual-link must be able to tolerate the loss of either an edge router or a WAN transport.

LAN Access

All remote sites support both wired and wireless LAN access.

Design Parameters

This design guide uses certain standard design parameters and references various network infrastructure services that are not located within the WAN. These parameters are listed in the following table.

Table 6 Universal design parameters

Network service	IP address
Domain name	cisco.local
Active Directory, DNS server, DHCP server	10.4.48.10
Cisco Secure Access Control System (ACS)	10.4.48.15
Network Time Protocol (NTP) server	10.4.48.17

Deploying the Transport Independent Design

DESIGN OVERVIEW

DMVPN Hub Routers

The most critical devices are the WAN routers that are responsible for reliable IP forwarding and QoS. The amount of bandwidth required at each site determines which model of router to use. The IWAN hybrid and dual Internet designs both require dual WAN aggregation routers to support the pair of DMVPN clouds that are required in order to provide resilient connections to all of the remote sites.

Cisco ASR 1000 Series Aggregation Services Routers represent the next-generation, modular, services-integrated Cisco routing platform. They are specifically designed for WAN aggregation, with the flexibility to support a wide range of packet-forwarding capabilities, system bandwidth performance, and scaling.

The Cisco ASR 1000 Series is fully modular from both hardware and software perspectives, and the routers have all the elements of a true carrier-class routing product that serves both enterprise and service-provider networks.

This design uses the following routers as DMVPN hub routers:

- Cisco 1002X Aggregation Services Router
- Cisco 1001X Aggregation Services Router
- Cisco 4451 Integrated Services Router

Remote Sites—DMVPN Spoke Router Selection

The actual WAN remote-site routing platforms remain unspecified because the specification is tied closely to the bandwidth required for a location and the potential requirement for the use of service module slots. The ability to implement this solution with a variety of potential router choices is one of the benefits of a modular design approach.

There are many factors to consider in the selection of the WAN remote-site routers. Among those, and key to the initial deployment, is the ability to process the expected amount and type of traffic. You also need to make sure that you have enough interfaces, enough module slots, and a properly licensed Cisco IOS Software image that supports the set of features that is required by the topology.

The DMVPN spoke routers at the WAN remote sites connect to the Internet directly through a router interface. More details about the security configuration of the remote-site routers connected to the Internet are discussed later in this guide. The single link DMVPN remote site is the most basic of building blocks for any remote location.

The first DMVPN connection is the primary WAN transport. You can add the second DMVPN link to an existing DMVPN single-link design in order to provide the resiliency either connecting on the same router or on an additional router. By adding an additional link, you provide the first level of high availability for the remote site. A failure in the primary link can be automatically detected by the router and traffic can be rerouted to the secondary path. It is mandatory to run dynamic routing when there are multiple paths. The routing protocols are tuned to ensure the proper path selection.

The dual-router, dual-link design continues to improve upon the level of high availability for the site. This design can tolerate the loss of the primary router and traffic can be rerouted via the secondary router (through the alternate path).

Figure 13 IWAN remote-site dual-link

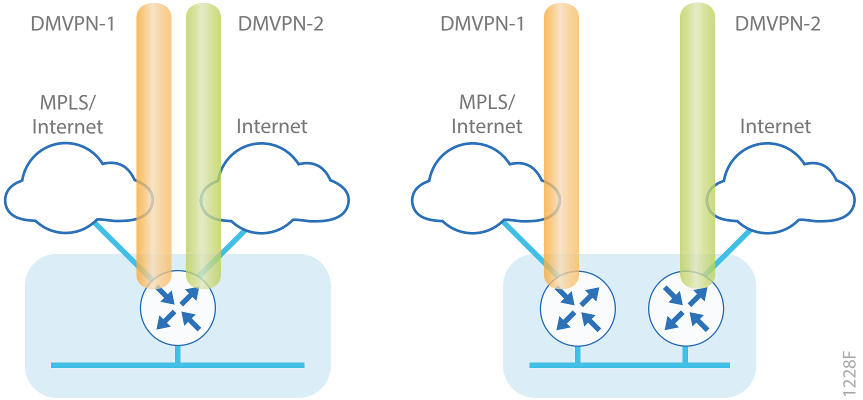

VRFs and Front Door VRF

Virtual route forwarding (VRF) is a technology used in computer networks that allows multiple instances of a routing table to co-exist within the same router at the same time. Because the routing instances are independent, you can use the same or overlapping IP Addresses without conflicting with each other. Often in an L3 VPN context, VRF is also defined as VPN Route Forwarding.

IWAN uses VRF to provide the following:

- Default route separation between user traffic and DMVPN tunnel establishment
- Control and data plane separation between inside and outside networks for security purposes

You can implement VRF in a network device by having distinct routing tables, also known as Forwarding Information Bases, one per VRF.

The simplest form of VRF implementation is VRF Lite. In this implementation, each router within the network participates in the virtual routing environment on a peer-by-peer basis. VRF Lite configurations are only locally significant.

The IP routing policy used in this design guide for the WAN remote sites does not allow direct Internet access for web browsing or other uses; any remote-site hosts that access the Internet must do so via the Internet edge at the primary site. The end hosts require a default route for all external and Internet destinations; however, this route must force traffic across the primary or secondary WAN transport DMVPN tunnels. DMVPN also has a default route requirement to establish tunnels between sites. The default route for the user traffic over DMVPN conflicts with the default route needed for DMVPN in order to establish tunnels between sites.

The multiple default route conundrum is solved through the use of VRFs on the router. A router can have multiple routing tables that are kept logically separate on the device. This separation is similar to a virtual router from the forwarding plane perspective. The global VRF corresponds to the traditional routing table, and additional VRFs are given names and route descriptors (RDs). Certain features on the router are VRF aware, including static routing and routing protocols, interface forwarding and IPsec tunneling. This set of features is used in conjunction with DMVPN to permit the use of multiple default routes for both the DMVPN hub routers and DMVPN spoke routers. This design uses global VRF for user traffic routing and a VRF for each WAN physical interface for DMVPN tunnel establishment. This combination of features is referred to as *FVRF*, because the VRF faces the WAN and the router internal LAN and DMVPN tunnel interfaces all remain in the global VRF. For more technical details regarding FVRF, see "Appendix B: Technical Feature Supplement."

Figure 14 *Front door VRF (FVRF)*

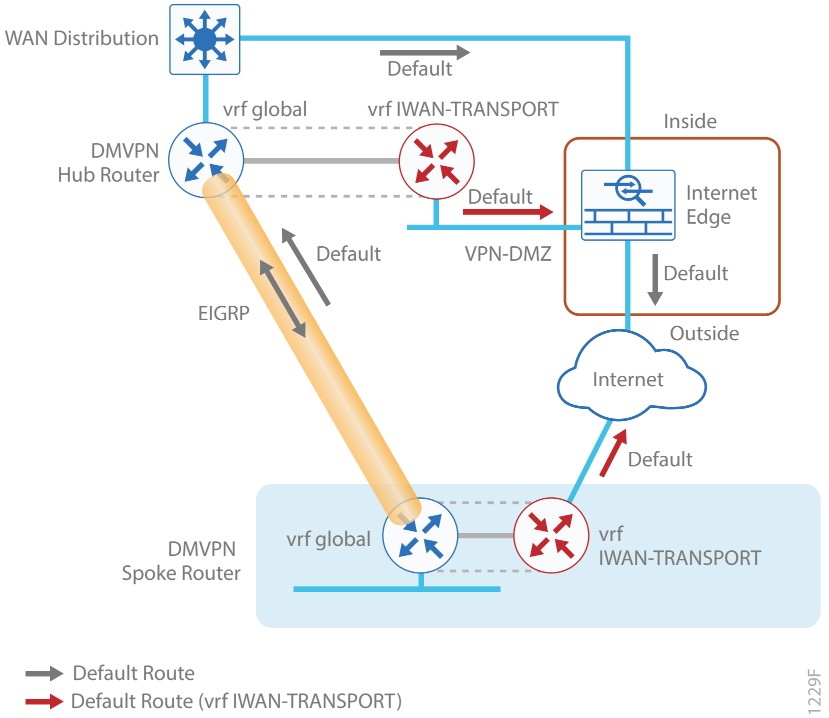

Design Details

In both the IWAN hybrid and IWAN dual Internet design models, the DMVPN hub routers must have sufficient IP-routing information in order to provide end-to-end reachability. Maintaining this routing information typically requires a routing protocol, and EIGRP or BGP are recommended for this purpose. However, a single end-to-end process named *IWAN-EIGRP* is used in this design guide for the WAN. You can also use Internal BGP (iBGP) as an overlay routing protocol and design guidance for iBGP will be added in the future.

At the WAN-aggregation site, you must connect the DMVPN hub routers to the WAN and configure default routing to build the DMVPN tunnels. The MPLS VPN hub uses default routing to the MPLS provider edge (PE) router, and the Internet VPN hubs use default routing to the DMZ-VPN that provides Internet connectivity. The DMVPN hub routers use FVRF and have a static default route with the IWAN-TRANSPORT VRF pointing to their respective next hops.

Figure 15 *IWAN hybrid design model: FVRF default routing*

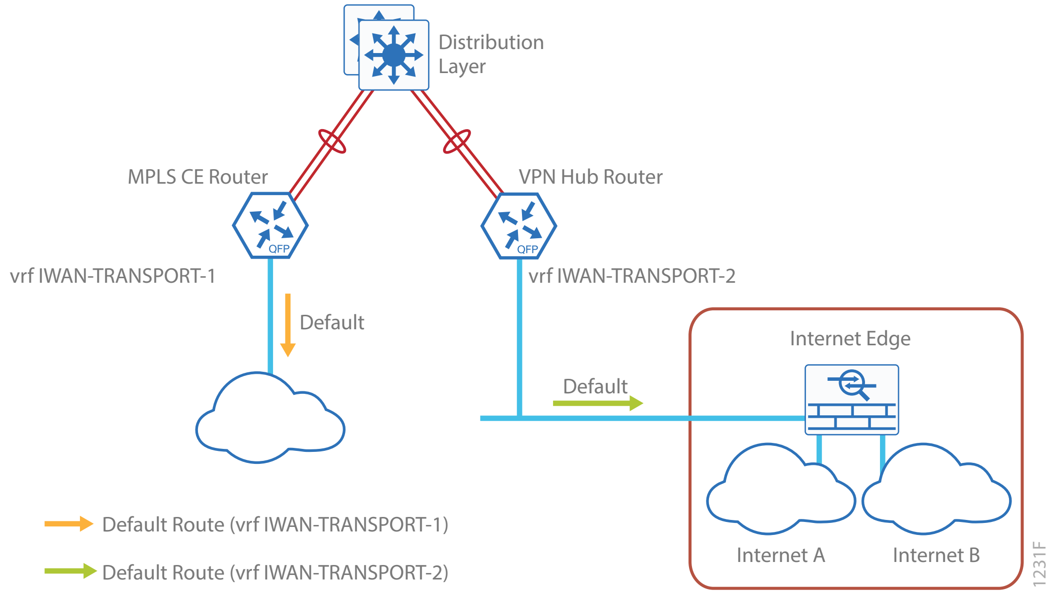

Figure 16 *IWAN dual Internet design model: FVRF default routing*

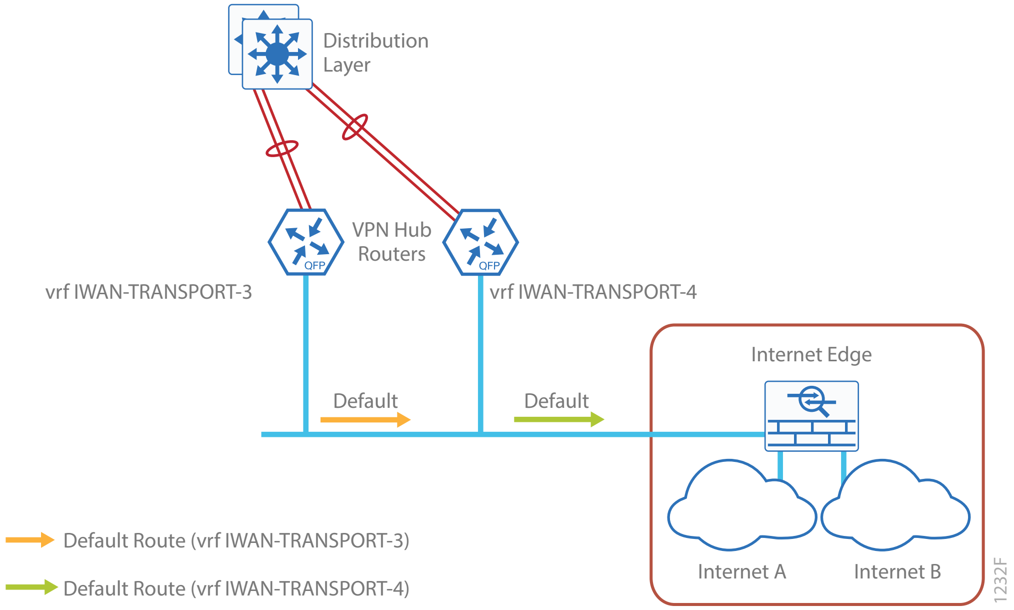

EIGRP

Cisco uses EIGRP as the primary routing protocol because it is easy to configure, does not require a large amount of planning, has flexible summarization and filtering, and can scale to large networks. As networks grow, the number of IP prefixes or routes in the routing tables grows as well. You should program IP summarization on links where logical boundaries exist, like distribution layer links to the wide area or to a core. By performing IP summarization, you can reduce the amount of bandwidth, processor, and memory necessary to carry large route tables, as well as reduce convergence time associated with a link failure.

With the advances in EIGRP, this guide uses EIGRP named mode. The use of named mode EIGRP allows related EIGRP configurations to be centrally located in the configuration. Named mode EIGRP includes features such as wide metrics, supporting larger multi-gigabit links. For added security, EIGRP neighbor authentication has been implemented to prevent unauthorized neighbor associations.

Tech Tip

With EIGRP named mode configuration, EIGRP Wide Metric support is on by default and backward compatible with existing routes.

In this design, the primary EIGRP process (AS 400) is referred to as *IWAN-EIGRP* and uses EIGRP named configuration.

The IWAN-EIGRP process is configured in the WAN-aggregation site in order to connect to the primary site LAN distribution layer, across the DMVPN tunnels and at all WAN remote sites, including those with distribution-layer LAN topologies.

Encryption

The primary goal of encryption is to provide data confidentiality, integrity, and authenticity by encrypting IP packets as the data travels across a network.

The encrypted payloads are then encapsulated with a new header (or multiple headers) and transmitted across the network. The additional headers introduce a certain amount of overhead to the overall packet length. The following table highlights the packet overhead associated with encryption based on the additional headers required for various combinations of IPsec and generic routing encapsulation (GRE).

Table 7   Overhead associated with IPsec and GRE

Encapsulation	Overhead
GRE only	24 bytes
IPsec (Transport Mode)	36 bytes
IPsec (Tunnel Mode)	52 bytes
IPsec (Transport Mode)+GRE	60 bytes
IPsec (Tunnel Mode)+GRE	76 bytes

There is a maximum transmission unit (MTU) parameter for every link in an IP network and typically the MTU is 1500 bytes. IP packets larger than 1500 bytes must be fragmented when transmitted across these links. Fragmentation is not desirable and can impact network performance. To avoid fragmentation, the original packet size plus overhead must be 1500 bytes or less, which means that the sender must reduce the original packet size. To account for other potential overhead, Cisco recommends that you configure tunnel interfaces with a 1400 byte MTU using the command **ip mtu 1400**.

There are dynamic methods for network clients to discover the path MTU, which allow the clients to reduce the size of packets they transmit. However, in many cases, these dynamic methods are unsuccessful, typically because security devices filter the necessary discovery traffic. This failure to discover the path MTU drives the need for a method that can reliably inform network clients of the appropriate packet size. The solution is to implement the **ip tcp adjust mss [size]** command on the WAN routers, which influences the TCP maximum segment size (MSS) value reported by end hosts.

The MSS defines the maximum amount of data that a host is willing to accept in a single TCP/IP datagram. The MSS value is sent as a TCP header option only in TCP SYN segments. Each side of a TCP connection reports its MSS value to the other side. The sending host is required to limit the size of data in a single TCP segment to a value less than or equal to the MSS reported by the receiving host.

The IP and TCP headers combine for 40 bytes of overhead, so the typical MSS value reported by network clients will be 1460 for standard 1500 byte MTU links. This design includes encrypted tunnels with a 1400 byte MTU, so the MSS used by endpoints should be configured to be 1360 to minimize any impact of fragmentation. In this solution, you implement the **ip tcp adjust mss 1360** command on all WAN facing tunnel interfaces.

IPsec security association (SA) anti-replay is a security service in which the decrypting router can reject duplicate packets and protect itself against replay attacks. Cisco QoS gives priority to high-priority packets. This prioritization may cause some low-priority packets to be discarded by IPsec anti-replay. Cisco IOS provides IPsec anti-replay to protect against an attacker duplicating encrypted packets. By expanding the IPsec anti-replay window you can allow the router to keep track of more than the default of 64 packets and reduce the chance that low-priority packets that are delayed by QoS are dropped by IPsec anti-replay.. In the IWAN solution, you implement the **crypto ipsec security-association replay window-size 1024** command in order to increase the window size on all DMVPN routers to the maximum setting.

IPsec uses a key exchange between the routers in order to encrypt/decrypt the traffic. You can set up secure authentication for exchanging these keys by using pre-shared keys or PKI certificates with a certificate authority. You can deploy IOS-CA in order to enroll, store, authenticate and distribute the keys to routers that request them. If a certificate authority is chosen, the certificates and keys can be distributed using the simple certificate enrollment protocol (SCEP) for automated certificate retrieval by the routers.

DMVPN

To address data security and privacy concerns, all IWAN traffic will be tunneled and encrypted using DMVPN.

All use cases in the Cisco IWAN design are dual-link. The dual-link use cases require a DMVPN dual-cloud design, each with a single hub router. Multiple DMVPN hub routers are supported, but the current version of PfR supports only a single hub router per link. The DMVPN routers use tunnel interfaces that support IP unicast as well as IP multicast and broadcast traffic, including the use of dynamic routing protocols. After the initial spoke-to-hub tunnel is active, it is possible to create dynamic spoke-to-spoke tunnels triggered by site-to-site (spoke-to-spoke) IP traffic flows.

The information required by a spoke to set up dynamic spoke-to-spoke tunnels and properly resolve other spokes is provided through the next-hop resolution protocol (NHRP) within DMVPN. Spoke-to-spoke tunnels allow for the optimal routing of traffic between locations without the indirect or hair-pin forwarding through the hub. Idle spoke-to-spoke tunnels gracefully time out after a period of inactivity.

It is common for a firewall to be placed between the DMVPN hub routers and the Internet. In many cases, the firewall may provide static or dynamic NAT from an internal RFC-1918 IP address (such as 192.168.146.10) to an Internet-routable IP address. The DMVPN solution works well with NAT because it uses IPsec transport mode, which is the standard mode even for spokes not behind NAT. DMVPN also supports a DMVPN hub behind static NAT. The IWAN DMVPN design requires the use of Internet Key Exchange version 2 (IKEv2) keepalives for dead peer detection (DPD), which is essential to facilitate fast reconvergence and for spoke registration to function properly in case a DMVPN hub is restarted. This design enables a spoke to detect that an encryption peer has failed and that the IKEv2 session with that peer is stale, which then allows a new one to be created. Without DPD, the IPsec security association (SA) must time out (the default is 60 minutes) and when the router cannot renegotiate a new SA, a new IKEv2 session is initiated. The IWAN design with the recommended IKEv2 DPD timers reduces this convergence time to 105 seconds.

Figure 17 *DMVPN dual-cloud*

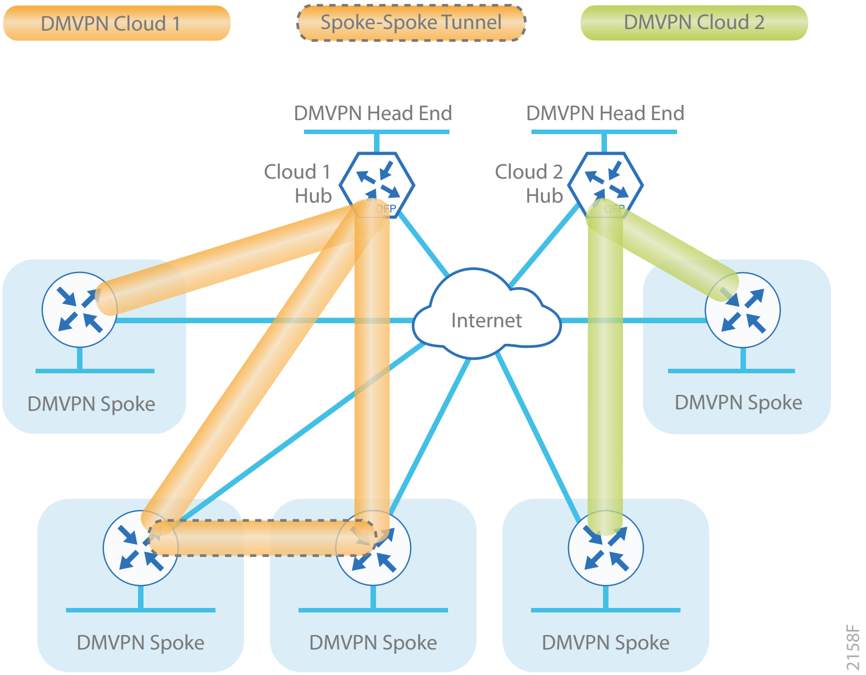

One of the key benefits of the DMVPN solution is that the spoke routers can use dynamically assigned addresses, often using DHCP from an Internet provider. The spoke routers can leverage an Internet default route for reachability to the hub routers and also other spoke addresses.

The DMVPN hub routers have static IP addresses assigned to their public-facing interfaces. This configuration is essential for proper operation as each of the spoke routers have these IP addresses embedded in their configurations.

DEPLOYMENT DETAILS

How to Read Commands

This guide uses the following conventions for commands that you enter at the command-line interface (CLI).

Commands to enter at a CLI prompt:

```
configure terminal
```

Commands that specify a value for a variable:

```
ntp server 10.10.48.17
```

Commands with variables that you must define:

```
class-map [highest class name]
```

Commands at a CLI or script prompt:

```
Router# enable
```

Long commands that line wrap are underlined. Enter them as one command:

```
police rate 10000 pps burst 10000
packets conform-action
```

Noteworthy parts of system output (or of device configuration files) are highlighted:

```
interface Vlan64
 ip address 10.5.204.5 255.255.255.0
```

The procedures in this section provide examples for most settings. The actual settings and values that you use are determined by your current network configuration. The following optional process is used for both the IWAN hybrid design model and the IWAN dual Internet design model.

PROCESS

Configuring an IOS Certificate Authority

1. Configure the IOS CA platform
2. Configure connectivity to the network
3. Configure certificate authority

Use this optional process if you want to deploy an IOS Certificate Authority (IOS CA) on a router with access from the internal network. Skip this process if you are using pre-shared keys or if you plan to use a different certificate authority. You can create a more complex CA environment, but the same basic reachability principles will apply for an IWAN enabled solution.

For this process, you configure an IOS CA by using a single internal LAN interface, which allows access from the hub routers and the remote sites. The remote sites access the IOS CA for authentication and for obtaining their certificate after a DMVPN tunnel has been established with pre-shared keys.

After the remote site has obtained the PKI certificate, it can be re-configured to use PKI certificates instead of PSKs for IKEV2 authentication to build the DMVPN encrypted tunnels with the hubs and other remote sites. It no longer needs direct access to the IOS CA, except when it needs a new PKI certificate before the current PKI certificate times out (usually at least 1 year lifetime). At that point, the remote site will use the current DMVPN tunnel to access the IOS CA for the new PKI certificate.

Figure 18 *IOS CA with internal LAN interface*

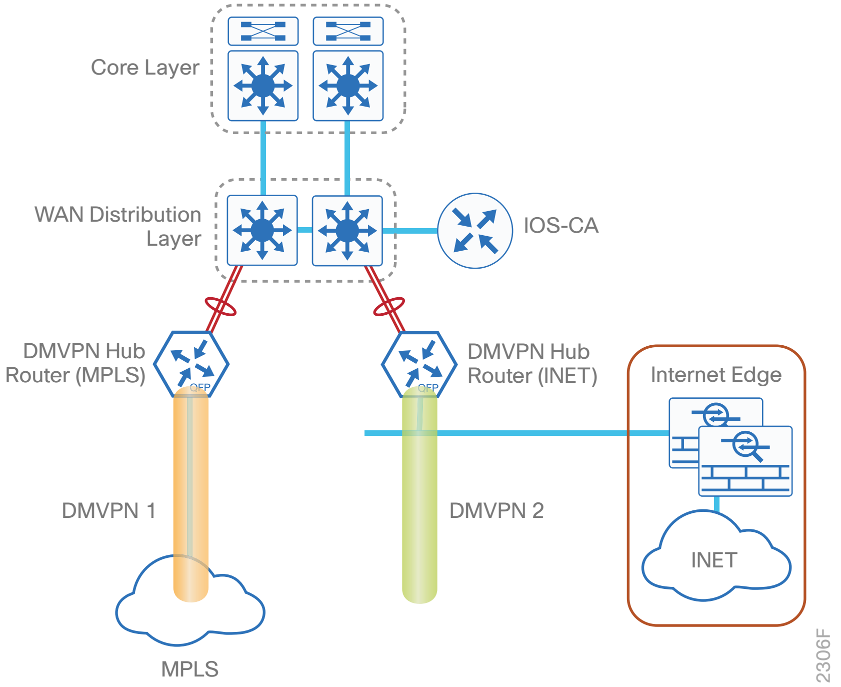

Procedure 1 Configure the IOS CA platform

To complete the base configuration for this router, follow the steps in “Configure the platform base features” in Appendix C.

Procedure 2 Configure connectivity to the network

Step 1: The internal address is an inside address that can be accessed from the hub site or a remote site if the site is already up and running with a DMVPN tunnel.

```
interface GigabitEthernet0/0
 description Internal
 ip address 10.6.24.11 255.255.255.224
 no shutdown
```

Step 2: Configure IP routing using a static route.

```
ip route 0.0.0.0 0.0.0.0 10.6.24.1
```

Procedure 3 Configure certificate authority

The following commands configure the CA on the router. This CA can be part of a PKI hierarchy, but only of IOS authorities, and the certificate from the root CA must be issued via SCEP.

Step 1: Configure the server.

```
crypto pki server IWAN-IOS-CA
 database level complete
 no database archive
 issuer-name CN=IWAN-IOS-CA.cisco.local L=SanJose St=CA C=US
```

Step 2: Configure the server to use SCEP for issuing certificates.

```
 grant auto
```

Step 3: Configure the lifetime for the issued certificates at 2 years. The time is in days.

```
 lifetime certificate 730
```

Step 4: Configure the lifetime for the certificate server signing certificate at 3 years. The time is in days.

```
 lifetime ca-certificate 1095
```

Step 5: Configure the location for certificate revocation lists.

Tech Tip

In order to force the parser to retain the embedded question mark within the specified location, enter CTRL+V prior to the question mark. If this action is not taken, CRL retrieval through HTTP returns an error message.

```
 cdp-url http://10.6.24.11/cgi-bin/pkiclient.exe?operation=GetCRL
 database url crl nvram:
```

Step 6: Start the server with the **no shutdown** command.

```
crypto pki server IWAN-IOS-CA
 no shutdown
%Some server settings cannot be changed after CA certificate generation.
% Please enter a passphrase to protect the private key
% or type Return to exit
Password: c1sco123
```

```
Re-enter password: c1sco123
% Generating 1024 bit RSA keys, keys will be non-exportable...
[OK] (elapsed time was 2 seconds)

% Certificate Server enabled.
IWAN-IOS-CA(cs-server)#
Dec 15 13:19:49.254: %PKI-6-CS_ENABLED: Certificate server now enabled.
```

The following trustpoint and rsa keypair are automatically generated when you start the server:

```
crypto pki trustpoint IWAN-IOS-CA
 revocation-check crl
 rsakeypair IWAN-IOS-CA
```

Tech Tip

For more information, including options for configuring certificates, see the following document: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_pki/configuration/15-mt/sec-pki-15-mt-book.pdf

PROCESS

Configuring DMVPN Hub Router

1. Configure the distribution switch
2. Configure the WAN aggregation platform
3. Configure IP multicast routing
4. Configure connectivity to the LAN
5. Configure the WAN-facing VRF
6. Connect to the MPLS WAN or Internet
7. Configure IKEv2 and IPsec
8. Configure the mGRE tunnel
9. Configure EIGRP

Use this process for the both the IWAN hybrid design model and the IWAN dual Internet design model, and repeat it for each DMVPN hub router.

Table 8 DMVPN hub router IP addresses

DMVPN cloud	Hostname	Loopback IP address	Port channel IP address
Hybrid–Primary WAN	VPN-MPLS-ASR1002X-1	10.6.32.241/32	10.6.32.2/30
Hybrid–Secondary WAN	VPN-INET-4451X-2	10.6.32.242/32	10.6.32.6/30
Dual Internet–Primary WAN	VPN-INET-ASR1002X-3	10.6.32.243/32	10.6.32.18/30
Dual Internet–Secondary WAN	VPN-INET-ASR1002X-4	10.6.32.244/32	10.6.32.22/30

Procedure 1 Configure the distribution switch

Reader Tip

This process assumes that the distribution switch has already been configured following the guidance in the Campus LAN Layer 2 Access with Simplified Distribution Deployment Guide. Only the procedures required to support the integration of the WAN aggregation router into the deployment are included.

The LAN distribution switch is the path to the organization's main campus and data center. A Layer 3 port-channel interface connects to the distribution switch to the WAN aggregation router and the internal routing protocol peers across this interface.

Tech Tip

As a best practice, use the same channel numbering on both sides of the link where possible.

Step 1: Configure the Layer 3 port-channel interface and assign the IP address.

```
interface Port-channel1
 description VPN-MPLS-ASR1002X-1
 no switchport
 ip address 10.6.32.1 255.255.255.252
 ip pim sparse-mode
 load-interval 30
 no shutdown
```

Step 2: Configure EtherChannel member interfaces.

Configure the physical interfaces to tie to the logical port-channel using the **channel-group** command. The number for the port-channel and channel-group must match. Not all router platforms can support link aggregation control protocol (LACP) in order to negotiate with the switch, so to keep the design consistent across the network, EtherChannel is configured statically, which also reduces startup times.

Also, apply the egress QoS macro that was defined in the platform configuration procedure in order to ensure traffic is prioritized appropriately.

```
interface GigabitEthernet1/0/1
 description VPN-MPLS-ASR1002X-1 Gig0/0/0

interface GigabitEthernet2/0/1
 description VPN-MPLS-ASR1002X-1 Gig0/0/1

interface range GigabitEthernet1/0/1, GigabitEthernet2/0/1
  no switchport
  channel-group 1 mode on
  logging event link-status
  logging event trunk-status
  logging event bundle-status
  load-interval 30
  no shutdown
  macro apply EgressQoS
```

Step 3: Allow the routing protocol to form neighbor relationships across the port channel interface.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Port-channel1
   no passive-interface
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
 exit-address-family
```

Step 4: If you had previously configured EIGRP stub routing on your WAN distribution switch, disable the feature.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
 no eigrp stub
 exit-address-family
```

Step 5: On the distribution layer switch, configure the Layer 3 interface connected to the LAN core to summarize the WAN network ranges.

Tech Tip

It is a best practice to summarize IP routes from the WAN distribution layer towards the core.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Port-channel38
   summary-address 10.6.32.0 255.255.248.0
   summary-address 10.7.0.0 255.255.0.0
   summary-address 10.255.240.0 255.255.248.0
  exit-af-interface
 exit-address-family
```

Procedure 2 Configure the WAN aggregation platform

Within this design, there are features and services that are common across all WAN aggregation routers. These are system settings that simplify and secure the management of the solution.

To complete the base configuration for this router, follow the steps in "Configure the platform base features" in Appendix C.

Procedure 3 Configure IP multicast routing

Optional

This optional procedure includes additional steps for configuring IP Multicast on a router. Skip this procedure if you do not want to use IP Multicast in your environment.

In this design, which is based on sparse mode multicast operation, Auto RP is used to provide a simple yet scalable way to provide a highly resilient RP environment.

Step 1: Enable IP Multicast routing on the platform in the global configuration mode.

```
ip multicast-routing
```

Step 2: The Cisco ASR1000 series and ISR4000 series routers require the **distributed** keyword.

```
ip multicast-routing distributed
```

Step 3: Configure every Layer 3 switch and router to discover the IP Multicast RP with autorp. Use the **ip pim autorp listener** command to allow for discovery across sparse mode links. This configuration provides for future scaling and control of the IP Multicast environment and can change based on network needs and design.

```
ip pim autorp listener
```

Step 4: Enable sparse mode multicast operation.

```
ip pim sparse-mode
```

Procedure 4 Configure connectivity to the LAN

Step 1: Configure IP unicast routing authentication key.

```
key chain LAN-KEY
 key 1
   key-string c1sco123
```

Step 2: Configure IP unicast routing using EIGRP named mode.

EIGRP is configured facing the LAN distribution or core layer. In this design, the port-channel interface and the loopback must be EIGRP interfaces. The loopback may remain a passive interface. Passive interfaces are used to prevent accidental peering and to reduce the EIGRP traffic on a network segment. The network range must include both interface IP addresses, either in a single network statement or in multiple network statements. This design uses a best practice of assigning the router ID to a loopback address.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface default
   passive-interface
  exit-af-interface
 network 10.6.0.0 0.1.255.255
 eigrp router-id 10.6.32.241
 nsf
 exit-address-family
```

Any links to adjacent distribution layers should be Layer 3 links or Layer 3 EtherChannels.

Step 3: Enable QoS support for port-channel interfaces.

```
platform qos port-channel-aggregate 1
```

Tech Tip

This only applies to ASR1000 routers. If there is a requirement to configure QoS on the port-channel interface, make sure to enable platform support before you create the port-channel interface on the router.

```
platform qos port-channel-aggregate [port-channel number]
```

If you apply this command globally for an existing port-channel-interface that already has been configured, you will receive an error:

```
"Port-channel 1 has been configured with non-aggregate mode already, please use
different interface number that port-channel interface hasn't been configured"
```

If you need to apply a QoS policy to an existing port-channel interface, you must first delete the existing port-channel interface and configure platform support for that port-channel interface number.

Step 4: Configure Layer 3 port-channel interface.

At the hub location where there are multiple border routers, the interface throughput delay setting should be set to influence the routing protocol path preference. Set the internal LAN path to 250000 microseconds (usec). The delay command is entered in 10 usec units.

```
interface Port-channel1
  ip address 10.6.32.2 255.255.255.252
  ip pim sparse-mode
  delay 25000
  no shutdown
```

Step 5: Configure EtherChannel member interfaces.

Configure the physical interfaces to tie to the logical port-channel using the channel-group command. The number for the port-channel and channel-group must match.

```
interface GigabitEthernet0/0/0
 description IW-WAN-D3750X Gig1/0/1

interface GigabitEthernet0/0/1
 description IW-WAN-D3750X Gig2/0/1

interface range GigabitEthernet0/0/0, GigabitEthernet0/0/1
  no ip address
  channel-group 1
  cdp enable
  no shutdown
```

Step 6: Configure the EIGRP interface.

Allow EIGRP to form neighbor relationships across the interface to establish peering adjacencies and exchange route tables. In this step, you configure EIGRP authentication by using the authentication key specified in the previous procedure.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Port-channel1
   no passive-interface
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
 exit-address-family
```

Procedure 5 Configure the WAN-facing VRF

Next, you create a WAN-facing VRF in order to support FVRF for DMVPN. The VRF name is arbitrary but it is useful to select a name that describes the VRF. The VRF must be enabled for IPv4.

Table 9 VRF assignments

IWAN design	Primary WAN VRF	Secondary WAN VRF
Hybrid	IWAN-TRANSPORT-1	IWAN-TRANSPORT-2
Dual Internet	IWAN-TRANSPORT-3	IWAN-TRANSPORT-4

This design uses VRF Lite, so the selection is only locally significant to the device. It is a best practice to use the same VRF/RD combination across multiple devices when using VRFs in a similar manner. However, this convention is not strictly required.

Step 1: Configure the primary WAN VRF.

Example: Primary WAN in IWAN hybrid design model

```
vrf definition IWAN-TRANSPORT-1
 address-family ipv4
```

Procedure 6 Connect to the MPLS WAN or Internet

Each IWAN DMVPN hub requires a connection to the WAN transport, which is either MPLS or Internet.

If you are using MPLS in this design, the DMVPN hub is connected to the service provider’s MPLS PE router. The IP addressing used between IWAN CE and MPLS PE routers must be negotiated with your MPLS carrier.

If you are using the Internet in this design, the DMVPN hub is connected through a Cisco Adaptive Security Appliance (ASA) 5500 using a DMZ interface specifically created and configured for a VPN termination router.

The IP address that you use for the Internet-facing interface of the DMVPN hub router must be an Internet-routable address. There are two possible methods for accomplishing this task:

- Assign a routable IP address directly to the router.
- Assign a non-routable RFC-1918 address directly to the router and use a static NAT on the Cisco ASA 5500 to translate the router IP address to a routable IP address.

This design assumes that the Cisco ASA 5500 is configured for static NAT for the DMVPN hub router.

Option 1: MPLS WAN physical WAN interface

The DMVPN design is using FVRF, so you must place the WAN interface into the VRF configured in the previous procedure.

Step 1: Enable the interface, select the VRF, and assign the IP address.

```
interface GigabitEthernet0/0/3
 vrf forwarding IWAN-TRANSPORT-1
 ip address 192.168.6.1 255.255.255.252
 no shutdown
```

Step 2: Configure the VRF-specific default routing.

The VRF created for FVRF must have its own default route to the MPLS. This default route points to the MPLS PE router's IP address and is used by DMVPN for tunnel establishment.

```
ip route vrf IWAN-TRANSPORT-1 0.0.0.0 0.0.0.0 192.168.6.2
```

Option 2: Internet WAN physical WAN interface

The DMVPN design is using FVRF, so you must place the WAN interface into the VRF configured in Procedure 5, "Configure the WAN-facing VRF."

Step 1: Enable the interface, select the VRF, and assign the IP address.

```
interface GigabitEthernet0/0/3
 vrf forwarding IWAN-TRANSPORT-2
 ip address 192.168.146.10 255.255.255.0
 no shutdown
```

Step 2: Configure the VRF-specific default routing.

The VRF created for FVRF must have its own default route to the Internet. This default route points to the Cisco ASA 5500's DMZ interface IP address.

```
ip route vrf IWAN-TRANSPORT-2 0.0.0.0 0.0.0.0 192.168.146.1
```

Figure 19 *Physical and logical views for DMZ connection*

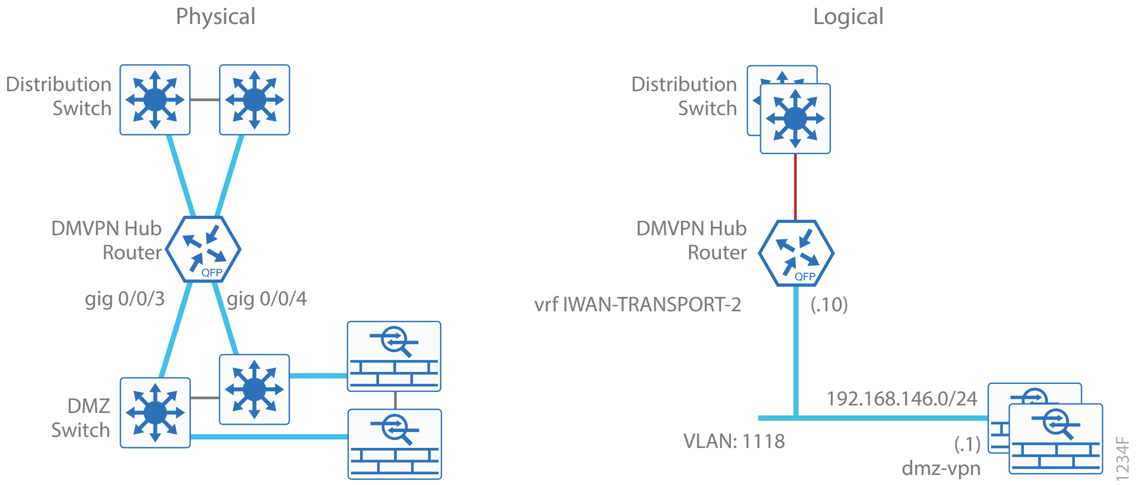

Procedure 7 Configure IKEv2 and IPsec

The parameters in the tables below are used in this procedure. Use the values in the table that represent the design model that you are configuring.

Table 10 *Crypto parameters: IWAN hybrid design model*

Parameter	Primary WAN	Secondary WAN
vrf	IWAN-TRANSPORT-1	IWAN-TRANSPORT-2
crypto ikev2 keyring	DMVPN-KEYRING-1	DMVPN-KEYRING-2
crypto ikev2 profile	FVRF-IKEv2-IWAN-TRANSPORT-1	FVRF-IKEv2-IWAN-TRANSPORT-2
crypto ipsec profile	DMVPN-PROFILE-TRANSPORT-1	DMVPN-PROFILE-TRANSPORT-2

Table 11 *Crypto parameters: IWAN dual Internet design model*

Parameter	Primary WAN	Secondary WAN
vrf	IWAN-TRANSPORT-3	IWAN-TRANSPORT-4
crypto ikev2 keyring	DMVPN-KEYRING-3	DMVPN-KEYRING-4
crypto ikev2 profile	FVRF-IKEv2-IWAN-TRANSPORT-3	FVRF-IKEv2-IWAN-TRANSPORT-4
crypto ipsec profile	DMVPN-PROFILE-TRANSPORT-3	DMVPN-PROFILE-TRANSPORT-4

IPsec uses a key exchange between the routers in order to encrypt/decrypt the traffic. These keys can be exchanged using pre-shared keys or PKI certificates with a certificate authority. It is also possible to use a combination of the two, which is useful during a migration from one method to the other.

You have two options for configuring key exchange: using pre-shared keys or using a certificate authority.

Option 1: Configure with pre-shared keys

Step 1: Configure the crypto keyring for pre-shared keys.

The crypto keyring defines a pre-shared key (or password) valid for IP sources that are reachable within a particular VRF. This key is a wildcard pre-shared key if it applies to any IP source. A wildcard key is configured using the 0.0.0.0 0.0.0.0 as the network/mask combination.

```
crypto ikev2 keyring [keyring name]
 peer ANY
  address 0.0.0.0 0.0.0.0
  pre-shared-key [password]
```

Example: Primary WAN in the IWAN hybrid design model

```
crypto ikev2 keyring DMVPN-KEYRING-1
 peer ANY
  address 0.0.0.0 0.0.0.0
  pre-shared-key c1sco123
```

Step 2: Configure the IKEv2 proposal.

The user-defined IKEv2 proposal includes only the following:

- Encryption with AES cipher with a 256-bit key
- Integrity with SHA and a 512-bit digest
- Pseudo-random function with SHA and a 512-bit digest
- Diffie-Hellman group: 14

```
crypto ikev2 proposal [proposal name]
 encryption [encryption algorithm]
 integrity [integrity hash algorithm]
 group [Diffie-Hellman group]
```

Example

```
crypto ikev2 proposal AES/CBC/256
 encryption aes-cbc-256
 integrity sha512
 group 14
```

The default IKEv2 proposal is also used.

A **show crypto ikev2 proposal** displays the details of the two proposals.

```
VPN-MPLS-ASR1002X-1# show crypto ikev2 proposal
IKEv2 proposal: AES/CBC/256
     Encryption : AES-CBC-256
     Integrity  : SHA512
     PRF        : SHA512
     DH Group   : DH_GROUP_2048_MODP/Group 14
IKEv2 proposal: default
    Encryption: AES-CBC-256 AES-CBC-192 AES-CBC-128
    Integrity : SHA512 SHA384 SHA256 SHA96 MD596
    PRF       : SHA512 SHA384 SHA256 SHA1 MD5
    DH Group  : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2
```

Step 3: Configure the IKEv2 profile.

The IKEv2 profile creates an association between an identity address, a VRF, and a crypto keyring. A wildcard address within a VRF is referenced with 0.0.0.0. The profile also defines what method of key sharing will be used on this router with **authentication local** and what methods will be accepted from remote locations with **authentication remote**. The **pre-share** keyword is used with the keyring defined above.

```
crypto ikev2 profile [profile name]
 match fvrf [vrf name]
 match identity remote address [IP address]
 authentication remote pre-share
 authentication local pre-share
 keyring local [keyring name]
```

Example: Primary WAN in the IWAN hybrid design model

```
crypto ikev2 profile FVRF-IKEv2-IWAN-TRANSPORT-1
 match fvrf IWAN-TRANSPORT-1
 match identity remote address 0.0.0.0
 authentication remote pre-share
 authentication local pre-share
 keyring local DMVPN-KEYRING-1
```

Step 4: Define the IPsec transform set.

A transform set is an acceptable combination of security protocols, algorithms, and other settings to apply to IPsec-protected traffic. Peers agree to use a particular transform set when protecting a particular data flow.

The IPsec transform set for DMVPN uses the following:

- ESP with the 256-bit AES encryption algorithm
- ESP with the SHA (HMAC variant) authentication algorithm

Because the DMVPN hub router is behind a NAT device, the IPsec transform must be configured for transport mode.

```
crypto ipsec transform-set [transform set] esp-aes 256 esp-sha-hmac
 mode transport
```

Example

```
crypto ipsec transform-set AES256/SHA/TRANSPORT esp-aes 256 esp-sha-hmac
 mode transport
```

Step 5: Configure the IPsec profile.

The IPsec profile creates an association between an IKEv2 profile and an IPsec transform-set.

```
crypto ipsec profile [profile name]
 set transform-set [transform set]
 set ikev2-profile [ikev2 profile name]
```

Example: Primary WAN in the IWAN hybrid design model

```
crypto ipsec profile DMVPN-PROFILE-TRANSPORT-1
 set transform-set AES256/SHA/TRANSPORT
 set ikev2-profile FVRF-IKEv2-IWAN-TRANSPORT-1
```

Step 6: Increase the IPsec anti-replay window size.

```
crypto ipsec security-association replay window-size [value]
```

Example

```
crypto ipsec security-association replay window-size 1024
```

Tech Tip

QoS queuing delays can cause anti-replay packet drops, so it is important to extend the window size in order to prevent the drops from occurring.

Increasing the anti-replay window size has no impact on throughput and security. The impact on memory is insignificant because only an extra 128 bytes per incoming IPsec SA is needed.

It is recommended that you use the maximum window size in order to minimize future anti-replay problems. On the ASR1K, ISR4K and ISRG2 router platforms, the maximum replay window size is 1024.

If you do not increase the window size, the router may drop packets and you may see the following error messages on the router CLI and in the log:

```
%CRYPTO-4-PKT_REPLAY_ERR: decrypt: replay check failed
```

Step 7: Proceed to Procedure 8, “Configure the mGRE tunnel.”

Option 2: Configure with a certificate authority

If you want to use a certificate authority, you will have to configure a pre-shared key on one of the hub border routers in order to allow each remote site to establish a DMVPN tunnel to the WAN aggregation site. After the first DMVPN tunnel at a remote site is established, the router will be able to authenticate to the CA and obtain a certificate. After obtaining the certificate, you can configure the remote site to use PKI.

The **crypto pki trustpoint** is the method of specifying the parameters associated with a CA. The router must authenticate to the CA first and then enroll with the CA to obtain its own identity certificate.

Step 1: The fingerprint command limits the responding CA. You can find this fingerprint by using **show crypto pki server** on the IOS CA.

```
IWAN-IOS-CA# show crypto pki server
Certificate Server IWAN-IOS-CA:
    Status: enabled
    State: enabled
    Server’s configuration is locked  (enter "shut" to unlock it)
    Issuer name: CN=IWAN-IOS-CA.cisco.local L=SanJose St=CA C=US
    CA cert fingerprint: 75BEF625 9A9876CF 6F341FE5 86D4A5D8
    Granting mode is: auto
    Last certificate issued serial number (hex): 11
    CA certificate expiration timer: 10:28:57 PST Nov 11 2017
    CRL NextUpdate timer: 09:47:47 PST Dec 4 2014
    Current primary storage dir: nvram:
    Current storage dir for .crl files: nvram:
    Database Level: Complete - all issued certs written as <serialnum>.cer
```

Step 2: Configure the PKI trust point.

```
crypto pki trustpoint [name]
 enrollment url [URL of IOS CA]
 serial-number none
 fqdn [fully qualified domain name of this router]
 ip-address [Loopback IP address of this router]
 fingerprint [fingerprint from IOS CA]
 revocation-check none
 rsakeypair [name] 2048 2048
```

Example: VPN-INET-ASR1002X-4

This example is from the secondary WAN hub router in the dual Internet design model. It can reach the IOS CA through the internal network at 10.6.24.11 using the default VRF.

```
crypto pki trustpoint IWAN-CA
 enrollment url http://10.6.24.11:80
 serial-number none
 fqdn VPN-INET-ASR1002X-4.cisco.local
 ip-address 10.6.32.244
 fingerprint 75BEF6259A9876CF6F341FE586D4A5D8
 revocation-check none
 rsakeypair IWAN-CA-KEYS 2048 2048
```

Step 3: Authenticate to the CA and obtain the CA's certificate

Exit the trustpoint configuration mode on the hub router and issue the following command to authenticate to the CA and get its certificate.

```
VPN-INET-ASR1002X-4 (config)# crypto pki authenticate IWAN-CA
Certificate has the following attributes:
       Fingerprint MD5: 75BEF625 9A9876CF 6F341FE5 86D4A5D8
      Fingerprint SHA1: 9C14D6F4 D1F08023 17A85669 52922632 C6B02928
Trustpoint Fingerprint: 75BEF625 9A9876CF 6F341FE5 86D4A5D8
Certificate validated - fingerprints matched.
```

Step 4: When the trustpoint CA certificate is accepted, enroll with the CA, enter a password for key retrieval, and obtain a certificate for this hub router.

```
VPN-INET-ASR1002X-4 (config)# crypto pki enroll IWAN-CA

% Start certificate enrollment ..
% Create a challenge password. You will need to verbally provide this
   password to the CA Administrator in order to revoke your certificate.
   For security reasons your password will not be saved in the configuration.
   Please make a note of it.

Password: c1sco123
Re-enter password: c1sco123

% The subject name in the certificate will include: VPN-INET-ASR1002X-4.cisco.
local
% The IP address in the certificate is 10.6.32.244
Request certificate from CA? [yes/no]: yes
% Certificate request sent to Certificate Authority
% The 'show crypto pki certificate verbose IWAN-CA' command will show the finger-
print.
```

Step 5: Configure the IKEv2 proposal.

The user-defined IKEv2 proposal includes only the following:

- Encryption with AES cipher with a 256-bit key
- Integrity with SHA and a 512-bit digest
- Pseudo-random function with SHA and a 512-bit digest
- Diffie-Hellman group: 14

```
crypto ikev2 proposal [proposal name]
 encryption [encryption algorithm]
 integrity [integrity hash algorithm]
 group [Diffie-Hellman group]
```

Example

```
crypto ikev2 proposal AES/CBC/256
 encryption aes-cbc-256
 integrity sha512
 group 14
```

The default IKEv2 proposal is also used.

A **show crypto ikev2 proposal** displays the details of the two proposals.

```
VPN-INET-ASR1002X-4# show crypto ikev2 proposal
IKEv2 proposal: AES/CBC/256
     Encryption : AES-CBC-256
     Integrity  : SHA512
     PRF        : SHA512
     DH Group   : DH_GROUP_2048_MODP/Group 14
IKEv2 proposal: default
    Encryption: AES-CBC-256 AES-CBC-192 AES-CBC-128
    Integrity : SHA512 SHA384 SHA256 SHA96 MD596
    PRF       : SHA512 SHA384 SHA256 SHA1 MD5
    DH Group  : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2
```

Step 6: Configure the IKEv2 profile.

The IKEv2 profile creates an association between an identity address, a VRF, and a crypto keyring. A wildcard address within a VRF is referenced with 0.0.0.0. The **identity local address** must match the crypto pki trustpoint's **ip-address** value from the step above.

Tech Tip

Use the **identity local address** in the ikev2 profile in order to avoid repeated CRYPTO-6-IKMP_NO_ID_CERT_ADDR_MATCH warning messages on the router.

The profile also defines what method of key sharing will be used on this router with **authentication local** and what methods will be accepted from remote locations with **authentication remote**. The **rsa-sig** keyword is used when certificates contain the encryption key.

```
crypto ikev2 profile [profile name]
 match fvrf [vrf name]
 match identity remote address [IP address]
 identity local address [Loopback IP address of this router]
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint [trustpoint name]
```

Example: VPN-INET-ASR1002X-4

```
crypto ikev2 profile FVRF-IKEv2-IWAN-TRANSPORT-4
 match fvrf IWAN-TRANSPORT-4
 match identity remote address 0.0.0.0
 identity local address 10.6.32.244
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint IWAN-CA
```

Step 7: Define the IPsec transform set.

A transform set is an acceptable combination of security protocols, algorithms, and other settings to apply to IPsec-protected traffic. Peers agree to use a particular transform set when protecting a particular data flow.

The IPsec transform set for DMVPN uses the following:

- ESP with the 256-bit AES encryption algorithm
- ESP with the SHA (HMAC variant) authentication algorithm

Because the DMVPN hub router is behind a NAT device, you must configure the IPsec transform for transport mode.

```
crypto ipsec transform-set [transform set] esp-aes 256 esp-sha-hmac
 mode transport
```

Example

```
crypto ipsec transform-set AES256/SHA/TRANSPORT esp-aes 256 esp-sha-hmac
 mode transport
```

Step 8: Configure the IPsec profile.

The IPsec profile creates an association between an IKEv2 profile and an IPsec transform-set.

```
crypto ipsec profile [profile name]
 set transform-set [transform set]
 set ikev2-profile [ikev2 profile name]
```

Example: Secondary WAN in the dual Internet design model

```
crypto ipsec profile DMVPN-PROFILE-TRANSPORT-4
 set transform-set AES256/SHA/TRANSPORT
 set ikev2-profile FVRF-IKEv2-IWAN-TRANSPORT-4
```

Step 9: Increase the IPsec anti-replay window size.

```
crypto ipsec security-association replay window-size [value]
```

Example

```
crypto ipsec security-association replay window-size 1024
```

Tech Tip

QoS queuing delays can cause anti-replay packet drops, so it is important to extend the window size to prevent the drops from occurring.

Increasing the anti-replay window size has no impact on throughput and security. The impact on memory is insignificant because only an extra 128 bytes per incoming IPsec SA is needed.

It is recommended that you use the maximum window size in order to minimize future anti-replay problems. On the Cisco ASR1K, ISR 4K and ISR G2 router platforms, the maximum replay window size is 1024.

If you do not increase the window size, the router may drop packets and you may see the following error message on the router CLI and in the log:

```
%CRYPTO-4-PKT_REPLAY_ERR: decrypt: replay check failed
```

Procedure 8 Configure the mGRE tunnel

The parameters in the table below are used in this procedure. Choose the rows that represent the design model that you are configuring. This procedure applies to the primary WAN hub router in the IWAN hybrid design model.

Table 12 DMVPN tunnel parameters

DMVPN cloud	Tunnel VRF	Tunnel number	Tunnel IP address	NHRP network ID/ tunnel key
Hybrid–Primary WAN	IWAN-TRANSPORT-1	10	10.6.34.1/23	101
Hybrid–Secondary WAN	IWAN-TRANSPORT-2	11	10.6.36.1/23	102
Dual Internet–Primary WAN	IWAN-TRANSPORT-3	20	10.6.38.1/23	201
Dual Internet–Secondary WAN	IWAN-TRANSPORT-4	21	10.6.40.1/23	202

Step 1: Configure the basic interface settings.

The tunnel number is arbitrary, but it is best to begin tunnel numbering at 10 or above, because other features deployed in this design may also require tunnels and they may select lower numbers by default.

The tunnel interface bandwidth setting should be set to match the bandwidth of the respective transport, which should correspond to the actual interface speed or, if you are using a subrate service, use the policed rate from the carrier.

Configure the **ip mtu** to 1400 and the **ip tcp adjust-mss** to 1360. There is a 40 byte difference that corresponds to the combined IP and TCP header length.

The tunnel interface throughput delay setting should be set to influence the routing protocol path preference. Set the primary WAN path to 10000 usec and the secondary WAN path to 200000 usec to prefer one over the other. The delay command is entered in 10 usec units.

```
interface Tunnel10
 bandwidth 1000000
 ip address 10.6.34.1 255.255.254.0
 no ip redirects
 ip mtu 1400
 ip tcp adjust-mss 1360
 delay 1000
```

Step 2: Configure the tunnel.

DMVPN uses mGRE tunnels. This type of tunnel requires a source interface only. Use the same source interface that you use to connect to the MPLS or Internet. Set the **tunnel vrf** command to the VRF defined previously for FVRF.

Enabling encryption on this interface requires that you apply the IPsec profile configured in the previous procedure.

```
interface Tunnel10
 tunnel source GigabitEthernet0/0/3
 tunnel mode gre multipoint
 tunnel key 101
 tunnel vrf IWAN-TRANSPORT-1
 tunnel protection ipsec profile DMVPN-PROFILE-TRANSPORT-1
```

Step 3: Configure NHRP.

The DMVPN hub router acts in the role of NHRP server for all of the spokes. NHRP is used by remote routers to determine the tunnel destinations for peers attached to the mGRE tunnel.

NHRP requires all devices within a DMVPN cloud to use the same network ID and authentication key. The NHRP cache holdtime should be configured to 600 seconds.

EIGRP (configured in the following procedure) relies on a multicast transport and requires NHRP to automatically add routers to the multicast NHRP mappings.

The **ip nhrp redirect** command allows the DMVPN hub to notify spoke routers that a more optimal path may exist to a destination network, which may be required for DMVPN spoke-spoke direct communications.

```
interface Tunnel10
 ip nhrp authentication cisco123
 ip nhrp map multicast dynamic
 ip nhrp network-id 101
 ip nhrp holdtime 600
 ip nhrp redirect
```

Step 4: (Optional) Enable PIM non-broadcast multiple access (NBMA) mode for the DMVPN tunnel.

Spoke-to-spoke DMVPN networks present a unique challenge because the spokes cannot directly exchange information with one another, even though they are on the same logical network. This inability to directly exchange information can also cause problems when running IP Multicast.

To resolve this issue requires a method where each remote PIM neighbor has its join messages tracked separately. A router in PIM NBMA mode treats each remote PIM neighbor as if it were connected to the router through a point-to-point link.

Tech Tip

Do not enable PIM on the Internet DMZ interface, as no multicast traffic should be requested from this interface.

```
interface Tunnel10
 ip pim nbma-mode
```

Procedure 9 Configure EIGRP

Step 1: Configure the EIGRP values for the mGRE tunnel interface.

Spoke-to-spoke DMVPN networks present a unique challenge because the spokes cannot directly exchange information with one another, even though they are on the same logical network. This limitation requires that the DMVPN hub router advertise routes from other spokes on the same network. This advertisement of these routes would normally be prevented by split horizon and can be overridden by the **no split-horizon** command.

The EIGRP hello interval is increased to 20 seconds and the EIGRP hold time is increased to 60 seconds in order to accommodate up to 2000 remote sites on a single DMVPN cloud. Increasing the EIGRP timers also slows down the routing convergence to improve network stability and the IWAN design allows PfR to initiate the fast failover, so changing the timers is recommended for all IWAN deployments.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel10
   hello-interval 20
   hold-time 60
   no passive-interface
   no split-horizon
  exit-af-interface
 exit-address-family
```

Step 2: Configure EIGRP neighbor authentication. Neighbor authentication enables the secure establishment of peering adjacencies and exchange route tables over the DMVPN tunnel interface.

```
key chain WAN-KEY
 key 1
  key-string c1sco123

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel10
   authentication mode md5
   authentication key-chain WAN-KEY
  exit-af-interface
 exit-address-family
```

Step 3: Configure EIGRP network summarization.

The IP assignments for the entire network were designed so they can be summarized within a few aggregate routes. As configured below, the **summary-address** command suppresses the more specific routes. If any network within the summary is present in the route table, the summary is advertised to the remote sites, which offers a measure of resiliency. If the various networks cannot be summarized, then EIGRP continues to advertise the specific routes.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel10
   summary-address 10.6.0.0 255.255.0.0
   summary-address 10.7.0.0 255.255.0.0
   summary-address 10.8.0.0 255.255.0.0
   summary-address 10.255.240.0 255.255.248.0
```

Step 4: Tag and filter the routes.

This design uses a single EIGRP autonomous system for the VPN and all of the WAN remote sites. Every remote site is dual-connected for resiliency. However, due to the multiple paths that exist within this topology, you must try to avoid routing loops and to prevent remote sites from becoming transit sites if WAN failures were to occur.

In this design, there are different IP subnets for each DMVPN network, and the EIGRP tags are clearly defined to help with readability and troubleshooting. When a design uses more than one data center, additional tags are required to identify the different DMVPN hub router locations.

The following logic is used to control the routing.

- Each DMVPN network will have an EIGRP route tag in order to prevent routes from being re-advertised over the other DMVPN networks at the remote sites.
- All prefixes that are advertised towards the WAN are uniquely tagged.
- All DMVPN learned WAN prefixes, except those that originate locally from a hub, are advertised towards the LAN and tagged.

Outbound distribute-lists are used to set tags on the DMVPN hub routers towards the WAN and LAN. The remote-site routers use the tags set towards the WAN in order to protect against becoming transit sites.

The following tables show specific route tags in use.

Table 13 Route tag information for IWAN hybrid hub routers

DMVPN hub	DMVPN tunnel key	Tag tunnel	Tag LAN
VPN-MPLS-AS-R1002X-1	101 (MPLS)	101 (All routes)	101 (WAN routes)
VPN-INET-4451X-2	102 (INET)	102 (All routes)	102 (WAN routes)

Table 14 Route tag information for IWAN dual Internet hub routers

DMVPN hub	DMVPN tunnel key	Tag tunnel	Tag LAN
VPN-INET-AS-R1002X-3	201 (INET1)	201 (All routes)	201 (WAN routes)
VPN-INET-AS-R1002X-4	202 (INET2)	202 (All routes)	202 (WAN routes)

The following examples show both DMVPN hub routers in the IWAN hybrid design model.

Example: VPN-MPLS-ASR1002X-1

```
route-map SET-TAG-ALL permit 10
 description Tag all routes advertised through the tunnel
 set tag 101

! All MPLS tunnel interfaces are in this IP address range
ip access-list standard DMVPN-1-SPOKES
 permit 10.6.34.0 0.0.1.255

route-map SET-TAG-DMVPN-1 permit 10
 description Tag all incoming routes advertised through LAN interface
 match ip route-source DMVPN-1-SPOKES
 set tag 101
```

```
route-map SET-TAG-DMVPN-1 permit 100
 description Advertise all other routes with no tag

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   topology base
    distribute-list route-map SET-TAG-DMVPN-1 out Port-channel1
    distribute-list route-map SET-TAG-ALL out Tunnel10
```

Example: VPN-INET-ISR4451X-2

```
route-map SET-TAG-ALL permit 10
 description tag all routes advertised through the tunnel
 set tag 102

! All INET tunnel interfaces are in this IP address range
ip access-list standard DMVPN-2-SPOKES
 permit 10.6.36.0 0.0.1.255

route-map SET-TAG-DMVPN-2 permit 10
description Tag all incoming routes advertised through LAN interface
 match ip route-source DMVPN-2-SPOKES
 set tag 102
route-map SET-TAG-DMVPN-2 permit 100
 description Advertise all other routes with no tag

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   topology base
    distribute-list route-map SET-TAG-DMVPN-2 out Port-channel2
    distribute-list route-map SET-TAG-ALL out Tunnel11
```

PROCESS

Configuring the Firewall and DMZ Switch

1. Configure the DMZ switch for DMVPN hub router
2. Configure firewall DMZ interface
3. Configure network address translation
4. Configure security policy

If necessary, configure the DMZ and firewall for the Internet WAN.

Procedure 1 Configure the DMZ switch for DMVPN hub router

Reader Tip

This procedure assumes that the switch has already been configured following the guidance in the Campus LAN Layer 2 Access with Simplified Distribution Deployment Guide. Only the procedures required to support the integration of the firewall into the deployment are included.

Step 1: Set the DMZ switch to be the spanning tree root for the VLAN that contains the DMVPN hub router.

```
vlan 1118
 name dmz-vpn

spanning-tree vlan 1118 root primary
```

Step 2: Configure the interface that is connected to the DMVPN hub routers. Repeat as necessary.

```
interface GigabitEthernet1/0/1
 description VPN-INET-4451X-2 (gig0/0/3)
 switchport access vlan 1118
 switchport host
 logging event link-status
 load-interval 30
 no shutdown
 macro apply EgressQoS
```

Step 3: Configure the interfaces that are connected to the appliances as a trunk.

```
interface GigabitEthernet1/0/48
 description IE-ASA5545Xa Gig0/1

interface GigabitEthernet2/0/48
 description IE-ASA5545Xb Gig0/1

interface range GigabitEthernet1/0/48, GigabitEthernet2/0/48
 switchport trunk allowed vlan add 1118
 switchport mode trunk
 logging event link-status
 logging event trunk-status
 load-interval 30
 no shutdown
 macro apply EgressQoS
```

Procedure 2 Configure firewall DMZ interface

The firewall’s DMZ is a portion of the network where, typically, traffic to and from other parts of the network is tightly restricted. Organizations place network services in a DMZ for exposure to the Internet. These servers are typically not allowed to initiate connections to the ‘inside’ network, except for specific circumstances.

The DMZ network is connected to the appliances on the appliances’ GigabitEthernet interface via a VLAN trunk to allow the greatest flexibility if new VLANs must be added to connect additional DMZs. The trunk connects the appliances to a 2960X access-switch stack to provide resiliency. The DMZ VLAN interfaces on the Cisco ASA are each assigned an IP address, which will be the default gateway for each of the VLAN subnets. The DMZ switch only offers Layer 2 switching capability; the DMZ switch’s VLAN interfaces do not have an IP address assigned, save for one VLAN interface with an IP address for management of the switch.

Figure 20 *DMZ VLAN topology and services*

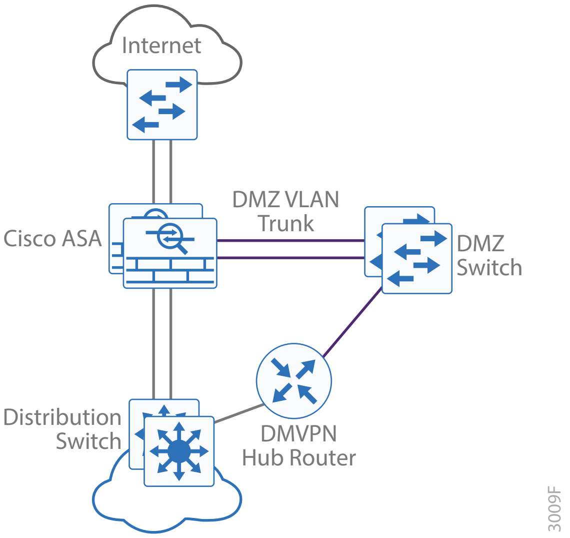

Tech Tip

By setting the DMZ connectivity as a VLAN trunk, you get the greatest flexibility.

Step 1: In **Configuration > Device Setup > Interfaces**, click the interface that is connected to the DMZ switch. (Example: GigabitEthernet0/1)

Step 2: Click **Edit**.

Step 3: Select **Enable Interface**, and then click **OK**.

Step 4: On the Interface pane, click **Add > Interface**.

Step 5: In the **Hardware Port** list choose the interface configured in Step 1. (Example: GigabitEthernet0/1)

Step 6: In the **VLAN ID** box, enter the VLAN number for the DMZ VLAN. (Example: 1118)

Step 7: In the **Subinterface ID** box, enter the VLAN number for the DMZ VLAN. (Example: 1118)

Step 8: Enter an **Interface Name**. (Example: dmz-dmvpn)

Step 9: In the **Security Level** box, enter a value of **75**.

Step 10: Enter the interface **IP Address**. (Example: 192.168.146.1)

Step 11: Enter the interface **Subnet Mask**, and then click **OK**. (Example: 255.255.255.0)

Step 12: Click **Apply**.

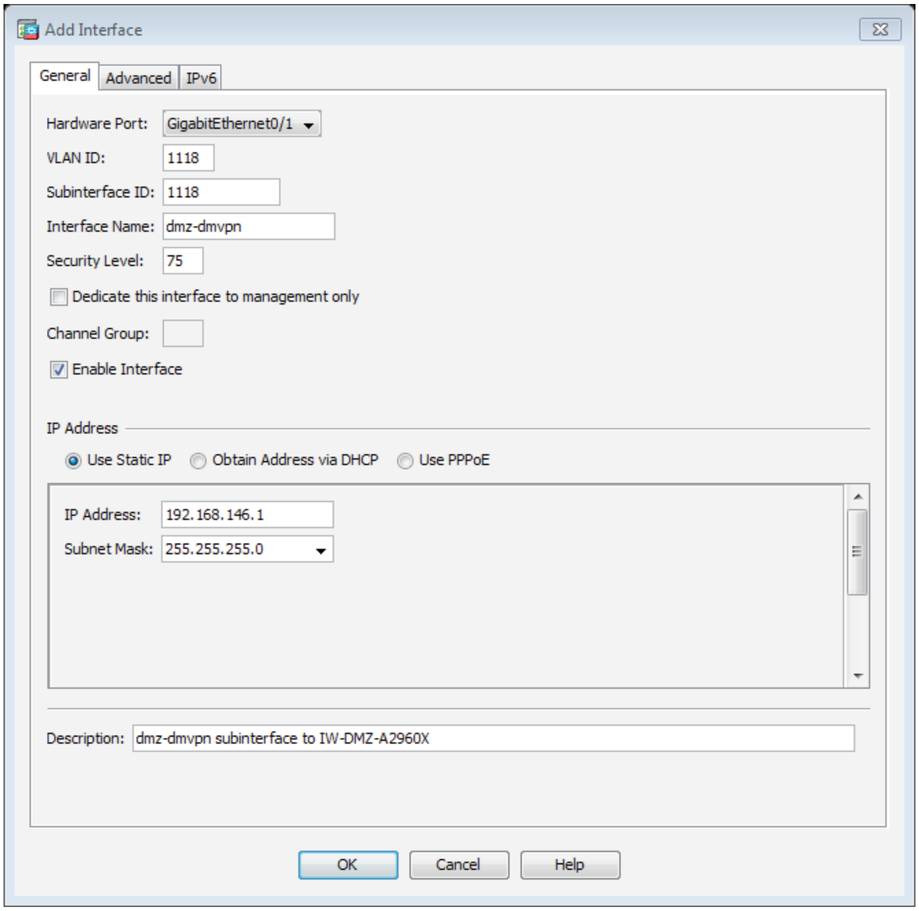

Hardware Port: GigabitEthernet0/1
VLAN ID: 1118
Subinterface ID: 1118
Interface Name: dmz-dmvpn
Security Level: 75
☐ Dedicate this interface to management only
Channel Group:
☑ Enable Interface

IP Address
● Use Static IP ○ Obtain Address via DHCP ○ Use PPPoE
IP Address: 192.168.146.1
Subnet Mask: 255.255.255.0

Description: dmz-dmvpn subinterface to IW-DMZ-A2960X

Step 13: In **Configuration > Device Management > High Availability >** click **Failover**.

Step 14: On the **Interfaces** tab, for the interface created in Step 4, enter the IP address of the standby unit in the **Standby IP address** column. (Example: 192.168.146.2)

Step 15: Select **Monitored**.

Step 16: Click **Apply**.

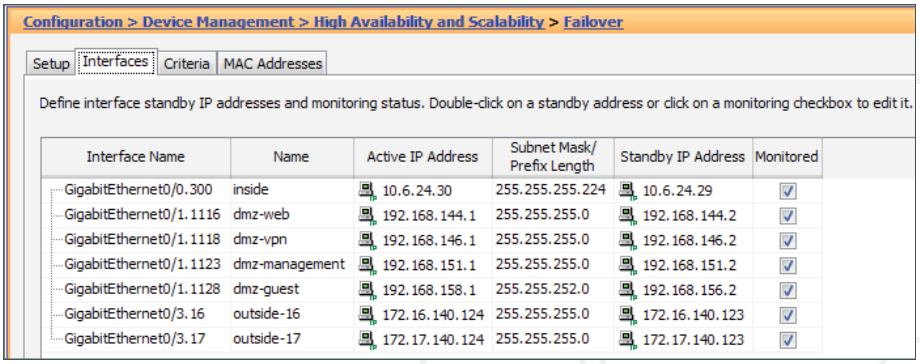
Configuration > Device Management > High Availability and Scalability > Failover
Setup
Interfaces
Criteria
MAC Addresses

Define interface standby IP addresses and monitoring status. Double-click on a standby address or click on a monitoring checkbox to edit it.

Interface Name	Name	Active IP Address	Subnet Mask/ Prefix Length	Standby IP Address	Monitored
GigabitEthernet0/0.300	inside	10.6.24.30	255.255.255.224	10.6.24.29	☑
GigabitEthernet0/1.1116	dmz-web	192.168.144.1	255.255.255.0	192.168.144.2	☑
GigabitEthernet0/1.1118	dmz-vpn	192.168.146.1	255.255.255.0	192.168.146.2	☑
GigabitEthernet0/1.1123	dmz-management	192.168.151.1	255.255.255.0	192.168.151.2	☑
GigabitEthernet0/1.1128	dmz-guest	192.168.158.1	255.255.252.0	192.168.156.2	☑
GigabitEthernet0/3.16	outside-16	172.16.140.124	255.255.255.0	172.16.140.123	☑
GigabitEthernet0/3.17	outside-17	172.17.140.124	255.255.255.0	172.17.140.123	☑

Procedure 3 Configure network address translation

The DMZ network uses private network (RFC 1918) addressing that is not Internet routable, so the firewall must translate the DMZ address of the DMVPN hub router to an outside public address.

The example DMZ address to public IP address mapping is shown in the following table.

Table 15 DMVPN NAT address mapping

DMVPN	DMVPN hub router DMZ address	DMVPN hub router public address (externally routable after NAT)
IWAN-TRANSPORT-2	192.168.146.10	172.16.140.1 (ISP-A)
IWAN-TRANSPORT-3	192.168.146.20	172.16.140.11 (ISP-A)
IWAN-TRANSPORT-4	192.168.146.21	172.17.140.11 (ISP-B)

First, to simplify the configuration of the security policy, you create the External DMZ network objects that are used in the firewall policies.

Table 16 External DMZ firewall network objects

Network object name	Object type	IP address	Description
outside-dmvpn-2-ISPa	Host	172.16.140.1	DMVPN hub router 2 on ISP A (outside)
outside-dmvpn-3-ISPa	Host	172.16.140.11	DMVPN hub router 3 on ISP A (outside)
outside-dmvpn-4-ISPb	Host	172.17.140.11	DMVPN hub router 4 on ISP B (outside)

Step 1: Navigate to **Configuration > Firewall > Objects > Network Objects/Groups**.

Step 2: Click **Add > Network Object**.

The Add Network Object dialog box appears.

Step 3: In the **Name** box, enter the name. (Example: outside-dmvpn-2-ISPa)

Step 4: In the **Type** list, choose **Host** or **Network**. (Example: Host)

Step 5: In the **IP Address** box, enter the address. (Example: 172.16.140.1)

Step 6: In the **Description** box, enter a useful description, and then click **OK**. (Example: DMVPN hub router 2 on ISP A)

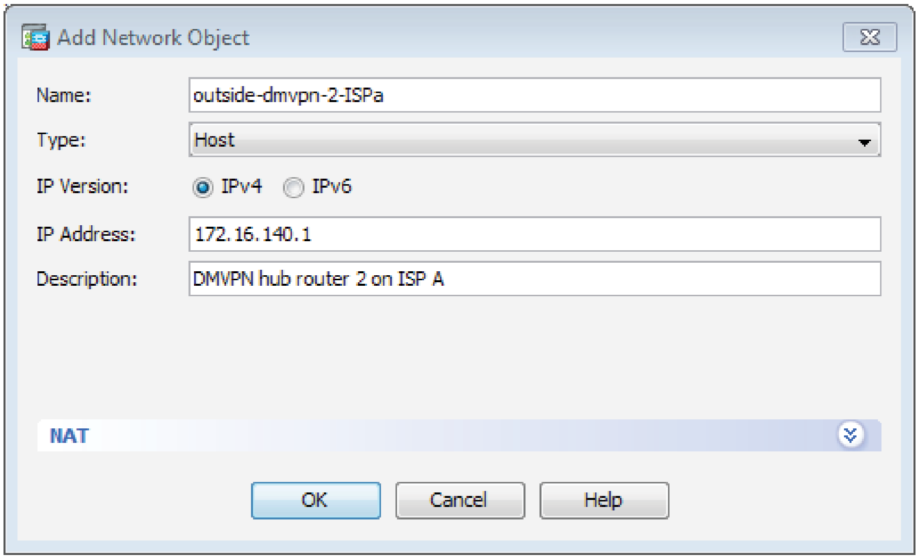

Name: outside-dmvpn-2-ISPa
Type: Host
IP Version: ● IPv4 ○ IPv6
IP Address: 172.16.140.1
Description: DMVPN hub router 2 on ISP A

Step 7: Repeat Step 2 through Step 6 for each object listed in the above table. If an object already exists, then skip to the next object listed in the table.

Step 8: After adding all of the objects listed, click **Apply** on the Network Objects/Groups pane.

Next, you add a network object for the private DMZ address of the DMVPN hub router.

Table 17 Private DMZ firewall network objects

Network object name	Object type	IP address	Description
dmz-dmvpn-2	Host	192.168.146.10	DMVPN hub router 2 on vpn-dmz
dmz-dmvpn-3	Host	192.168.146.20	DMVPN hub router 3 on vpn-dmz
dmz-dmvpn-4	Host	192.168.146.21	DMVPN hub router 4 on vpn-dmz

Step 9: Navigate to **Configuration > Firewall > Objects > Network Objects/Groups**.

Step 10: Click **Add > Network Object**.

The Add Network Object dialog box appears.

Step 11: In the **Name** box, enter the name. (Example: dmz-dmvpn-2)

Step 12: In the **Type** list, choose **Host** or **Network**. (Example: Host)

Step 13: In the **IP Address** box, enter the address. (Example: 192.168.146.10)

Step 14: In the **Description** box, enter a useful description, and then click **OK**. (Example: DMVPN hub router 2 on vpn-dmz)

Step 15: Click the two down arrows. The NAT pane expands.

Step 16: Select **Add Automatic Address Translation Rules**.

Step 17: In the **Translated Address** list, choose the network object created previously. (Example: outside-dm-vpn-2-ISPa)

Step 18: Select **Use one-to-one address translation**, and then click **OK**.

Step 19: Repeat Step 10 through Step 18 for each object listed in the above table. If an object already exists, then skip to the next object listed in the table.

Step 20: After adding all of the objects listed, on the Network Objects/Groups pane, click **Apply**.

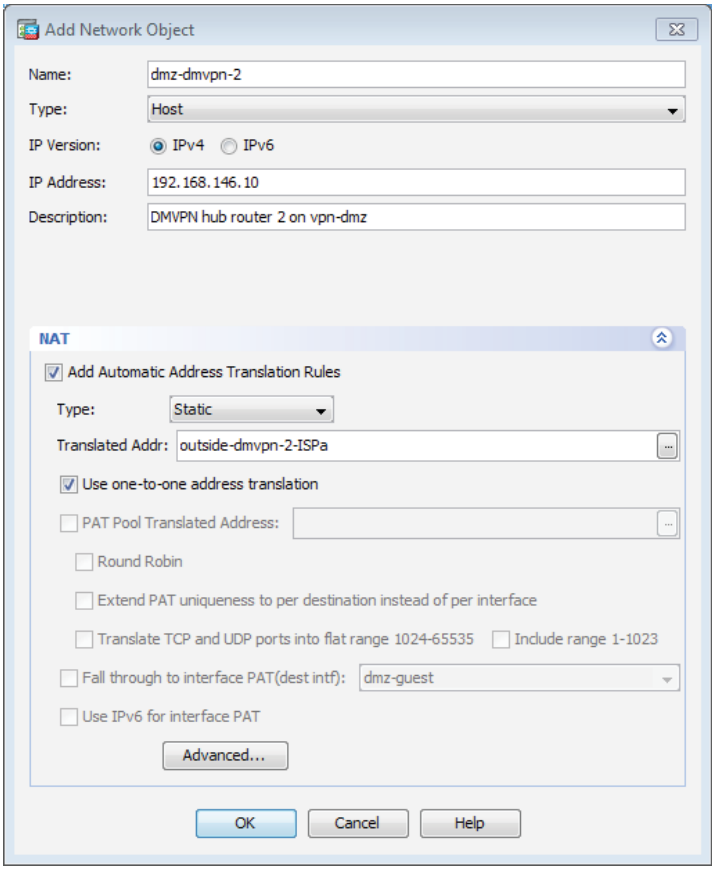

Name: dmz-dmvpn-2
Type: Host
IP Version: ● IPv4 ○ IPv6
IP Address: 192.168.146.10
Description: DMVPN hub router 2 on vpn-dmz

☑ Add Automatic Address Translation Rules
Type: Static
Translated Addr: outside-dmvpn-2-ISPa
☑ Use one-to-one address translation
☐ PAT Pool Translated Address:
☐ Round Robin
☐ Extend PAT uniqueness to per destination instead of per interface
☐ Translate TCP and UDP ports into flat range 1024-65535 ☐ Include range 1-1023
☐ Fall through to interface PAT(dest intf): dmz-guest
☐ Use IPv6 for interface PAT

Procedure 4 Configure security policy

The VPN DMZ provides an additional layer of protection to lower the likelihood of certain types of misconfiguration of the DMVPN routers exposing the business network to the Internet. A filter allows only DMVPN related traffic to reach the DMVPN hub routers from the DMVPN spoke routers on the Internet.

Step 1: Navigate to **Configuration > Firewall > Access Rules**.

Table 18 Firewall policy rules for DMZ-VPN DMVPN hub routers

Interface	Action	Source	Destination	Service	Description	Logging enable/ level
Any	Permit	any4	dmz-vpn-network	udp/4500	(required) Allow (non500-ISAKMP) traffic to the DMVPN hub routers	Selected/ Default
Any	Permit	any4	dmz-vpn-network	udp/isakmp	(required) Allow ISAKMP (UDP500) traffic to the DMVPN hub routers	Selected/ Default
Any	Permit	any4	dmz-vpn-network	Esp	(required) Allow ESP IP protocol 50 IPsec traffic to the DMVPN hub routers	Selected/ Default
Any	Permit	any4	dmz-vpn-network	icmp/echo	(optional) Allow remote ping diagnostic traffic [ICMP Type 0, Code 0]	Selected/ Default
Any	Permit	any4	dmz-vpn-network	icmp/echo reply	(optional) Allow remote pings reply diagnostic traffic [ICMP Type 8, Code 0]	Selected/ Default
Any	Permit	any4	dmz-vpn-network	icmp/time-exceeded	(optional) ICMP Type 11, Code 0	Selected/ Default
Any	Permit	any4	dmz-vpn-network	icmp/port-unreachable	(optional) ICMP Type 3, Code 3	Selected/ Default
Any	Permit	any4	dmz-vpn-network	>udp/1023	(optional) UDP high ports	Selected/ Default

Step 2: Click the rule that denies traffic from the DMZ toward other networks.

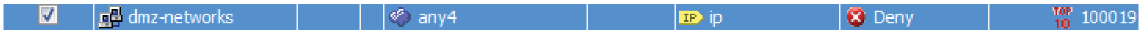

☑ dmz-networks any4 ip Deny 100019

Caution

Be sure to perform this step for *every* rule listed in the previous table. Inserting the rules above the DMZ-to-any rule keeps the added rules in the same order as listed, which is essential for the proper execution of the security policy.

Step 3: Click **Add > Insert**.

The Add Access Rule dialog box appears.

Step 4: In the **Interface** list, choose the interface. (Example: Any)

Step 5: For the **Action** option, select the action. (Example: Permit)

Step 6: In the **Source** box, choose the source. (Example: any4)

Step 7: In the **Destination** box, choose the destination. (Example: dmz-vpn-network)

Step 8: In the **Service** box, enter the service. (Example: udp/4500)

Step 9: In the **Description** box, enter a useful description. (Example: Allow (non500-ISAKMP) traffic to the DM-VPN hub routers)

Step 10: Select or clear **Enable Logging**. (Example: Selected)

Step 11: In the **Logging Level** list, choose the logging level value, and then click **OK**. (Example: Default)

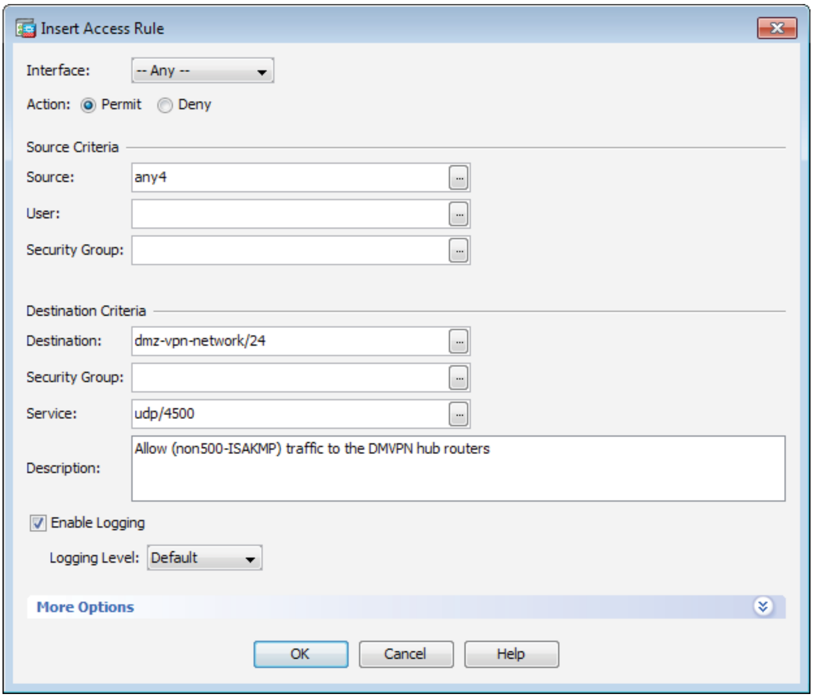

Interface: -- Any --

Action: ● Permit ○ Deny

Source: any4

User:

Security Group:

Destination: dmz-vpn-network/24

Security Group:

Service: udp/4500

Description: Allow (non500-ISAKMP) traffic to the DMVPN hub routers

☑ Enable Logging

Logging Level: Default

Step 12: Repeat Step 2 through Step 11 for all rules listed in the above table.

Step 13: After adding all of the rules in the order listed, click **Apply** on the Access Rules pane.

Figure 21 *Firewall rules summary*

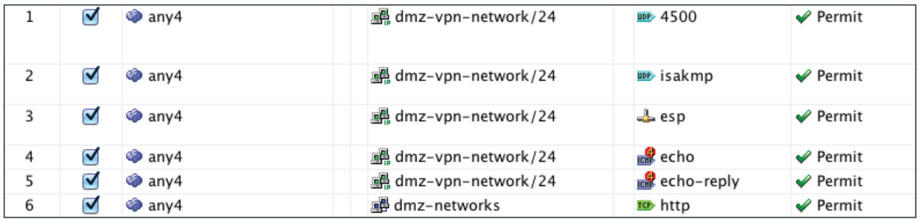

1	☑	any4	dmz-vpn-network/24	4500	✓ Permit
2	☑	any4	dmz-vpn-network/24	isakmp	✓ Permit
3	☑	any4	dmz-vpn-network/24	esp	✓ Permit
4	☑	any4	dmz-vpn-network/24	echo	✓ Permit
5	☑	any4	dmz-vpn-network/24	echo-reply	✓ Permit
6	☑	any4	dmz-networks	http	✓ Permit

PROCESS

Configuring Remote-Site DMVPN Router

1. Configure the WAN remote site router
2. Configure IP multicast routing
3. Configure the WAN-facing VRF
4. Connect to the MPLS WAN or Internet
5. Configure IKEv2 and IPsec
6. Configure the mGRE Tunnel
7. Configure EIGRP
8. Configure IP multicast routing
9. Connect router to access layer switch
10. Configure access layer routing

These procedures describe configuring a single-router, dual-link design. You also use them when configuring the first router of a dual-router, dual-link design.

Refer to the following flowchart to help you navigate through the required procedures for your environment.

Figure 22 Remote-site DMVPN router configuration flowchart

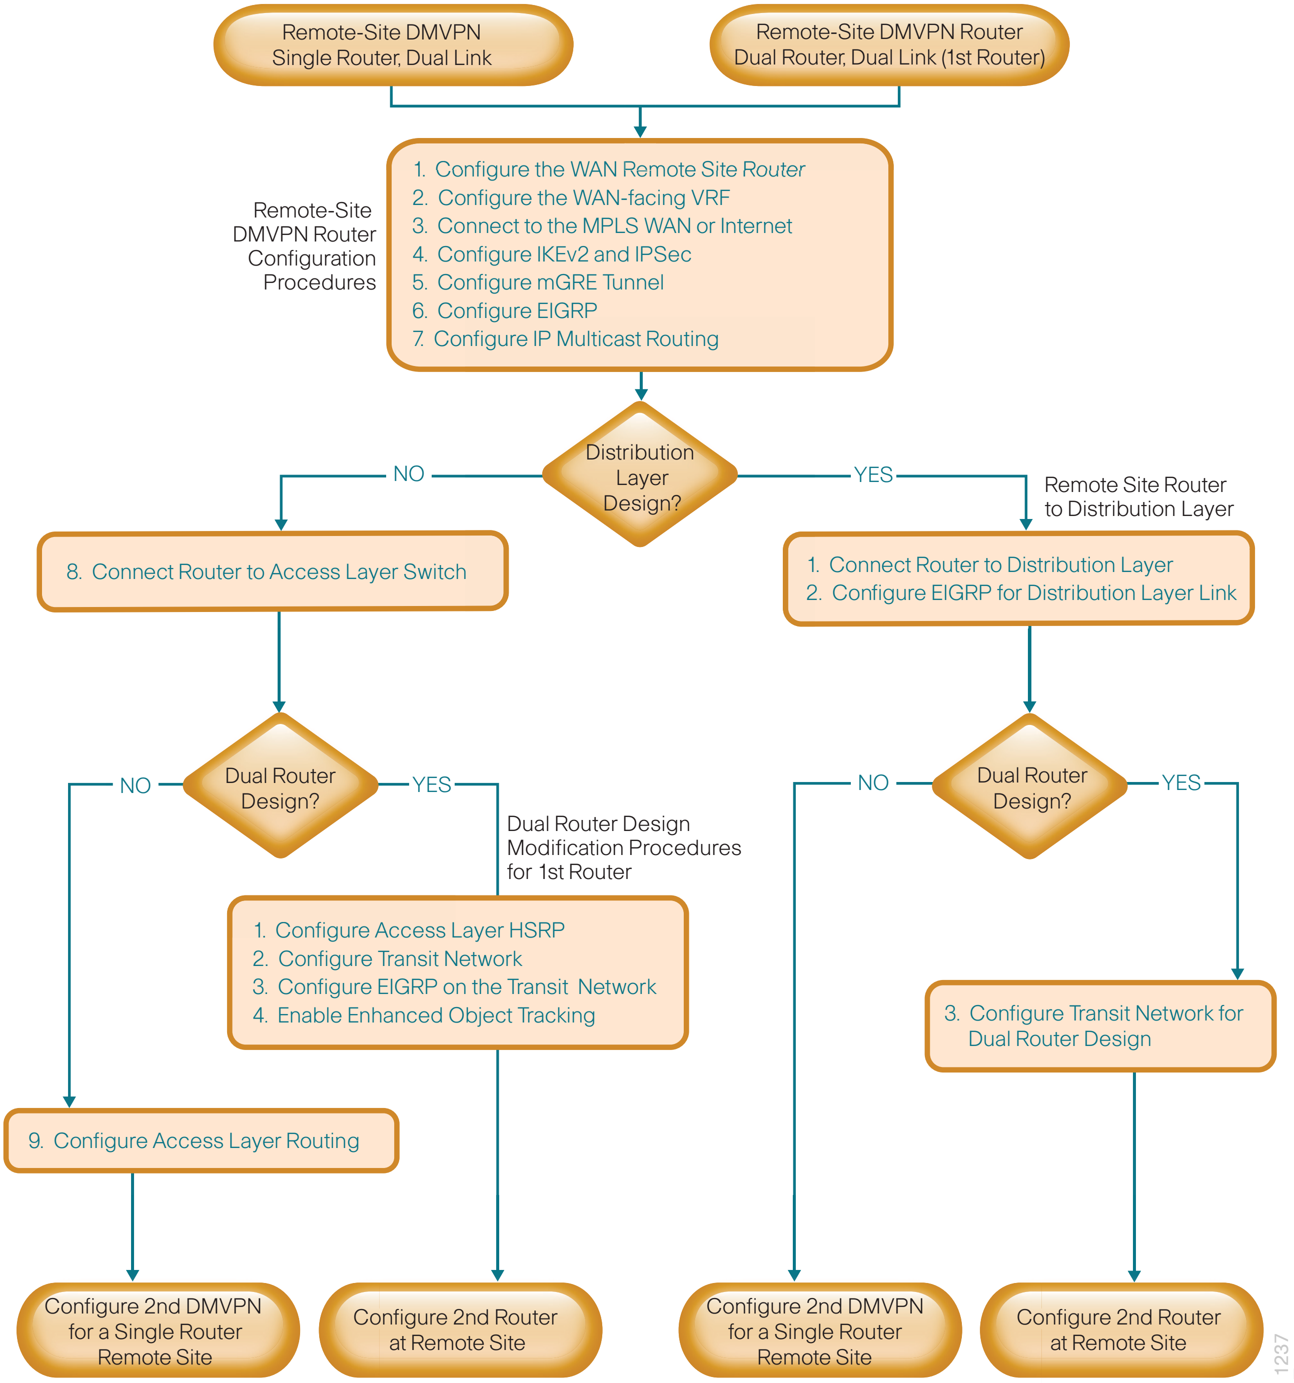
Remote-Site DMVPN
Single Router, Dual Link
Remote-Site DMVPN Router
Dual Router, Dual Link (1st Router)
Remote-Site
DMVPN Router
Configuration
Procedures
1. Configure the WAN Remote Site Router
2. Configure the WAN-facing VRF
3. Connect to the MPLS WAN or Internet
4. Configure IKEv2 and IPSec
5. Configure mGRE Tunnel
6. Configure EIGRP
7. Configure IP Multicast Routing
Distribution Layer Design?
NO
YES
Remote Site Router
to Distribution Layer
8. Connect Router to Access Layer Switch
1. Connect Router to Distribution Layer
2. Configure EIGRP for Distribution Layer Link
Dual Router Design?
NO
YES
Dual Router Design
Modification Procedures
for 1st Router
1. Configure Access Layer HSRP
2. Configure Transit Network
3. Configure EIGRP on the Transit Network
4. Enable Enhanced Object Tracking
Dual Router Design?
NO
YES
3. Configure Transit Network for
Dual Router Design
9. Configure Access Layer Routing
Configure 2nd DMVPN
for a Single Router
Remote Site
Configure 2nd Router
at Remote Site
Configure 2nd DMVPN
for a Single Router
Remote Site
Configure 2nd Router
at Remote Site
1237

Procedure 1 Configure the WAN remote site router

Within this design, there are features and services that are common across all WAN remote site routers. These are system settings that simplify and secure the management of the solution.

To complete the base configuration for this router, follow the steps in "Configure the platform base features" in Appendix C.

Procedure 2 Configure IP multicast routing

Optional

This optional procedure includes additional steps for configuring IP Multicast on a router. Skip this procedure if you do not want to use IP Multicast in your environment.

In this design, which is based on sparse mode multicast operation, Auto RP is used to provide a simple yet scalable way to provide a highly resilient RP environment.

Step 1: Enable IP Multicast routing on the platform in the global configuration mode.

```
ip multicast-routing
```

Step 2: Every Layer 3 switch and router must be configured to discover the IP Multicast RP with autorp. Use the **ip pim autorp listener** command to allow for discovery across sparse mode links. This configuration provides for future scaling and control of the IP Multicast environment and can change based on network needs and design.

```
ip pim autorp listener
```

Step 3: All Layer 3 interfaces in the network must be enabled for sparse mode multicast operation.

```
ip pim sparse-mode
```

Procedure 3 Configure the WAN-facing VRF

You create a WAN-facing VRF in order to support FVRF for DMVPN. The VRF name is arbitrary, but it is useful to select a name that describes the VRF. The VRF must be enabled for IPv4.

Table 19 VRF assignments

IWAN design model	Primary WAN VRF	Secondary WAN VRF
Hybrid	IWAN-TRANSPORT-1	IWAN-TRANSPORT-2
Dual Internet	IWAN-TRANSPORT-3	IWAN-TRANSPORT-4

This design uses VRF Lite, so the selection is only locally significant to the device. It is a best practice to use the same VRF/RD combination across multiple devices when using VRFs in a similar manner. However, this convention is not strictly required.

Step 1: Configure the primary WAN VRF.

Example: Primary WAN in the IWAN hybrid design model

```
vrf definition IWAN-TRANSPORT-1
 address-family ipv4
```

Procedure 4 Connect to the MPLS WAN or Internet

The remote sites that are using DMVPN can use either static or dynamically assigned IP addresses. Cisco tested the design with static addresses for MPLS connections and DHCP assigned external addresses for Internet connections, which also provides a dynamically configured default route.

If you are using MPLS in this design, the DMVPN spoke router is connected to the service provider's MPLS PE router. The IP addressing used between IWAN CE and MPLS PE routers must be negotiated with your MPLS carrier.

The DMVPN spoke router connects directly to the Internet without a separate firewall. This connection is secured in two ways. Because the Internet interface is in a separate VRF, no traffic can access the global VRF except traffic sourced through the DMVPN tunnel. This design provides implicit security. Additionally, an IP access list permits only the traffic required for an encrypted tunnel, as well as DHCP and various ICMP protocols for troubleshooting.

Option 1: MPLS WAN Physical WAN Interface

The DMVPN design uses FVRF, so you must place this interface into the VRF configured in the previous procedure.

Step 1: Enable the interface, select VRF, and assign the IP address.

Example: Primary WAN in IWAN hybrid design model

```
interface GigabitEthernet0/0
 vrf forwarding IWAN-TRANSPORT-1
 ip address 192.168.6.5 255.255.255.252
 no shutdown
```

Do not enable PIM on this interface because no multicast traffic should be requested from this interface.

Step 2: Configure the VRF-specific default routing.

The VRF created for FVRF must have its own default route to the Internet. This default route points to the MPLS PE router's IP address and is used by DMVPN for tunnel establishment.

```
ip route vrf IWAN-TRANSPORT-1 0.0.0.0 0.0.0.0 192.168.6.6
```

Option 2: Internet WAN Physical WAN Interface

The DMVPN design uses FVRF, so you must place this interface into the VRF configured in the previous procedure.

Step 1: Enable the interface, select VRF, and enable DHCP.

Example: Primary WAN in the IWAN dual Internet design model

```
interface GigabitEthernet0/0
 vrf forwarding IWAN-TRANSPORT-3
 ip address dhcp
 no cdp enable
 no shutdown
```

Do not enable PIM on this interface, because no multicast traffic should be requested from this interface.

Step 2: Configure and apply the access list.

The IP access list must permit the protocols specified in the following table. The access list is applied inbound on the WAN interface, so filtering is done on traffic destined to the router.

Table 20 Required DMVPN protocols

Name	Protocol	Usage
non500-isakmp	UDP 4500	IPsec using NAT-T
isakmp	UDP 500	ISAKMP
esp	IP 50	IPsec
bootpc	UDP 68	DHCP

Example: Primary WAN in the IWAN dual Internet design model

```
interface GigabitEthernet0/0
 ip access-group ACL-INET-PUBLIC in


ip access-list extended ACL-INET-PUBLIC
 permit udp any any eq non500-isakmp
 permit udp any any eq isakmp
 permit esp any any
 permit udp any any eq bootpc
```

The additional protocols listed in the following table may assist in troubleshooting, but are not explicitly required to allow DMVPN to function properly.

Table 21 Optional protocols: DMVPN spoke router

Name	Protocol	Usage
icmp echo	ICMP Type 0, Code 0	Allow remote pings
icmp echo-reply	ICMP Type 8, Code 0	Allow ping replies (from your requests)
icmp ttl-exceeded	ICMP Type 11, Code 0	Allow traceroute replies (from your requests)
icmp port-unreachable	ICMP Type 3, Code 3	Allow traceroute replies (from your requests)
UDP high ports	UDP > 1023, TTL=1	Allow remote traceroute

The additional optional entries for an access list to support ping are as follows:

```
permit icmp any any echo
permit icmp any any echo-reply
```

The additional optional entries for an access list to support traceroute are as follows:

```
permit icmp any any ttl-exceeded      ! for traceroute (sourced)
permit icmp any any port-unreachable  ! for traceroute (sourced)
permit udp any any gt 1023 ttl eq 1   ! for traceroute (destination)
```

Procedure 5 Configure IKEv2 and IPsec

The parameters in the tables below are used in this procedure. Choose the table that represents the design model that you are configuring. This procedure applies to the Primary WAN.

Table 22 Crypto parameters: IWAN hybrid design model

Parameter	Primary WAN	Secondary WAN
vrf	IWAN-TRANSPORT-1	IWAN-TRANSPORT-2
crypto ikev2 keyring	DMVPN-KEYRING-1	DMVPN-KEYRING-2
crypto ikev2 profile	FVRF-IKEv2-IWAN-TRANSPORT-1	FVRF-IKEv2-IWAN-TRANSPORT-2
crypto ipsec profile	DMVPN-PROFILE-TRANSPORT-1	DMVPN-PROFILE-TRANSPORT-2

Table 23 Crypto parameters: IWAN dual Internet design model

Parameter	Primary WAN	Secondary WAN
vrf	IWAN-TRANSPORT-3	IWAN-TRANSPORT-4
crypto ikev2 keyring	DMVPN-KEYRING-3	DMVPN-KEYRING-4
crypto ikev2 profile	FVRF-IKEv2-IWAN-TRANSPORT-3	FVRF-IKEv2-IWAN-TRANSPORT-4
crypto ipsec profile	DMVPN-PROFILE-TRANSPORT-3	DMVPN-PROFILE-TRANSPORT-4

IPsec uses a key exchange between the routers in order to encrypt/decrypt the traffic. These keys can be exchanged using pre-shared keys or PKI certificates with a certificate authority. It is also possible to use a combination of the two, which is useful during a migration from one method to the other. Choose one of the two options below as your method of key exchange.

Option 1: Configure with Pre-Shared Keys

Step 1: Configure the crypto keyring for pre-shared keys.

The crypto keyring defines a pre-shared key (or password) valid for IP sources that are reachable within a particular VRF. This key is a wildcard pre-shared key if it applies to any IP source. A wildcard key is configured using the 0.0.0.0 0.0.0.0 network/mask combination.

```
crypto ikev2 keyring [keyring name]
 peer ANY
  address 0.0.0.0 0.0.0.0
  pre-shared-key [password]
```

Example: Primary WAN in the IWAN hybrid design model

```
crypto ikev2 keyring DMVPN-KEYRING-1
 peer ANY
  address 0.0.0.0 0.0.0.0
  pre-shared-key c1sco123
```

Step 2: Configure the IKEv2 proposal.

The user-defined IKEv2 proposal includes only the following:

- Encryption with AES cipher with a 256-bit key
- Integrity with SHA and a 512-bit digest
- Pseudo-random function with SHA and a 512-bit digest
- Diffie-Hellman group: 14

```
crypto ikev2 proposal [proposal name]
 encryption [encryption algorithm]
 integrity [integrity hash algorithm]
 group [Diffie-Hellman group]
```

Example

```
crypto ikev2 proposal AES/CBC/256
 encryption aes-cbc-256
 integrity sha512
 group 14
```

The default IKEv2 proposal is also used.

A **show crypto ikev2 proposal** displays the details of the two proposals.

```
RS11-2921# show crypto ikev2 proposal
IKEv2 proposal: AES/CBC/256
     Encryption : AES-CBC-256
     Integrity  : SHA512
     PRF        : SHA512
     DH Group   : DH_GROUP_2048_MODP/Group 14
IKEv2 proposal: default
    Encryption: AES-CBC-256 AES-CBC-192 AES-CBC-128
    Integrity : SHA512 SHA384 SHA256 SHA96 MD596
    PRF       : SHA512 SHA384 SHA256 SHA1 MD5
    DH Group  : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2
```

Step 3: Configure the IKEv2 profile.

The IKEv2 profile creates an association between an identity address, a VRF, and a crypto keyring. A wildcard address within a VRF is referenced with 0.0.0.0. The profile also defines what method of key sharing will be used on this router with **authentication local** and what methods will be accepted from remote locations with **authentication remote**. The **pre-share** keyword is used with the keyring defined above.

DPD is essential in order to facilitate fast reconvergence and for spoke registration to function properly in case a DMVPN hub is restarted. The IWAN design recommends you set the remote site DPD timer to 40 seconds with a 5 second retry. Moving the DPD timer into the **crypto ikev2 profile** ensures the command will be used immediately, rather than waiting for the first 24 hour refresh cycle if the command is entered in the global configuration.

```
crypto ikev2 profile [profile name]
 match fvrf [vrf name]
 match identity remote address [IP address]
 authentication remote pre-share
 authentication local pre-share
 keyring local [keyring name]
 dpd [interval in seconds] [retry interval] on-demand
```

Example: Primary WAN in the IWAN hybrid design model

```
crypto ikev2 profile FVRF-IKEv2-IWAN-TRANSPORT-1
 match fvrf IWAN-TRANSPORT-1
 match identity remote address 0.0.0.0
 authentication remote pre-share
 authentication local pre-share
 keyring local DMVPN-KEYRING-1
 dpd 40 5 on-demand
```

Step 4: Define the IPsec transform set.

A transform set is an acceptable combination of security protocols, algorithms, and other settings to apply to IPsec-protected traffic. Peers agree to use a particular transform set when protecting a particular data flow.

The IPsec transform set for DMVPN uses the following:

- ESP with the 256-bit AES encryption algorithm
- ESP with the SHA (HMAC variant) authentication algorithm

Because the DMVPN hub router is behind a NAT device, the IPsec transform must be configured for transport mode.

```
crypto ipsec transform-set [transform set] esp-aes 256 esp-sha-hmac
 mode transport
```

Example

```
crypto ipsec transform-set AES256/SHA/TRANSPORT esp-aes 256 esp-sha-hmac
 mode transport
```

Step 5: Configure the IPsec profile.

The IPsec profile creates an association between an IKEv2 profile and an IPsec transform-set.

```
crypto ipsec profile [profile name]
 set transform-set [transform set]
 set ikev2-profile [ikev2 profile name]
```

Example: Primary WAN in the IWAN hybrid design model

```
crypto ipsec profile DMVPN-PROFILE-TRANSPORT-1
 set transform-set AES256/SHA/TRANSPORT
 set ikev2-profile FVRF-IKEv2-IWAN-TRANSPORT-1
```

Step 6: Increase the IPsec anti-replay window size.

```
crypto ipsec security-association replay window-size 1024
```

Tech Tip

QoS queuing delays can cause anti-replay packet drops, so it is important to extend the window size to prevent the drops from occurring.

Increasing the anti-replay window size has no impact on throughput and security. The impact on memory is insignificant because only an extra 128 bytes per incoming IPsec SA are needed.

It is recommended that you use the 1024 window size in order to minimize future anti-replay problems.

If you do not increase the window size, the router may drop packets and you may see the following error message on the router CLI and in the log:

```
%CRYPTO-4-PKT_REPLAY_ERR: decrypt: replay check failed
```

Step 7: Proceed to Procedure 6, “Configure the mGRE Tunnel.”

Option 2: Configure with a certificate authority

If you want to use a certificate authority, you will have to configure a pre-shared key on one of the hub border routers in order to allow each remote site to establish a DMVPN tunnel to the WAN aggregation site. After the first DMVPN tunnel at a remote site is established, the router will be able to authenticate to the CA and obtain a certificate. After obtaining the certificate, you can configure the remote site to use PKI.

The **crypto pki trustpoint** is the method of specifying the parameters associated with a CA. The router must authenticate to the CA first and then enroll with the CA in order to obtain its own identity certificate.

Step 1: The fingerprint command limits the responding CA. You can find this fingerprint by using **show crypto pki server** on the IOS CA.

```
IWAN-IOS-CA# show crypto pki server
Certificate Server IWAN-IOS-CA:
    Status: enabled
    State: enabled
    Server’s configuration is locked  (enter "shut" to unlock it)
    Issuer name: CN=IWAN-IOS-CA.cisco.local L=SanJose St=CA C=US
    CA cert fingerprint: 75BEF625 9A9876CF 6F341FE5 86D4A5D8
    Granting mode is: auto
    Last certificate issued serial number (hex): 11
    CA certificate expiration timer: 10:28:57 PST Nov 11 2017
    CRL NextUpdate timer: 09:47:47 PST Dec 4 2014
    Current primary storage dir: nvram:
    Current storage dir for .crl files: nvram:
    Database Level: Complete - all issued certs written as <serialnum>.cer
```

Step 2: Configure the PKI trust point.

```
crypto pki trustpoint [name]
 enrollment url [URL of IOS CA]
 serial-number none
 fqdn [fully qualified domain name of this router]
 ip-address [loopback IP address of this router]
 fingerprint [fingerprint from IOS CA]
 revocation-check none
 rsakeypair [name] 2048 2048
```

Example: RS14-2921-1

This example is from the primary WAN remote site router in the dual Internet design model. After the DMVPN tunnel is established with pre-shared keys, it can reach the IOS CA through the internal network at 10.6.24.11 using the default VRF.

```
crypto pki trustpoint IWAN-CA
 enrollment url http://10.6.24.11:80
 serial-number none
 fqdn RS14-2921-1.cisco.local
 ip-address 10.255.243.14
 fingerprint 75BEF6259A9876CF6F341FE586D4A5D8
 revocation-check none
 rsakeypair IWAN-CA-KEYS 2048 2048
```

Step 3: Authenticate to the CA and obtain the CA's certificate

Exit the trustpoint configuration mode on the hub router and issue the following command to authenticate to the CA and get its certificate.

```
RS14-2921-1(config)# crypto pki authenticate IWAN-CA
Certificate has the following attributes:
       Fingerprint MD5: 75BEF625 9A9876CF 6F341FE5 86D4A5D8
      Fingerprint SHA1: 9C14D6F4 D1F08023 17A85669 52922632 C6B02928
Trustpoint Fingerprint: 75BEF625 9A9876CF 6F341FE5 86D4A5D8
Certificate validated - fingerprints matched.
```

Step 4: When the trustpoint CA certificate is accepted, enroll with the CA, enter a password for key retrieval and obtain a certificate for this hub router.

```
RS14-2921-1(config)# crypto pki enroll IWAN-CA
% Start certificate enrollment ..
% Create a challenge password. You will need to verbally provide this
   password to the CA Administrator in order to revoke your certificate.
   For security reasons your password will not be saved in the configuration.
   Please make a note of it.

Password: c1sco123
Re-enter password: c1sco123

% The subject name in the certificate will include: RS14-2921-1.cisco.local
% The IP address in the certificate is 10.255.243.14
```

```
Request certificate from CA? [yes/no]: yes
% Certificate request sent to Certificate Authority
% The 'show crypto pki certificate verbose IWAN-CA' commandwill show the finger-
print.
```

Step 5: Configure the IKEv2 proposal.

The user-defined IKEv2 proposal includes only the following:

- Encryption with AES cipher with a 256-bit key
- Integrity with SHA and a 512-bit digest
- Pseudo-random function with SHA and a 512-bit digest
- Diffie-Hellman group: 14

```
crypto ikev2 proposal [proposal name]
 encryption [encryption algorithm]
 integrity [integrity hash algorithm]
 group [Diffie-Hellman group]
```

Example

```
crypto ikev2 proposal AES/CBC/256
 encryption aes-cbc-256
 integrity sha512
 group 14
```

The default IKEv2 proposal is also used.

A **show crypto ikev2 proposal** displays the details of the two proposals.

```
RS14-2921-1# show crypto ikev2 proposal
IKEv2 proposal: AES/CBC/256
     Encryption : AES-CBC-256
     Integrity  : SHA512
     PRF        : SHA512
     DH Group   : DH_GROUP_2048_MODP/Group 14
IKEv2 proposal: default
    Encryption: AES-CBC-256 AES-CBC-192 AES-CBC-128
    Integrity : SHA512 SHA384 SHA256 SHA96 MD596
    PRF       : SHA512 SHA384 SHA256 SHA1 MD5
    DH Group  : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2
```

Step 6: Configure the IKEv2 profile.

The IKEv2 profile creates an association between an identity address, a VRF, and a crypto keyring. A wildcard address within a VRF is referenced with 0.0.0.0. The **identity local address** must match the crypto pki trustpoint's **ip-address** value from the step above.

Tech Tip

Use the **identity local address** in the ikev2 profile in order to avoid repeated CRYPTO-6-IKMP_NO_ID_CERT_ADDR_MATCH warning messages on the router.

The profile also defines what method of key sharing will be used on this router with **authentication local** and what methods will be accepted from remote locations with **authentication remote**. The **rsa-sig** keyword is used when certificates contain the encryption key.

DPD is essential in order to facilitate fast reconvergence and for spoke registration to function properly in case a DMVPN hub is restarted. The IWAN design recommends you set the remote site DPD timer to 40 seconds with a 5 second retry.

```
crypto ikev2 profile [profile name]
 match fvrf [vrf name]
 match identity remote address [IP address]
 identity local address [Loopback IP address of this router]
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint [trustpoint name]
 dpd [interval in seconds] [retry interval] on-demand
```

Example: RS14-2921-1

```
crypto ikev2 profile FVRF-IKEv2-IWAN-TRANSPORT-3
 match fvrf IWAN-TRANSPORT-3
 match identity remote address 0.0.0.0
 identity local address 10.255.243.14
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint IWAN-CA
 dpd 40 5 on-demand
```

Step 7: Define the IPsec transform set.

A transform set is an acceptable combination of security protocols, algorithms, and other settings to apply to IPsec-protected traffic. Peers agree to use a particular transform set when protecting a particular data flow.

The IPsec transform set for DMVPN uses the following:

- ESP with the 256-bit AES encryption algorithm
- ESP with the SHA (HMAC variant) authentication algorithm

Because the DMVPN hub router is behind a NAT device, the IPsec transform must be configured for transport mode.

```
crypto ipsec transform-set [transform set] esp-aes 256 esp-sha-hmac
 mode transport
```

Example

```
crypto ipsec transform-set AES256/SHA/TRANSPORT esp-aes 256 esp-sha-hmac
 mode transport
```

Step 8: Configure the IPsec profile.

The IPsec profile creates an association between an IKEv2 profile and an IPsec transform-set.

```
crypto ipsec profile [profile name]
 set transform-set [transform set]
 set ikev2-profile [ikev2 profile name]
```

Example: Primary WAN in the dual Internet design model

```
crypto ipsec profile DMVPN-PROFILE-TRANSPORT-3
 set transform-set AES256/SHA/TRANSPORT
 set ikev2-profile FVRF-IKEv2-IWAN-TRANSPORT-3
```

Step 9: Increase the IPsec anti-replay window size.

```
crypto ipsec security-association replay window-size [value]
```

Example

```
crypto ipsec security-association replay window-size 1024
```

Tech Tip

QoS queuing delays can cause anti-replay packet drops, so it is important to extend the window size in order to prevent the drops from occurring.

Increasing the anti-replay window size has no impact on throughput and security. The impact on memory is insignificant because only an extra 128 bytes per incoming IPsec SA is needed.

It is recommended that you use the maximum window size in order to minimize future anti-replay problems. On the Cisco ASR 1K, ISR 4K and ISR G2 router platforms, the maximum replay window size is 1024.

If you do not increase the window size, the router may drop packets and you may see the following error message on the router CLI and in the log:

```
%CRYPTO-4-PKT_REPLAY_ERR: decrypt: replay check failed
```

Procedure 6 Configure the mGRE Tunnel

The parameters in the table below are used in this procedure. Choose the rows that represent the design model that you are configuring. This procedure applies to the Primary WAN remote site router in the IWAN hybrid design model.

Table 24 DMVPN tunnel parameters

Design model	Tunnel VRF	Tunnel number	Tunnel network	NHRP network ID/ tunnel key
Hybrid—Primary WAN	IWAN-TRANSPORT-1	10	10.6.34.0/23	101
Hybrid—Secondary WAN	IWAN-TRANSPORT-2	11	10.6.36.0/23	102
Dual Internet—Primary WAN	IWAN-TRANSPORT-3	20	10.6.38.0/23	201
Dual Internet—Secondary WAN	IWAN-TRANSPORT-4	21	10.6.40.0/23	202

Step 1: Configure basic interface settings.

The tunnel number is arbitrary, but it is best to begin tunnel numbering at 10 or above, because other features deployed in this design may also require tunnels and they may select lower numbers by default.

You must set the bandwidth to match the bandwidth of the respective transport which corresponds to the actual interface speed. Or, if you are using a subrate service, use the policed rate from the carrier. QoS and PfR require the correct bandwidth setting in order to operate properly.

Configure the **ip mtu** to 1400 and the **ip tcp adjust-mss** to 1360. There is a 40 byte difference, which corresponds to the combined IP and TCP header length.

The tunnel interface throughput delay setting should be set to influence the routing protocol path preference. Set the primary WAN path to 10000 usec and the secondary WAN path to 200000 usec to prefer one over the other. The delay command is entered in 10 usec units.

```
interface Tunnel10
 bandwidth 300000
 ip address 10.6.34.11 255.255.254.0
 no ip redirects
 ip mtu 1400
 ip tcp adjust-mss 1360
 delay 1000
```

Step 2: Configure the tunnel.

DMVPN uses mGRE tunnels. This type of tunnel requires a source interface only. The source interface should be the same interface used in to connect to the MPLS or Internet. The **tunnel vrf** command should be set to the VRF defined previously for FVRF.

Enabling encryption on this interface requires the application of the IPsec profile configured in the previous procedure.

```
interface Tunnel10
 tunnel source GigabitEthernet0/0
 tunnel mode gre multipoint
 tunnel key 101
 tunnel vrf IWAN-TRANSPORT-1
 tunnel protection ipsec profile DMVPN-PROFILE-TRANSPORT-1
```

Step 3: Configure NHRP.

The DMVPN hub router is the NHRP server for all of the spokes. Remote routers use NHRP in order to determine the tunnel destinations for peers attached to the mGRE tunnel.

The spoke router requires an additional configuration statement in order to define the NHRP server. This statement includes the NBMA definition for the DMVPN hub router tunnel endpoint. EIGRP relies on a multicast transport. Spoke routers require the NHRP multicast keyword in this statement.

When hub BRs are added for horizontal scaling or a second data center is added as a transit site, spoke routers require additional NHS statements for each BR in their environment. The configuration details are covered in subsequent sections of this guide.

The value used for the next hop server (NHS) is the mGRE tunnel address for the DMVPN hub router. The NBMA entry must be set to either the MPLS DMVPN hub router's actual public address or the outside NAT value of the DMVPN hub, as configured on the Cisco ASA 5500. This design uses the values shown in the following tables.

Table 25 DMVPN tunnel NHRP parameters: IWAN hybrid design model

	Transport 1	Transport 2
VRF	IWAN-TRANSPORT-1	IWAN-TRANSPORT-2
DMVPN hub public address (actual)	192.168.6.1	192.168.146.10
DMVPN hub public address (externally routable after NAT)	n/a (MPLS)	172.16.140.1
DMVPN hub tunnel IP address (NHS)	10.6.34.1	10.6.36.1
Tunnel number	10	11
NHRP network ID	101	102

Table 26 DMVPN tunnel NHRP parameters: IWAN dual Internet design model

	Transport 3	Transport 4
VRF	IWAN-TRANSPORT-3	IWAN-TRANSPORT-4
DMVPN hub public address (actual)	192.168.146.20	192.168.146.21
DMVPN hub public address (externally routable after NAT)	172.16.140.11	172.17.140.11
DMVPN hub tunnel IP address (NHS)	10.6.38.1	10.6.40.1
Tunnel number	20	21
NHRP network ID	201	202

NHRP requires all devices within a DMVPN cloud to use the same network ID and authentication key. The NHRP cache holdtime should be configured to 600 seconds.

This design supports DMVPN spoke routers that receive their external IP addresses through DHCP. It is possible for these routers to acquire different IP addresses after a reload. When the router attempts to register with the NHRP server, it may appear as a duplicate to an entry already in the cache and be rejected. The **registration no-unique** option allows you to overwrite existing cache entries. This feature is only required on NHRP clients (DMVPN spoke routers). The **if-state nhrp** option ties the tunnel line-protocol state to the reachability of the NHRP NHS, and if the NHS is unreachable, the tunnel line-protocol state changes to down.

This feature is used in conjunction with EOT.

```
interface Tunnel10
 ip nhrp authentication cisco123
 ip nhrp network-id 101
 ip nhrp holdtime 600
 ip nhrp nhs 10.6.34.1 nbma 192.168.6.1 multicast
 ip nhrp registration no-unique
 ip nhrp shortcut
 if-state nhrp
```

By default, NHRP will not install shortcuts for paths not seen in the Routing Information Base (RIB) of the router. In a location with a single router and multiple WAN transports, only the preferred path is in the RIB. If you have a remote site location with more than one WAN transport, you need to disable the **nhrp route-watch** feature on each of the tunnel interfaces in order to allow NHRP to install the non-preferred shortcut path and allow PfR to maintain this information.

```
interface Tunnel10
 no nhrp route-watch
```

Procedure 7 Configure EIGRP

A single EIGRP process runs on the DMVPN spoke router. All interfaces on the router are EIGRP interfaces, but only the DMVPN tunnel interface is non-passive. The network range must include all interface IP addresses either in a single network statement or in multiple network statements. This design uses a best practice of assigning the router ID to a loopback address. All DMVPN spoke routers should run EIGRP stub routing to improve network stability and reduce resource utilization.

Step 1: Configure an EIGRP process for DMVPN using EIGRP named mode on the spoke router.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface default
   passive-interface
  exit-af-interface
  af-interface Tunnel10
   no passive-interface
  exit-af-interface
network 10.6.34.0 0.0.1.255
network 10.7.0.0 0.0.255.255
network 10.255.0.0 0.0.255.255
eigrp router-id [IP address of Loopback0]
eigrp stub connected summary redistributed
exit-address-family
```

Step 2: Configure EIGRP values for the mGRE tunnel interface.

The EIGRP hello interval is increased to 20 seconds and the EIGRP hold time is increased to 60 seconds in order to accommodate up to 2000 remote sites on a single DMVPN cloud. Increasing the EIGRP timers also slows down the routing convergence in order to improve network stability and the IWAN design allows PfR to initiate the fast failover, so changing the timers is recommended for all IWAN deployments.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel10
   hello-interval 20
   hold-time 60
  exit-af-interface
 exit-address-family
```

Step 3: Configure EIGRP neighbor authentication. Neighbor authentication enables the secure establishment of peering adjacencies and exchange route tables over the DMVPN tunnel interface.

```
key chain WAN-KEY
 key 1
  key-string c1sco123

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel10
   authentication mode md5
   authentication key-chain WAN-KEY
  exit-af-interface
 exit-address-family
```

Step 4: Configure EIGRP network summarization.

The remote-site LAN networks must be advertised. The IP assignment for the remote sites was designed so that all of the networks in use can be summarized within a single aggregate route. As configured below, the **summary-address** command suppresses the more specific routes. If any network within the summary is present in the route table, the summary is advertised to the DMVPN hub, which offers a measure of resiliency. If the various LAN networks cannot be summarized, then EIGRP continues to advertise the specific routes.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel10
```

```
   summary-address [summary network] [summary mask]
  exit-af-interface
 exit-address-family
```

Step 5: Tag and filter the routes.

This design uses a single EIGRP autonomous system for the hub sites, WAN, and all of the WAN remote sites. Every remote site is dual-connected for resiliency. However, due to the multiple paths that exist within this topology, you must try to avoid routing loops and to prevent remote sites from becoming transit sites if WAN failures were to occur.

In this design, there are different IP subnets for each DMVPN network, and the EIGRP tags are defined for each segment in order to help with readability and troubleshooting. When a design uses more than one data center, additional tags are required in order to identify the different DMVPN hub router locations.

The following logic is used to control the routing.

- Each DMVPN network will have an EIGRP route tag to prevent routes from being re-advertised over the other DMVPN networks.
- All prefixes that are advertised towards the WAN are uniquely tagged.
- All DMVPN learned WAN prefixes, except those that originate locally from a hub, are advertised towards the LAN and tagged.
- The design always uses DMVPN hub routers in pairs. Each DMVPN hub router blocks DMVPN WAN routes from the LAN that are tagged with the opposite hub's tags.
- Remote sites can learn routes from other remote sites when NHRP uses shortcut routing to bypass the hub locations. The spoke-to-spoke routes are blocked using tag filtering to prevent issues with PfR.

Outbound distribute-lists are used to set tags towards the WAN and LAN. The remote-site routers use the tags set towards the WAN in order to protect against becoming transit sites.

The remote-site routers use an inbound distribute-list in order to limit which routes are accepted for installation into the route table. These routers are configured to only accept routes that do not originate from the WAN sources. To accomplish this task, the remote site router must explicitly tag the DMVPN learned WAN routes. The following tables show specific route tags in use.

Table 27 Route tag information for IWAN hybrid remote site routers

Tunnel: transport	Tunnel key	Match inbound	Block outbound
Tunnel10: DMVPN1	101 (MPLS)	101 (All routes)	101 (WAN routes)
Tunnel11: DMVPN2	102 (INET)	102 (All routes)	102 (WAN routes)

Table 28 Route tag information for IWAN dual Internet remote site routers

Tunnel: transport	Tunnel key	Match inbound	Block outbound
Tunnel20 DMVPN3	201 (INET1)	201 (All routes)	201 (WAN routes)
Tunnel21: DMVPN4	202 (INET2)	202 (All routes)	202 (WAN routes)

The following examples show single and dual router designs in the IWAN hybrid design model.

Example: Single Router Site with DMVPN1 and DMVPN2

```
route-map DMVPN1-BR-IN Permit 10
 description Match tagged routes inbound
 match tag 101

route-map DMVPN2-BR-IN Permit 10
 description Match tagged routes inbound
 match tag 102

route-map BLOCK-LEARNED deny 10
 description Block learned routes outbound
 Match tag 101 102

route-map BLOCK-LEARNED permit 20
 description Advertise all other routes outbound

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  topology base
   distribute-list route-map DMVPN1-BR-IN in Tunnel10
   distribute-list route-map DMVPN2-BR-IN in Tunnel11
   distribute-list route-map BLOCK-LEARNED out Tunnel10
   distribute-list route-map BLOCK-LEARNED out Tunnel11
```

Example: First Router of a Dual Router Site with DMVPN1

```
route-map DMVPN1-BR-IN Permit 10
 description Match tagged routes inbound
 match tag 101

route-map BLOCK-LEARNED deny 10
 description Block learned routes outbound
 Match tag 101 102

route-map BLOCK-LEARNED permit 20
 description Advertise all other routes outbound

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  topology base
   distribute-list route-map DMVPN1-BR-IN in Tunnel10
   distribute-list route-map BLOCK-LEARNED out Tunnel10
```

Procedure 8 Configure IP multicast routing

Optional

This optional procedure includes additional steps for configuring IP Multicast for a DMVPN tunnel on a router with IP Multicast already enabled. Skip this procedure if you do not want to use IP Multicast in your environment.

Step 1: Configure PIM on the DMVPN tunnel interface.

Enable IP PIM sparse mode on the DMVPN tunnel interface.

```
interface Tunnel10
 ip pim sparse-mode
```

Step 2: Enable PIM non-broadcast multiple access mode for the DMVPN tunnel.

Spoke-to-spoke DMVPN networks present a unique challenge because the spokes cannot directly exchange information with one another, even though they are on the same logical network. This inability to directly exchange information can also cause problems when running IP Multicast.

To resolve the NBMA issue, you need to implement a method where each remote PIM neighbor has its join messages tracked separately. A router in PIM NBMA mode treats each remote PIM neighbor as if it were connected to the router through a point-to-point link.

```
interface Tunnel10
 ip pim nbma-mode
```

Step 3: Configure the DR priority for the DMVPN spoke router.

Proper multicast operation across a DMVPN cloud requires that the hub router assumes the role of PIM designated router (DR). Spoke routers should never become the DR. You can prevent that by setting the DR priority to 0 for the spokes.

```
interface Tunnel10
 ip pim dr-priority 0
```

Procedure 9 Connect router to access layer switch

Optional

If you are using a remote-site distribution layer, skip to the "Deploying an IWAN Remote-Site Distribution Layer" section of this guide.

Reader Tip

This guide includes only the steps needed in order to complete the access layer configuration. For complete access layer configuration details, refer to the Campus LAN Layer 2 Access with Simplified Distribution Deployment Guide.

Layer 2 EtherChannels are used to interconnect the CE router to the access layer in the most resilient method possible. If your access layer device is a single fixed configuration switch, a simple Layer 2 trunk between the router and switch is used.

In the access layer design, the remote sites use collapsed routing, with 802.1Q trunk interfaces to the LAN access layer. The VLAN numbering is locally significant only.

Option 1: Layer 2 EtherChannel from router to access layer switch

Step 1: Configure port-channel interface on the router.

```
interface Port-channel1
 description EtherChannel link to RS12-A2960X
 no shutdown
```

Step 2: Configure EtherChannel member interfaces on the router.

Configure the physical interfaces to tie to the logical port-channel using the channel-group command. The number for the port-channel and channel-group must match. Not all router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet0/1
 description RS12-A2960X Gig1/0/47

interface GigabitEthernet0/2
 description RS12-A2960X Gig2/0/47

interface range GigabitEthernet0/1, GigabitEthernet0/2
  no ip address
  channel-group 1
  no shutdown
```

Step 3: Configure EtherChannel member interfaces on the access layer switch.

Connect the router EtherChannel uplinks to separate switches in the access layer switch stack.

The physical interfaces that are members of a Layer 2 EtherChannel are configured prior to configuring the logical port-channel interface. Doing the configuration in this order allows for minimal configuration and reduces errors because most of the commands entered to a port-channel interface are copied to its members interfaces and do not require manual replication.

Configure two physical interfaces to be members of the EtherChannel. Also, apply the egress QoS macro that was defined in the LAN switch platform configuration procedure to ensure traffic is prioritized appropriately.

Not all connected router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet1/0/47
 description Link to RS12-2911-1 Gig0/1

interface GigabitEthernet2/0/47
 description Link to RS12-2911-1 Gig0/2

interface range GigabitEthernet1/0/24, GigabitEthernet2/0/24
 switchport
 macro apply EgressQoS
 channel-group 1 mode on
 logging event link-status
 logging event trunk-status
 logging event bundle-status
```

Step 4: Configure EtherChannel trunk on the access layer switch.

An 802.1Q trunk is used which allows the router to provide the Layer 3 services to all the VLANs defined on the access layer switch. The VLANs allowed on the trunk are pruned to only the VLANs that are active on the access layer switch. When using EtherChannel the interface type will be port-channel and the number must match the channel group configured in Step 3. DHCP Snooping and address resolution protocol (ARP) inspection are set to trust.

```
interface Port-channel1
 description EtherChannel link to RS12-2911-1
 switchport trunk allowed vlan 64,69
 switchport mode trunk
 ip arp inspection trust
 spanning-tree portfast trunk
 ip dhcp snooping trust
 load-interval 30
 no shutdown
```

The Cisco Catalyst 3750 Series Switch requires the **switchport trunk encapsulation dot1q** command

Option 2: Layer 2 trunk from router to access layer switch

Step 1: Enable the physical interface on the router.

```
interface GigabitEthernet0/2
 description RS11-A2960X Gig1/0/48
 no ip address
 no shutdown
```

Step 2: Configure the trunk on the access layer switch.

Use an 802.1Q trunk for the connection, which allows the router to provide the Layer 3 services to all the VLANs defined on the access layer switch. The VLANs allowed on the trunk are pruned to only the VLANs that are active on the access switch. DHCP Snooping and ARP inspection are set to trust.

```
interface  GigabitEthernet1/0/48
 description Link to RS11-2921 Gig0/2
 switchport trunk allowed vlan 64,69
 switchport mode trunk
 ip arp inspection trust
 spanning-tree portfast trunk
 logging event link-status
 logging event trunk-status
 ip dhcp snooping trust
 load-interval 30
 no shutdown
 macro apply EgressQoS
```

The Cisco Catalyst 3750 Series Switch requires the **switchport trunk encapsulation dot1q** command

Procedure 10 Configure access layer routing

Optional

If you are using a dual router design, skip to the “Modifying the First Router for Dual Router Design” section of this guide.

Step 1: Create subinterfaces and assign VLAN tags.

After the physical interface or port-channel has been enabled, then the appropriate data or voice subinterfaces can be mapped to the VLANs on the LAN switch. The subinterface number does not need to equate to the 802.1Q tag, but making them the same simplifies the overall configuration. The subinterface portion of the configuration should be repeated for all data or voice VLANs.

```
interface [type][number].[sub-interface number]
 encapsulation dot1Q [dot1q VLAN tag]
```

Step 2: Configure IP settings for each subinterface.

This design uses an IP addressing convention with the default gateway router assigned an IP address and IP mask combination of **N.N.N.1 255.255.255.0** where N.N.N is the IP network and 1 is the IP host.

When you are using a centralized DHCP server, your routers with LAN interfaces connected to a LAN using DHCP for end-station IP addressing must use an IP helper.

If the remote-site router is the first router of a dual-router design, then HSRP is configured at the access layer. This requires a modified IP configuration on each subinterface.

```
interface [type][number].[sub-interface number]
 ip address [LAN network 1] [LAN network 1 netmask]
 ip helper-address 10.4.48.10
 ip pim sparse-mode
```

Example: Layer 2 EtherChannel

```
interface Port-channel1
 no ip address
 no shutdown

 interface Port-channel1.64
 description Data
 encapsulation dot1Q 64
 ip address 10.7.18.1 255.255.255.0
 ip helper-address 10.4.48.10
 ip pim sparse-mode

interface Port-channel1.69
 description Voice
 encapsulation dot1Q 69
 ip address 10.7.19.1 255.255.255.0
 ip helper-address 10.4.48.10
 ip pim sparse-mode
```

Example: Layer 2 Link

```
interface GigabitEthernet0/2
 no ip address
 no shutdown

interface GigabitEthernet0/2.64
 description Data
 encapsulation dot1Q 64
 ip address 10.7.2.1 255.255.255.0
 ip helper-address 10.4.48.10
 ip pim sparse-mode

interface GigabitEthernet0/2.69
 description Voice
 encapsulation dot1Q 69
 ip address 10.7.3.1 255.255.255.0
 ip helper-address 10.4.48.10
 ip pim sparse-mode
```

PROCESS

Adding Second DMVPN for a Single-Router Remote Site

1. Configure the WAN-facing VRF
2. Connect to the Internet
3. Configure IKEv2 and IPsec
4. Configure the mGRE Tunnel
5. Configure EIGRP
6. Configure IP multicast routing

Use this set of procedures for either of the following topologies: single router with a hybrid (MPLS+INET) remote site or a dual Internet (INET+INET) remote site.

This set of procedures includes the additional steps necessary to add a second DMVPN link to a remote-site router that has already been configured with a DMVPN link in the “Configuring Remote-Site DMVPN Router” process in this guide.

The following flowchart details how to add the second DMVPN to an existing remote-site router.

Figure 23 Adding second DMVPN configuration flowchart

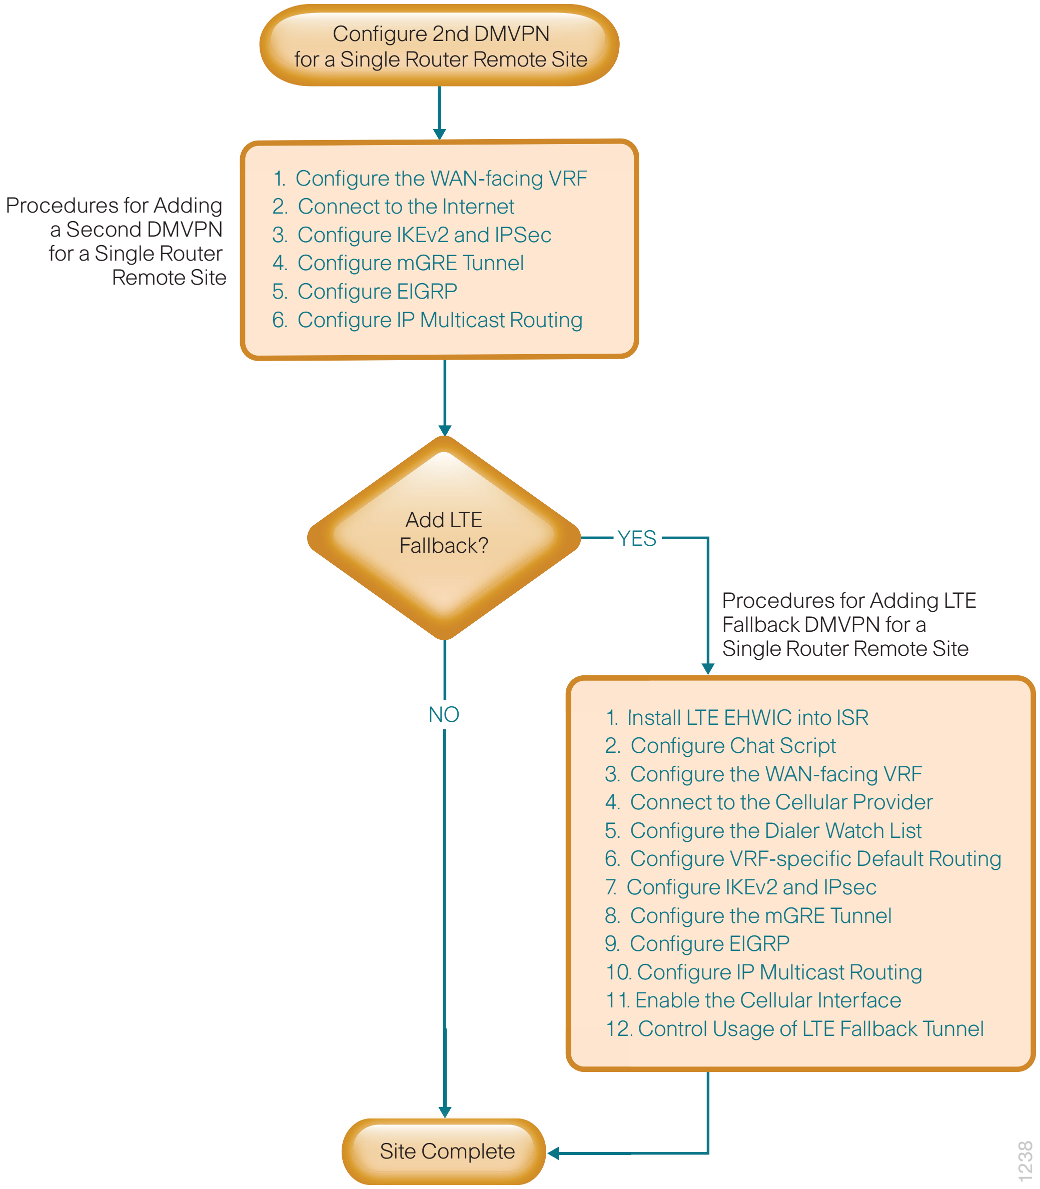
Configure 2nd DMVPN for a Single Router Remote Site
Procedures for Adding a Second DMVPN for a Single Router Remote Site
1. Configure the WAN-facing VRF
2. Connect to the Internet
3. Configure IKEv2 and IPSec
4. Configure mGRE Tunnel
5. Configure EIGRP
6. Configure IP Multicast Routing
Add LTE Fallback?
YES
NO
Procedures for Adding LTE Fallback DMVPN for a Single Router Remote Site
1. Install LTE EHWIC into ISR
2. Configure Chat Script
3. Configure the WAN-facing VRF
4. Connect to the Cellular Provider
5. Configure the Dialer Watch List
6. Configure VRF-specific Default Routing
7. Configure IKEv2 and IPsec
8. Configure the mGRE Tunnel
9. Configure EIGRP
10. Configure IP Multicast Routing
11. Enable the Cellular Interface
12. Control Usage of LTE Fallback Tunnel
Site Complete
1238

Procedure 1 Configure the WAN-facing VRF

You create a WAN-facing VRF in order to support FVRF for DMVPN. The VRF name is arbitrary, but it is useful to select a name that describes the VRF. The VRF must be enabled for IPv4.

Table 29 VRF assignments

IWAN design model	Primary WAN VRF	Secondary WAN VRF
Hybrid	IWAN-TRANSPORT-1	IWAN-TRANSPORT-2
Dual Internet	IWAN-TRANSPORT-3	IWAN-TRANSPORT-4

This design uses VRF Lite, so the selection is only locally significant to the device. It is a best practice to use the same VRF/RD combination across multiple devices when using VRFs in a similar manner. However, this convention is not strictly required.

Step 1: Configure the secondary WAN VRF.

Example: Secondary WAN in the hybrid design model

```
vrf definition IWAN-TRANSPORT-2
 address-family ipv4
```

Procedure 2 Connect to the Internet

The remote sites that are using DMVPN can use either static or dynamically assigned IP addresses. Cisco tested the design with DHCP assigned external addresses for Internet connections, which also provides a dynamically configured default route.

The DMVPN spoke router connects directly to the Internet without a separate firewall. This connection is secured in two ways. Because the Internet interface is in a separate VRF, no traffic can access the global VRF except traffic sourced through the DMVPN tunnel. This design provides implicit security. Additionally, an IP access list permits only the traffic required for an encrypted tunnel, as well as DHCP and various ICMP protocols for troubleshooting.

Step 1: Enable the interface, select VRF and enable DHCP.

The DMVPN design uses FVRF, so you must place this interface into the VRF configured in the previous procedure.

```
interface GigabitEthernet0/1
 ip vrf forwarding IWAN-TRANSPORT-2
 ip address dhcp
 no cdp enable
 no shutdown
```

Step 2: Configure and apply the access list.

The IP access list must permit the protocols specified in the following table. The access list is applied inbound on the WAN interface, so filtering is done on traffic destined to the router.

Table 30 Required DMVPN protocols

Name	Protocol	Usage
non500-isakmp	UDP 4500	IPsec via NAT-T
isakmp	UDP 500	ISAKMP
esp	IP 50	IPsec
bootpc	UDP 68	DHCP

Example

```
interface GigabitEthernet0/1
 ip access-group ACL-INET-PUBLIC in

ip access-list extended ACL-INET-PUBLIC
 permit udp any any eq non500-isakmp
 permit udp any any eq isakmp
 permit esp any any
 permit udp any any eq bootpc
```

The additional protocols listed in the following table may assist in troubleshooting, but are not explicitly required to allow DMVPN to function properly.

Table 31 Optional protocols: DMVPN spoke router

Name	Protocol	Usage
icmp echo	ICMP Type 0, Code 0	Allow remote pings
icmp echo-reply	ICMP Type 8, Code 0	Allow ping replies (from your requests)
icmp ttl-exceeded	ICMP Type 11, Code 0	Allow traceroute replies (from your requests)
icmp port-unreachable	ICMP Type 3, Code 3	Allow traceroute replies (from your requests)
UDP high ports	UDP > 1023, TTL=1	Allow remote traceroute

The additional optional entries for an access list to support ping are as follows:

```
permit icmp any any echo
permit icmp any any echo-reply
```

The additional optional entries for an access list to support traceroute are as follows:

```
permit icmp any any ttl-exceeded        ! for traceroute (sourced)
permit icmp any any port-unreachable    ! for traceroute (sourced)
permit udp any any gt 1023 ttl eq 1     ! for traceroute (destination)
```

Procedure 3 Configure IKEv2 and IPsec

This procedure uses the parameters in the tables below. Choose the table that represents the design model that you are configuring. This procedure applies to the Secondary WAN.

Table 32 Crypto parameters: IWAN hybrid design model

Parameter	Primary WAN	Secondary WAN
vrf	IWAN-TRANSPORT-1	IWAN-TRANSPORT-2
crypto ikev2 keyring	DMVPN-KEYRING-1	DMVPN-KEYRING-2
crypto ikev2 profile	FVRF-IKEv2-IWAN-TRANSPORT-1	FVRF-IKEv2-IWAN-TRANSPORT-2
crypto ipsec profile	DMVPN-PROFILE-TRANSPORT-1	DMVPN-PROFILE-TRANSPORT-2

Table 33 Crypto parameters: IWAN dual Internet design model

Parameter	Primary WAN	Secondary WAN
vrf	IWAN-TRANSPORT-3	IWAN-TRANSPORT-4
crypto ikev2 keyring	DMVPN-KEYRING-3	DMVPN-KEYRING-4
crypto ikev2 profile	FVRF-IKEv2-IWAN-TRANSPORT-3	FVRF-IKEv2-IWAN-TRANSPORT-4
crypto ipsec profile	DMVPN-PROFILE-TRANSPORT-3	DMVPN-PROFILE-TRANSPORT-4

IPsec uses a key exchange between the routers in order to encrypt/decrypt the traffic. These keys can be exchanged using pre-shared keys or PKI certificates with a certificate authority. It is also possible to use a combination of the two, which is useful during a migration from one method to the other. Choose one of two options below as your method of key exchange.

Option 1: Configure with Pre-Shared Keys

Step 1: Configure the crypto keyring for pre-shared keys.

The crypto keyring defines a pre-shared key (or password) valid for IP sources that are reachable within a particular VRF. This key is a wildcard pre-shared key if it applies to any IP source. A wildcard key is configured using the 0.0.0.0 0.0.0.0 network/mask combination.

```
crypto ikev2 keyring [keyring name]
 peer ANY
  address 0.0.0.0 0.0.0.0
  pre-shared-key [password]
```

Example: Secondary WAN in the hybrid design model

```
crypto ikev2 keyring DMVPN-KEYRING-2
 peer ANY
  address 0.0.0.0 0.0.0.0
  pre-shared-key c1sco123
```

Step 2: Configure the IKEv2 profile.

The default IKEv2 proposal is used, which includes the following:

- Encryption with AES cipher with a 256-bit key
- Integrity with SHA with 512-bit digest
- Pseudo-random function with SHA with 512-bit digest
- Diffie-Hellman group: 5

```
RS11-2921# show crypto ikev2 proposal
IKEv2 proposal: default
    Encryption: AES-CBC-256 AES-CBC-192 AES-CBC-128
    Integrity : SHA512 SHA384 SHA256 SHA96 MD596
    PRF       : SHA512 SHA384 SHA256 SHA1 MD5
    DH Group  : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2
```

The IKEv2 profile creates an association between an identity address, a VRF, and a crypto keyring. A wildcard address within a VRF is referenced with 0.0.0.0. The profile also defines what method of key sharing will be used on this router with **authentication local** and what methods will be accepted from remote locations with **authentication remote**.

The **pre-share** keyword is used with the keyring defined above.

```
crypto ikev2 profile [profile name]
 match fvrf [vrf name]
 match identity remote address [IP address]
 authentication remote pre-share
 authentication local pre-share
 keyring local [keyring name]
```

Example: Secondary WAN in the hybrid design model

```
crypto ikev2 profile FVRF-IKEv2-IWAN-TRANSPORT-2
 match fvrf IWAN-TRANSPORT-2
 match identity remote address 0.0.0.0
 authentication remote pre-share
 authentication local pre-share
 keyring local DMVPN-KEYRING-2
```

Step 3: Configure the IPsec profile.

The IPsec profile creates an association between an IKEv2 profile and an IPsec transform-set.

```
crypto ipsec profile [profile name]
 set transform-set [transform set]
 set ikev2-profile [ikev2 profile name]
```

Example: Secondary WAN in the hybrid design model

```
crypto ipsec profile DMVPN-PROFILE-TRANSPORT-2
 set transform-set AES256/SHA/TRANSPORT
 set ikev2-profile FVRF-IKEv2-IWAN-TRANSPORT-2
```

Step 4: Proceed to Procedure 4, “Configure the mGRE Tunnel.”

Option 2: Configure with a certificate authority

The crypto pki trustpoint is the method of specifying the parameters associated with a CA. This router has already authenticated to the CA and enrolled in order to obtain its identity certificate in a previous procedure.

Step 1: Configure the IKEv2 profile.

The default IKEv2 proposal is used, which includes the following:

- Encryption with AES cipher with a 256-bit key
- Integrity with SHA with 512-bit digest
- Pseudo-random function with SHA with 512-bit digest
- Diffie-Hellman group: 5

```
RS11-2921# show crypto ikev2 proposal
IKEv2 proposal: default
    Encryption: AES-CBC-256 AES-CBC-192 AES-CBC-128
    Integrity : SHA512 SHA384 SHA256 SHA96 MD596
    PRF       : SHA512 SHA384 SHA256 SHA1 MD5
    DH Group  : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2
```

The IKEv2 profile creates an association between an identity address, a VRF, and a crypto keyring. A wildcard address within a VRF is referenced with 0.0.0.0. The profile also defines what method of key sharing will be used on this router with **authentication local** and what methods will be accepted from remote locations with **authentication remote**. The **rsa-sig** keyword is used when certificates contain the encryption key.

```
crypto ikev2 profile [profile name]
 match fvrf [vrf name]
 match identity remote address [IP address]
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint [trustpoint name]
```

Example: Secondary WAN in the hybrid design model

```
crypto ikev2 profile FVRF-IKEv2-IWAN-TRANSPORT-2
 match fvrf IWAN-TRANSPORT-2
 match identity remote address 0.0.0.0
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint IWAN-CA
```

Step 2: Configure the IPsec profile.

The IPsec profile creates an association between an IKEv2 profile and an IPsec transform-set.

```
crypto ipsec profile [profile name]
 set transform-set [transform set]
 set ikev2-profile [ikev2 profile name]
```

Example: Secondary WAN in the hybrid design model

```
crypto ipsec profile DMVPN-PROFILE-TRANSPORT-2
 set transform-set AES256/SHA/TRANSPORT
 set ikev2-profile FVRF-IKEv2-IWAN-TRANSPORT-2
```

Procedure 4 Configure the mGRE Tunnel

This procedure uses the parameters in the table below. Choose the rows that represent the design model that you are configuring. This procedure applies to the secondary WAN.

Table 34 DMVPN tunnel parameters

Design model	Tunnel VRF	Tunnel number	Tunnel network	NHRP network ID/ tunnel key
Hybrid–Primary WAN	IWAN-TRANSPORT-1	10	10.6.34.0/23	101
Hybrid–Secondary WAN	IWAN-TRANSPORT-2	11	10.6.36.0/23	102
Dual Internet–Primary WAN	IWAN-TRANSPORT-3	20	10.6.38.0/23	201
Dual Internet–Secondary WAN	IWAN-TRANSPORT-4	21	10.6.40.0/23	202

Step 1: Configure the basic interface settings.

The tunnel number is arbitrary, but it is best to begin tunnel numbering at 10 or above, because other features deployed in this design may also require tunnels and they may select lower numbers by default.

The bandwidth setting must be set to match the bandwidth of the respective transport, which corresponds to the actual interface speed. Or, if you are using a subrate service, use the policed rate from the carrier. QoS and PfR require the correct bandwidth setting in order to operate properly.

Configure the **ip mtu** to 1400 and the **ip tcp adjust-mss** to 1360. There is a 40 byte difference, which corresponds to the combined IP and TCP header length.

The tunnel interface throughput delay setting should be set to influence the routing protocol path preference. Set the primary WAN path to 10000 usec and the secondary WAN path to 200000 usec to prefer one over the other. The delay command is entered in 10 usec units.

```
interface Tunnel11
 bandwidth 200000
 ip address 10.6.36.11 255.255.254.0
 no ip redirects
 ip mtu 1400
 ip tcp adjust-mss 1360
 delay 20000
```

Step 2: Configure the tunnel.

DMVPN uses mGRE tunnels. This type of tunnel requires a source interface only. Use the same source interface that you use to connect to the Internet. Set the **tunnel vrf** command should be set to the VRF defined previously for FVRF.

Enabling encryption on this interface requires the application of the IPsec profile configured in the previous procedure.

```
interface Tunnel11
 tunnel source GigabitEthernet0/1
 tunnel mode gre multipoint
 tunnel key 102
 tunnel vrf IWAN-TRANSPORT-2
 tunnel protection ipsec profile DMVPN-PROFILE-TRANSPORT-2
```

Step 3: Configure NHRP.

The DMVPN hub router is the NHRP server for all of the spokes. NHRP is used by remote routers to determine the tunnel destinations for peers attached to the mGRE tunnel.

The spoke router requires several additional configuration statements to define the NHRP server and NHRP map statements for the DMVPN hub router mGRE tunnel IP address. EIGRP (configured in the following Procedure 5) relies on a multicast transport. Spoke routers require the NHRP static multicast mapping.

When hub BRs are added for horizontal scaling or a second data center is added as a transit site, spoke routers require additional NHS statements for each BR in their environment. The configuration details are covered in subsequent sections of this guide.

The value used for the NHS is the mGRE tunnel address for the DMVPN hub router. The map entries must be set to the outside NAT value of the DMVPN hub, as configured on the Cisco ASA 5500.

This design uses the values shown in the following tables.

Table 35 DMVPN tunnel NHRP parameters: IWAN hybrid design model

	Transport 1	Transport 2
VRF	IWAN-TRANSPORT-1	IWAN-TRANSPORT-2
DMVPN hub public address (actual)	192.168.6.1	192.168.146.10
DMVPN hub public address (externally routable after NAT)	n/a (MPLS)	172.16.140.1
DMVPN hub tunnel IP address (NHS)	10.6.34.1	10.6.36.1
Tunnel number	10	11
NHRP network ID	101	102

Table 36 DMVPN tunnel NHRP parameters: IWAN dual Internet design model

	Transport 3	Transport 4
VRF	IWAN-TRANSPORT-3	IWAN-TRANSPORT-4
DMVPN hub public address (actual)	192.168.146.20	192.168.146.21
DMVPN hub public address (externally routable after NAT)	172.16.140.11	172.17.140.11
DMVPN hub tunnel IP address (NHS)	10.6.38.1	10.6.40.1
Tunnel number	20	21
NHRP network ID	201	202

NHRP requires all devices within a DMVPN cloud to use the same network ID and authentication key. The NHRP cache holdtime should be configured to 600 seconds.

This design supports DMVPN spoke routers that receive their external IP addresses through DHCP. It is possible for these routers to acquire different IP addresses after a reload. When the router attempts to register with the NHRP server, it may appear as a duplicate to an entry already in the cache and be rejected. The **registration no-unique** option allows you to overwrite existing cache entries. This feature is only required on NHRP clients (DMVPN spoke routers). The **if-state nhrp** option ties the tunnel line-protocol state to the reachability of the NHRP NHS, and if the NHS is unreachable the tunnel line-protocol state changes to down. This feature is used in conjunction with EOT.

```
interface Tunnel11
 ip nhrp authentication cisco123
 ip nhrp network-id 102
 ip nhrp holdtime 600
 ip nhrp nhs 10.6.36.1 nbma 172.16.140.1 multicast
 ip nhrp registration no-unique
 ip nhrp shortcut
 if-state nhrp
```

By default, NHRP will not install shortcuts for paths not seen in the RIB of the router. In a location with a single router and multiple WAN transports, only the preferred path is in the RIB. If you have a remote site location with more than one WAN transport, you need to disable the **nhrp route-watch** feature on each of the tunnel interfaces in order to allow NHRP to install the non-preferred shortcut path.

```
interface Tunnel11
 no nhrp route-watch
```

Procedure 5 Configure EIGRP

A single EIGRP process runs on the DMVPN spoke router, which has already been enabled during the first DMVPN tunnel's configuration. All interfaces on the router are EIGRP interfaces, but only the DMVPN tunnel interfaces are non-passive. The network range must include all interface IP addresses either in a single network statement or in multiple network statements.

Step 1: Add the network range for the secondary DMVPN tunnel and configure as non-passive.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   af-interface Tunnel11
    no passive-interface
   exit-af-interface
 network 10.6.36.0 0.0.1.255
exit-address-family
```

Step 2: Configure EIGRP values for the mGRE tunnel interface.

The EIGRP hello interval is increased to 20 seconds and the EIGRP hold time is increased to 60 seconds in order to accommodate up to 2000 remote sites on a single DMVPN cloud. Increasing the EIGRP timers also slows down the routing convergence to improve network stability and the IWAN design allows PfR to initiate the fast failover, so changing the timers is recommended for all IWAN deployments.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   af-interface Tunnel11
    hello-interval 20
    hold-time 60
   exit-af-interface
 exit-address-family
```

Step 3: Configure EIGRP neighbor authentication. Neighbor authentication enables the secure establishment of peering adjacencies and exchange route tables over the DMVPN tunnel interface.

```
key chain WAN-KEY
 key 1
  key-string c1sco123

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel11
   authentication mode md5
   authentication key-chain WAN-KEY
  exit-af-interface
 exit-address-family
```

Step 4: Configure EIGRP route summarization.

The remote-site LAN networks must be advertised. The IP assignment for the remote sites was designed so that all of the networks in use can be summarized within a single aggregate route. As configured below, the **summary-address** command suppresses the more specific routes. If any network within the summary is present in the route table, the summary is advertised to the DMVPN hub, which offers a measure of resiliency. If the various LAN networks cannot be summarized, EIGRP continues to advertise the specific routes.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel11
   summary-address [summary network] [summary mask]
  exit-af-interface
 exit-address-family
```

Step 5: Tag and filter the routes.

This design uses a single EIGRP autonomous system for the hub sites, WAN, and all of the WAN remote sites. Every remote site is dual-connected for resiliency. However, due to the multiple paths that exist within this topology, you must try to avoid routing loops and to prevent remote sites from becoming transit sites if WAN failures were to occur.

In this design, there are different IP subnets for each DMVPN network, and the EIGRP tags are defined for each segment in order to help with readability and troubleshooting. When a design uses more than one data center, additional tags are required in order to identify the different DMVPN hub router locations.

The following logic is used to control the routing.

- Each DMVPN network will have an EIGRP route tag to prevent routes from being re-advertised over the other DMVPN networks.
- All prefixes that are advertised towards the WAN are uniquely tagged.
- All DMVPN learned WAN prefixes, except those that originate locally from a hub, are advertised towards the LAN and tagged.
- The design always uses DMVPN hub routers in pairs. Each DMVPN hub router blocks DMVPN WAN routes from the LAN that are tagged with the opposite hub's tags.
- Remote sites can learn routes from other remote sites when NHRP uses shortcut routing to bypass the hub locations. The spoke-to-spoke routes are blocked using tag filtering to prevent issues with PfR.

Outbound distribute-lists are used to set tags towards the WAN and LAN. The remote-site routers use the tags set towards the WAN in order to protect against becoming transit sites.

The remote-site routers use an inbound distribute-list in order to limit which routes are accepted for installation into the route table. These routers are configured to only accept routes that do not originate from the WAN sources. To accomplish this task, the remote site router must explicitly tag the DMVPN learned WAN routes. The following tables show specific route tags in use.

Table 37 Route tag information for IWAN hybrid remote site routers

Tunnel: transport	Tunnel key	Match inbound	Block outbound
Tunnel10: DMVPN1	101 (MPLS)	101 (All routes)	101 (WAN routes)
Tunnel11: DMVPN2	102 (INET)	102 (All routes)	102 (WAN routes)

Table 38 Route tag information for IWAN dual Internet remote site routers

Tunnel: transport	Tunnel key	Match inbound	Block outbound
Tunnel20 DMVPN3	201 (INET1)	201 (All routes)	201 (WAN routes)
Tunnel21: DMVPN4	202 (INET2)	202 (All routes)	202 (WAN routes)

The following example shows a single router design in the IWAN hybrid design model.

Example: Single Router Site with DMVPN1 and DMVPN2

```
route-map DMVPN1-BR-IN Permit 10
 description Match tagged routes inbound
 match tag 101

route-map DMVPN2-BR-IN Permit 10
 description Match tagged routes inbound
 match tag 102

route-map BLOCK-LEARNED deny 10
 description Block learned routes outbound
 Match tag 101 102

route-map BLOCK-LEARNED permit 20
 description Advertise all other routes outbound

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  topology base
   distribute-list route-map DMVPN1-BR-IN in Tunnel10
   distribute-list route-map DMVPN2-BR-IN in Tunnel11
   distribute-list route-map BLOCK-LEARNED out Tunnel10
   distribute-list route-map BLOCK-LEARNED out Tunnel11
```

Procedure 6 Configure IP multicast routing

Optional

This optional procedure includes additional steps for configuring IP Multicast for a DMVPN tunnel on a router with IP Multicast already enabled. Skip this procedure if you do not want to use IP Multicast in your environment.

Step 1: Configure PIM on the DMVPN tunnel interface.

Enable IP PIM sparse mode on the DMVPN tunnel interface.

```
interface Tunnel11
 ip pim sparse-mode
```

Step 2: Enable PIM NBMA mode for the DMVPN tunnel.

Spoke-to-spoke DMVPN networks present a unique challenge because the spokes cannot directly exchange information with one another, even though they are on the same logical network. This inability to directly exchange information can also cause problems when running IP Multicast.

To resolve the NBMA issue, you need to implement a method where each remote PIM neighbor has its join messages tracked separately. A router in PIM NBMA mode treats each remote PIM neighbor as if it were connected to the router through a point-to-point link.

```
interface Tunnel11
 ip pim nbma-mode
```

Step 3: Configure the DR priority for the DMVPN spoke router.

Proper multicast operation across a DMVPN cloud requires that the hub router assumes the role of PIM DR. Spoke routers should never become the DR. You can prevent that by setting the DR priority to 0 for the spokes.

```
interface Tunnel11
 ip pim dr-priority 0
```

PROCESS

Adding LTE fallback DMVPN for a single-router remote site

1. Install LTE EHWIC into ISR
2. Configure chat script
3. Configure the WAN-facing VRF
4. Connect to the cellular provider
5. Configure the dialer watch list
6. Configure VRF-specific default routing
7. Configure IKEv2 and IPsec
8. Configure the mGRE Tunnel
9. Configure EIGRP
10. Configure IP multicast routing
11. Enable the cellular interface
12. Control usage of LTE fallback tunnel

This set of procedures includes the additional steps necessary to add a third fallback DMVPN link to a remote-site router that has already been configured primary and secondary DMVPN links by using the following processes in this guide:

- “Configuring Remote-Site DMVPN Router”
- “Adding Second DMVPN for a Single-Router Remote Site”

This section includes only the additional procedures for adding the LTE fallback DMVPN to the running remote-site router.

This section is specific to cellular LTE devices used to test this document. There are other Cisco products that share common configuration with the devices mentioned that may have different packages (Cisco Enhanced High-Speed WAN Interface Card [EHWIC] vs. router) and different carriers, such as Verizon, T-Mobile or Sprint. You must get a data service account from your service provider. You should receive a SIM card that you install on the LTE EHWIC, no matter the carrier.

There are vendor specific variations of 4G/LTE HWICs, some with geographically specific firmware. The table below shows the version of the 4G/LTE card validated in this guide and the version of firmware tested. Additional specific geographic and carrier information for the various Cisco cellular WAN access interfaces can be found online at: http://www.cisco.com/c/en/us/products/routers/networking_solutions_products_genericcontent0900aecd80601f7e.html

Table 39 *GSM 4G/LTE specific HWICs*

Part number	Modem	Carrier	Firmware version	Firmware date	Remote site
EHWIC-4G-LTE-A	MC7700	AT&T	SWI9200X_03.05.10.02	2012/02/25 11:58:38	RS51

Procedure 1 Install LTE EHWIC into ISR

Figure 24 *LTE EHWIC SIM card installation*

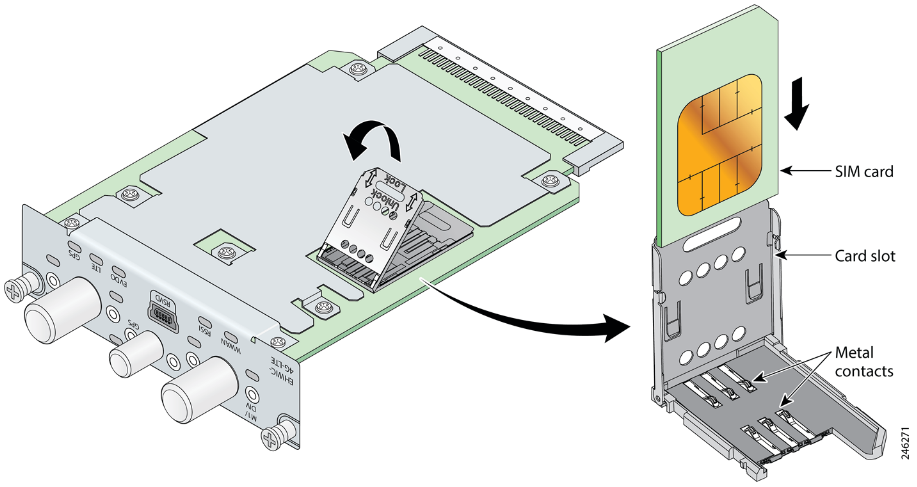

Step 1: Insert the SIM card into the EHWIC.

Step 2: Power down the Integrated Services G2 router.

Step 3: Insert and fasten the LTE EHWIC into the router.

Step 4: Power up the router, and then begin configuration.

Procedure 2 Configure chat script

Chat scripts are strings of text used to send commands for modem dialing, to log in to remote systems, and to initialize asynchronous devices connected to an asynchronous line. The 4G WAN interface should be treated just like any other asynchronous interface.

The following chat script shows the required information to connect to the Verizon or AT&T LTE network. It uses an LTE-specific dial string and a timeout value of 30 seconds. Note that your carrier may require a different chat script.

Step 1: Create the chat script.

```
chat-script [Script-Name] [Script]
```

Example

```
chat-script LTE "" "AT!CALL1" TIMEOUT 30 "OK"
```

Step 2: Apply the chat script to the asynchronous line.

```
line [Cellular-Interface-Number]
  script dialer [Script-Name]
```

Example

For the interface cellular0/1/0, the matching line would be as follows.

```
line 0/1/0
  script dialer LTE
```

Procedure 3 Configure the WAN-facing VRF

You create a WAN-facing VRF in order to support FVRF for DMVPN. The VRF name is arbitrary, but it is useful to select a name that describes the VRF. The VRF must be enabled for IPv4.

Table 40 VRF assignments

IWAN design model	Primary WAN VRF	Secondary WAN VRF	LTE Fallback VRF
Hybrid	IWAN-TRANSPORT-1	IWAN-TRANSPORT-2	IWAN-TRANSPORT-3
Dual Internet	IWAN-TRANSPORT-3	IWAN-TRANSPORT-4	IWAN-TRANSPORT-1

This design uses VRF Lite, so the selection is only locally significant to the device. It is a best practice to use the same VRF/RD combination across multiple devices when using VRFs in a similar manner. However, this convention is not strictly required.

Step 1: Configure the LTE fallback VRF.

Example: LTE fallback in the IWAN hybrid design model

```
vrf definition IWAN-TRANSPORT-3
 address-family ipv4
```

Procedure 4 Connect to the cellular provider

You add the cellular interface to a dialer watch group and to the VRF. You set the bandwidth value to match the minimum uplink speed of the chosen technology, as shown in the following table. Configure the interface as administratively down until the rest of the configuration steps are complete.

Table 41 4G encapsulation and bandwidth parameters

Cellular keyword	Encapsulation	Cellular script name (created previously)	Downlink speed (Kbps)	Uplink speed (Kbps)
LTE	Direct IP (SLIP)	LTE	8000 to 12,000 (range)	2000 to 5000 (range)

Tech Tip

LTE cellular interfaces use Direct IP encapsulation. When configuring Direct IP encapsulation, use the serial line Internet protocol (SLIP) keyword.

Step 1: Configure the cellular interface.

```
interface Cellular [Interface-Number]
 bandwidth [bandwidth (Kbps)]
 vrf forwarding IWAN-TRANSPORT-3
 ip address negotiated
 no ip unreachables
 ip virtual-reassembly in
 encapsulation [encapsulation type]
 dialer in-band
 dialer idle-timeout 0
 dialer string [Chat Script Name]
 dialer watch-group 1
 no peer default ip address
 async mode interactive
 shutdown
```

Example: LTE Bandwidth and Encapsulation

```
interface Cellular0/1/0
 bandwidth 2000
 ip vrf forwarding IWAN-TRANSPORT-3
 ip address negotiated
 no ip unreachables
 ip virtual-reassembly in
 encapsulation slip
 dialer in-band
 dialer idle-timeout 0
 dialer string LTE
 dialer watch-group 1
 no peer default ip address
 async mode interactive
 shutdown
```

Step 2: Configure and apply the access list.

The IP access list must permit the protocols specified in the following table. The access list is applied inbound on the WAN interface, so filtering is done on traffic destined to the router.

Table 42 Required DMVPN protocols

Name	Protocol	Usage
non500-isakmp	UDP 4500	IPsec via NAT-T
isakmp	UDP 500	ISAKMP
esp	IP 50	IPsec

```
interface Cellular0/1/0
 ip access-group ACL-INET-PUBLIC-4G in

ip access-list extended ACL-INET-PUBLIC-4G
 permit udp any any eq non500-isakmp
 permit udp any any eq isakmp
 permit esp any any
 permit udp any any eq bootpc
```

The additional protocols listed in the following table may assist in troubleshooting but are not explicitly required to allow DMVPN to function properly.

Table 43 Optional protocols: DMVPN spoke router

Name	Protocol	Usage
icmp echo	ICMP Type 0, Code 0	Allow remote pings
icmp echo-reply	ICMP Type 8, Code 0	Allow ping replies (from your requests)
icmp ttl-exceeded	ICMP Type 11, Code 0	Allow traceroute replies (from your requests)
icmp port-unreachable	ICMP Type 3, Code 3	Allow traceroute replies (from your requests)
UDP high ports	UDP > 1023, TTL=1	Allow remote traceroute

The additional optional entries for an access list to support ping are as follows:

```
permit icmp any any echo
permit icmp any any echo-reply
```

The additional optional entries for an access list to support traceroute are as follows:

```
permit icmp any any ttl-exceeded       ! for traceroute (sourced)
permit icmp any any port-unreachable   ! for traceroute (sourced)
permit udp any any gt 1023 ttl eq 1    ! for traceroute (destination)
```

Procedure 5 Configure the dialer watch list

The *dialer watch-list* is a construct that allows the activation of the dialer script and associated cellular interface when the specified route no longer exists in the routing table. In this procedure, the dialer-watch list activates the cellular interface when the specified phantom route is missing from the routing table.

This design uses the IANA-specified loopback address of 127.0.0.255, which should never appear in the routing table under normal circumstances. The absence of this route in the routing table causes the cellular interface to become active and stay active until the interface is brought down.

Step 1: Assign a phantom route to the **dialer watch-list**. Use the same value as the **dialer watch-group** in the previous procedure.

```
dialer watch-list 1 ip 127.0.0.255 255.255.255.255
dialer watch-list 1 delay route-check initial 60
dialer watch-list 1 delay connect 1
```

Procedure 6 Configure VRF-specific default routing

The remote sites using 3G or 4G DMVPN use negotiated IP addresses for the cellular interfaces. Unlike DHCP, the negotiation does not automatically set a default route. This step must be completed manually.

Step 1: Configure a VRF-specific default route for the cellular interface.

```
ip route vrf IWAN-TRANSPORT-3 0.0.0.0 0.0.0.0 Cellular0/1/0
```

Procedure 7 Configure IKEv2 and IPsec

This procedure uses the parameters in the following table. Choose the column that represents the design model that you are configuring. This procedure applies to the LTE Fallback DMVPN using either the IWAN hybrid design model or the IWAN dual Internet design model.

Table 44 Crypto parameters: LTE fallback DMVPN

Parameter	IWAN hybrid design model	IWAN dual Internet design model
vrf	IWAN-TRANSPORT-3	IWAN-TRANSPORT-1
crypto ikev2 keyring	DMVPN-KEYRING-3	DMVPN-KEYRING-1
crypto ikev2 profile	FVRF-IKEv2-IWAN-TRANSPORT-3	FVRF-IKEv2-IWAN-TRANSPORT-1
crypto ipsec profile	DMVPN-PROFILE-TRANSPORT-3	DMVPN-PROFILE-TRANSPORT-1

IPsec uses a key exchange between the routers in order to encrypt/decrypt the traffic. These keys can be exchanged using pre-shared keys or PKI certificates with a certificate authority. It is also possible to use a combination of the two, which is useful during a migration from one method to the other. Choose one of two options below as your method of key exchange.

Option 1: Configure with Pre-Shared Keys

Step 1: Configure the crypto keyring for pre-shared keys.

The crypto keyring defines a pre-shared key (or password) valid for IP sources that are reachable within a particular VRF. This key is a wildcard pre-shared key if it applies to any IP source. A wildcard key is configured using the 0.0.0.0 0.0.0.0 network/mask combination.

```
crypto ikev2 keyring [keyring name]
 peer ANY
  address 0.0.0.0 0.0.0.0
  pre-shared-key [password]
```

Example: LTE fallback in the hybrid design model

```
crypto ikev2 keyring DMVPN-KEYRING-3
 peer ANY
  address 0.0.0.0 0.0.0.0
  pre-shared-key c1sco123
```

Step 2: Configure the IKEv2 profile.

The IKEv2 profile creates an association between an identity address, a VRF, and a crypto keyring. A wildcard address within a VRF is referenced with 0.0.0.0. The profile also defines what method of key sharing will be used on this router with **authentication local** and what methods will be accepted from remote locations with **authentication remote**. The **pre-share** keyword is used with the keyring defined above.

DPD is essential in order to facilitate fast reconvergence and for spoke registration to function properly in case a DMVPN hub is restarted. The IWAN design recommends you set the remote site DPD timer to 40 seconds with a 5 second retry.

```
crypto ikev2 profile [profile name]
 match fvrf [vrf name]
 match identity remote address [IP address]
 authentication remote pre-share
 authentication local pre-share
 keyring local [keyring name]
 dpd [interval in seconds] [retry interval] on-demand
```

Example: LTE fallback in the hybrid design model

```
crypto ikev2 profile FVRF-IKEv2-IWAN-TRANSPORT-3
 match fvrf IWAN-TRANSPORT-1
 match identity remote address 0.0.0.0
 authentication remote pre-share
 authentication local pre-share
 keyring local DMVPN-KEYRING-3
 dpd 40 5 on-demand
```

Step 3: Configure the IPsec profile.

The IPsec profile creates an association between an IKEv2 profile and an IPsec transform-set.

```
crypto ipsec profile [profile name]
 set transform-set [transform set]
 set ikev2-profile [ikev2 profile name]
```

Example: LTE fallback in the hybrid design model

```
crypto ipsec profile DMVPN-PROFILE-TRANSPORT-3
 set transform-set AES256/SHA/TRANSPORT
 set ikev2-profile FVRF-IKEv2-IWAN-TRANSPORT-3
```

Step 4: Skip to the next procedure.

Option 2: Configure with a Certificate Authority

The crypto pki trustpoint is the method of specifying the parameters associated with a CA. This router has already authenticated to the CA and enrolled to obtain its identity certificate in a previous procedure.

Step 1: Configure the IKEv2 profile.

The IKEv2 profile creates an association between an identity address, a VRF, and a crypto keyring. A wildcard address within a VRF is referenced with 0.0.0.0. The profile also defines what method of key sharing will be used on this router with **authentication local** and what methods will be accepted from remote locations with **authentication remote**. The **rsa-sig** keyword is used when certificates contain the encryption key.

DPD is essential in order to facilitate fast reconvergence and for spoke registration to function properly in case a DMVPN hub is restarted. The IWAN design recommends you set the remote site DPD timer to 40 seconds with a 5 second retry.

```
crypto ikev2 profile [profile name]
 match fvrf [vrf name]
 match identity remote address [IP address]
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint [trustpoint name]
 dpd [interval in seconds] [retry interval] on-demand
```

Example: LTE fallback in the hybrid design model

```
crypto ikev2 profile FVRF-IKEv2-IWAN-TRANSPORT-3
 match fvrf IWAN-TRANSPORT-3
 match identity remote address 0.0.0.0
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint IWAN-CA
 dpd 40 5 on-demand
```

Step 2: Configure the IPsec profile.

The IPsec profile creates an association between an IKEv2 profile and an IPsec transform-set.

```
crypto ipsec profile [profile name]
 set transform-set [transform set]
 set ikev2-profile [ikev2 profile name]
```

Example: LTE fallback in the hybrid design model

```
crypto ipsec profile DMVPN-PROFILE-TRANSPORT-3
 set transform-set AES256/SHA/TRANSPORT
 set ikev2-profile FVRF-IKEv2-IWAN-TRANSPORT-3
```

Procedure 8 Configure the mGRE Tunnel

This procedure uses the parameters in the table below. Choose the rows that represent the design model that you are configuring.

Table 45 LTE fallback DMVPN tunnel parameters

Design model	Tunnel VRF	Tunnel number	Tunnel network	NHRP network ID/ tunnel key
Hybrid	IWAN-TRANSPORT-3	20	10.6.38.0/23	201
Dual Internet	IWAN-TRANSPORT-1	10	10.6.34.0/23	101

Step 1: Configure basic interface settings.

The tunnel number is arbitrary, but it is best to begin tunnel numbering at 10 or above, because other features deployed in this design may also require tunnels and they may select lower numbers by default.

The bandwidth setting must be set to match the bandwidth of the respective transport which corresponds to the actual interface speed. Or, if you are using a subrate service, use the policed rate from the carrier. QoS and PfR require the correct bandwidth setting to operate properly.

Configure the **ip mtu** to 1400 and the **ip tcp adjust-mss** to 1360. There is a 40 byte difference, which corresponds to the combined IP and TCP header length.

The tunnel interface throughput delay setting is not needed because this is a tertiary path which will only be used when the other two paths are not available.

```
interface Tunnel20
 bandwidth 8000
 ip address 10.6.38.51
 no ip redirects
 ip mtu 1400
 ip tcp adjust-mss 1360
```

Step 2: Configure the tunnel.

DMVPN uses mGRE tunnels. This type of tunnel requires a source interface only. Use the same source interface that you use to connect to the Internet. Set the **tunnel vrf** command should be set to the VRF defined previously for FVRF.

Enabling encryption on this interface requires the application of the IPsec profile configured in the previous procedure.

```
interface Tunnel20
 tunnel source Cellular0/1/0
 tunnel mode gre multipoint
 tunnel key 201
 tunnel vrf IWAN-TRANSPORT-3
 tunnel protection ipsec profile DMVPN-PROFILE-TRANSPORT-3
```

Step 3: Configure NHRP.

The DMVPN hub router is the NHRP server for all of the spokes. NHRP is used by remote routers to determine the tunnel destinations for peers attached to the mGRE tunnel.

The spoke router requires several additional configuration statements to define the NHRP server and NHRP map statements for the DMVPN hub router mGRE tunnel IP address. EIGRP (configured in the following Procedure 5) relies on a multicast transport. Spoke routers require the NHRP static multicast mapping.

When hub BRs are added for horizontal scaling or a second data center is added as a transit site, each spoke router will require additional NHS statements for each BR in their environment. The configuration details are covered in subsequent sections of this guide.

The value used for the NHS is the mGRE tunnel address for the DMVPN hub router. The map entries must be set to the outside NAT value of the DMVPN hub, as configured on the Cisco ASA 5500. This design uses the values shown in the following tables.

Table 46 DMVPN tunnel NHRP parameters: IWAN hybrid design model

	LTE fallback
VRF	IWAN-TRANSPORT-3
DMVPN hub public address (actual)	192.168.146.20
DMVPN hub public address (externally routable after NAT)	172.16.140.11
DMVPN hub tunnel IP address (NHS)	10.6.38.1
Tunnel number	20
NHRP network ID	201

Table 47 LTE fallback DMVPN tunnel NHRP parameters: IWAN dual Internet design model

	LTE fallback
VRF	IWAN-TRANSPORT-1
DMVPN hub public address (actual)	192.168.146.10
DMVPN hub public address (externally routable after NAT)	172.16.140.1
DMVPN hub tunnel IP address (NHS)	10.6.34.1
Tunnel number	10
NHRP network ID	101

NHRP requires all devices within a DMVPN cloud to use the same network ID and authentication key. The NHRP cache holdtime should be configured to 600 seconds.

This design supports DMVPN spoke routers that receive their external IP addresses through DHCP. It is possible for these routers to acquire different IP addresses after a reload. When the router attempts to register with the NHRP server, it may appear as a duplicate to an entry already in the cache and be rejected. The **registration no-unique** option allows you to overwrite existing cache entries. This feature is only required on NHRP clients (DMVPN spoke routers). The **if-state nhrp** option ties the tunnel line-protocol state to the reachability of the NHRP NHS, and if the NHS is unreachable, the tunnel line-protocol state changes to down.

```
interface Tunnel20
 ip nhrp authentication cisco123
 ip nhrp network-id 201
 ip nhrp holdtime 600
 ip nhrp nhs 10.6.38.1 nbma 172.16.140.11 multicast
 ip nhrp registration no-unique
 ip nhrp shortcut
 if-state nhrp
```

Procedure 9 Configure EIGRP

A single EIGRP process runs on the DMVPN spoke router, which has already been enabled during the configuration of the first DMVPN tunnel. All interfaces on the router are EIGRP interfaces, but only the DMVPN tunnel interfaces are non-passive. The network range must include all interface IP addresses either in a single network statement or in multiple network statements.

Step 1: Add the network range for the LTE Fallback DMVPN tunnel and configure as non-passive.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   af-interface Tunnel20
    no passive-interface
   exit-af-interface
 network 10.6.38.0 0.0.1.255
exit-address-family
```

Step 2: Configure EIGRP values for the mGRE tunnel interface.

The EIGRP hello interval is increased to 20 seconds and the EIGRP hold time is increased to 60 seconds in order to accommodate up to 2000 remote sites on a single DMVPN cloud. Increasing the EIGRP timers also slows down the routing convergence to improve network stability and the IWAN design allows PfR to initiate the fast failover, so changing the timers is recommended for all IWAN deployments.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   af-interface Tunnel20
    hello-interval 20
    hold-time 60
   exit-af-interface
 exit-address-family
```

Step 3: Configure EIGRP neighbor authentication. Neighbor authentication enables the secure establishment of peering adjacencies and exchange route tables over the DMVPN tunnel interface.

```
key chain WAN-KEY
 key 1
  key-string c1sco123
!
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel20
   authentication mode md5
   authentication key-chain WAN-KEY
  exit-af-interface
 exit-address-family
```

Step 4: Configure EIGRP route summarization.

The remote-site LAN networks must be advertised. The IP assignment for the remote sites was designed so that all of the networks in use can be summarized within a single aggregate route. The summary address as configured below suppresses the more specific routes. If any network within the summary is present in the route table, the summary is advertised to the DMVPN hub, which offers a measure of resiliency. If the various LAN networks cannot be summarized, EIGRP continues to advertise the specific routes.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel20
   summary-address [summary network] [summary mask]
  exit-af-interface
 exit-address-family
```

Procedure 10 Configure IP multicast routing

Optional

This optional procedure includes additional steps for configuring IP Multicast for a DMVPN tunnel on a router with IP Multicast already enabled. Skip this procedure if you do not want to use IP Multicast in your environment.

Step 1: Configure PIM on the DMVPN tunnel interface.

Enable IP PIM sparse mode on the DMVPN tunnel interface.

```
interface Tunnel20
 ip pim sparse-mode
```

Step 2: Enable PIM NBMA mode for the DMVPN tunnel.

Spoke-to-spoke DMVPN networks present a unique challenge because the spokes cannot directly exchange information with one another, even though they are on the same logical network. This inability to directly exchange information can also cause problems when running IP Multicast.

To resolve the NBMA issue, you need to implement a method where each remote PIM neighbor has its join messages tracked separately. A router in PIM NBMA mode treats each remote PIM neighbor as if it were connected to the router through a point-to-point link.

```
interface Tunnel20
 ip pim nbma-mode
```

Step 3: Configure the DR priority for the DMVPN spoke router.

Proper multicast operation across a DMVPN cloud requires that the hub router assumes the role of PIM DR. Spoke routers should never become the DR. You can prevent that by setting the DR priority to 0 for the spokes.

```
interface Tunnel20
 ip pim dr-priority 0
```

Procedure 11 Enable the cellular interface

The 4G/LTE portion of the router configuration is essentially complete.

Step 1: Enable the cellular interface to bring up the DMVPN tunnel.

```
interface Cellular0/1/0
 no shutdown
```

Procedure 12 Control usage of LTE fallback tunnel

Many 4G/LTE service providers do not offer a mobile data plan with unlimited usage. More typically, you will need to select a usage-based plan with a bandwidth tier that aligns with the business requirements for the remote site. To minimize recurring costs of the 4G/LTE solution, it is a best practice to limit the use of the wireless WAN specifically to a backup-only path.

The remote-site router can use EOT to track the status of the DMVPN hub routers for the primary and secondary links. If both become unreachable, then the router can use the Embedded Event Manager (EEM) to dynamically enable the cellular interface.

Step 1: Configure EOT to track the interface state of primary and secondary tunnels.

This step links the status of each interface to a basic EOT object.

```
track 10 interface Tunnel10 line-protocol
track 11 interface Tunnel11 line-protocol
```

Step 2: Configure composite object tracking.

A track list using Boolean OR is Up if either basic object is Up, and changes state to Down only when both basic objects are Down. This logic permits either the primary or secondary DMVPN tunnel to fail without enabling the LTE fallback tunnel. Both the primary and secondary tunnels must be down before the LTE fallback tunnel is enabled.

A short delay of 20 seconds is added when the primary or secondary tunnels are restored before shutting down the cellular interface.

```
track 20 list boolean or
 object 10
 object 11
 delay up 20
```

Step 3: Configure EEM scripting to enable or disable the cellular interface.

An event-tracking EEM script monitors the state of an object and runs router Cisco IOS commands for that particular state. It is also a best practice to generate syslog messages that provide status information regarding EEM.

```
event manager applet [EEM script name]
 event track [object number] state [tracked object state]
 action [sequence 1]  cli command "[command 1]"
 action [sequence 2]  cli command "[command 2]"
 action [sequence 3]  cli command "[command 3]"
 action [sequence …]  cli command "[command …]"
 action [sequence N] syslog msg "[syslog message test]"
```

Example: EEM script to enable the cellular interface.

```
event manager applet ACTIVATE-LTE
 event track 20 state down
 action 1  cli command "enable"
 action 2  cli command "configure terminal"
 action 3  cli command "interface cellular0/1/0"
 action 4  cli command "no shutdown"
 action 5  cli command "end"
 action 99 syslog msg "Both tunnels down - Activating 4G interface"
```

Example: EEM script to disable the cellular interface.

```
event manager applet DEACTIVATE-LTE
 event track 20 state up
 action 1  cli command "enable"
 action 2  cli command "configure terminal"
 action 3  cli command "interface cellular0/1/0"
 action 4  cli command "shutdown"
 action 5  cli command "end"
 action 99 syslog msg "Connectivity Restored - Deactivating 4G interface "
```

PROCESS

Modifying the First Router for Dual Router Design

1. Configure access layer HSRP
2. Configure transit network
3. Configure EIGRP on the transit network
4. Enable enhanced object tracking

This process is required when the first router has already been configured using the "Configuring Remote-Site DMVPN Router" process.

Procedure 1 Configure access layer HSRP

You need to configure HSRP to enable the use of a Virtual IP (VIP) as a default gateway that is shared between two routers. The HSRP active router is the router connected to the primary WAN link and the HSRP standby router is the router connected to the secondary WAN link.

Step 1: Configure the HSRP active router with a standby priority that is higher than the HSRP standby router.

The router with the higher standby priority value is elected as the HSRP active router. The preempt option allows a router with a higher priority to become the HSRP active, without waiting for a scenario where there is no router in the HSRP active state. The following table shows the relevant HSRP parameters for the router configuration.

Table 48 WAN remote-site HSRP parameters (dual router)

Router	HSRP role	VIP	Real IP address	HSRP priority	PIM DR priority
Primary	Active	.1	.2	110	110
Secondary	Standby	.1	.3	105	105

The assigned IP addresses override those configured in the previous procedure, so the default gateway IP address remains consistent across locations with single or dual routers.

The dual-router access-layer design requires a modification for resilient multicast. The PIM DR should be on the HSRP active router. The DR is normally elected based on the highest IP address, and has no awareness of the HSRP configuration. In this design, the HSRP active router has a lower real IP address than the HSRP standby router, which requires a modification to the PIM configuration. The PIM DR election can be influenced by explicitly setting the DR priority on the LAN-facing subinterfaces for the routers.

Tech Tip

The HSRP priority and PIM DR priority are shown in the previous table to be the same value; however you are not required to use identical values.

Step 2: This procedure should be repeated for all data or voice subinterfaces.

```
interface [type][number].[sub-interface number]
 encapsulation dot1Q [dot1q VLAN tag]
 ip address [LAN network 1 address] [LAN network 1 netmask]
 ip helper-address 10.4.48.10
 ip pim sparse-mode
 ip pim dr-priority 110
 standby version 2
 standby 1 ip [LAN network 1 gateway address]
 standby 1 priority 110
 standby 1 preempt
 standby 1 authentication md5 key-string c1sco123
```

Example: Layer 2 Link

```
interface GigabitEthernet0/2
 no ip address
 no shutdown

interface GigabitEthernet0/2.64
 description Data
 encapsulation dot1Q 64
 ip address 10.7.18.2 255.255.255.0
 ip helper-address 10.4.48.10
 ip pim dr-priority 110
 ip pim sparse-mode
 standby version 2
 standby 1 ip 10.7.18.1
 standby 1 priority 110
 standby 1 preempt
 standby 1 authentication md5 key-string c1sco123

interface GigabitEthernet0/2.69
 description Voice
 encapsulation dot1Q 69
 ip address 10.7.19.2 255.255.255.0
 ip helper-address 10.4.48.10
 ip pim dr-priority 110
 ip pim sparse-mode
 standby version 2
 standby 1 ip 10.7.19.1
 standby 1 priority 110
 standby 1 preempt
 standby 1 authentication md5 key-string c1sco123
```

Procedure 2 Configure transit network

The transit network is configured between the two routers. This network is used for router-router communication and to avoid hair-pinning. The transit network should use an additional subinterface on the router's physical interface that is already being used for data or voice.

Step 1: Configure the transit network between the two routers.

At a remote site location where there are multiple border routers, the interface throughput delay setting should be set to influence the routing protocol path preference. Set the transit network LAN path to 250000 usec. The delay command is entered in 10 usec units.

There are no end stations connected to this network, so HSRP and DHCP are not required.

```
interface [type][number].[sub-interface number]
 encapsulation dot1Q [dot1q VLAN tag]
 ip address [transit net address] [transit net netmask]
 ip pim sparse-mode
 delay [in 10 usecs]
```

Example

```
interface GigabitEthernet0/2.99
 description Transit Net
 encapsulation dot1Q 99
 ip address 10.7.16.9 255.255.255.252
 ip pim sparse-mode
 delay 25000
```

Step 2: Add transit network VLAN to the access layer switch.

If the VLAN does not already exist on the access layer switch, configure it now.

```
vlan 99
 name Transit-net
```

Step 3: Add transit network VLAN to existing access layer switch trunk.

```
interface GigabitEthernet1/0/48
 switchport trunk allowed vlan add 99
```

Procedure 3 Configure EIGRP on the transit network

A single EIGRP process runs on the DMVPN spoke router, which has already been enabled when configuring the DMVPN tunnel. All interfaces on the router are EIGRP interfaces, but only the DMVPN tunnel interface and transit network are non-passive. The network range must include all interface IP addresses either in a single network statement or in multiple network statements.

EIGRP stub routers normally do not exchange routes with other stub routers, which is problematic in a dual router remote site. It is useful to maintain the benefits of EIGRP stub routing in this type of design. This requires the configuration of stub route leaking between the two remote-site routers for full route reachability. This is needed so that if a router loses its DMVPN link, it knows to how to reach the WAN from the other router at the branch.

In order to prevent the remote site from becoming a transit site during certain failure conditions when stub route leaking is enabled, you must also configure a distribute-list on the tunnel interfaces to control route advertisement.

Step 1: Configure the transit network subinterface as non-passive.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   af-interface Port-channel1.99
    no passive-interface
   exit-af-interface
 exit-address-family
```

Step 2: Configure EIGRP neighbor authentication. Neighbor authentication enables the secure establishment of peering adjacencies and exchange route tables over the transit network interface.

```
key chain LAN-KEY
 key 1
  key-string c1sco123

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Port-channel1.99
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
 exit-address-family
```

Step 3: Configure stub route leaking.

A simple route-map statement with no match statements matches all routes, which permits full route leaking between two routers configured as EIGRP stub.

```
route-map STUB-LEAK-ALL permit 100
 description Leak all routes

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  eigrp stub connected summary redistributed leak-map STUB-LEAK-ALL
 exit-address-family
```

Procedure 4 Enable enhanced object tracking

The HSRP active router remains the active router unless the router is reloaded or fails. Having the HSRP router remain as the active router can lead to undesired behavior. If the primary WAN transport were to fail, the HSRP active router would learn an alternate path through the transit network to the HSRP standby router and begin to forward traffic across the alternate path. This is sub-optimal routing, and you can address it by using EOT.

The HSRP active router can track the state of its DMVPN tunnel interface. If the tunnel line-protocol state changes to down, this implies that the path to the primary site is no longer viable. This is a benefit of using the **if-state nhrp** feature with a DMVPN tunnel configuration.

This procedure is valid only on the router connected to the primary transport.

Step 1: Configure EOT.

A tracked object is created based on tunnel line-protocol state. If the tunnel is up, the tracked object status is Up; if the tunnel is down, the tracked object status is Down. A short delay is added after the tunnel interface comes back up in order to ensure that routing has converged properly before changing the HSRP active router.

```
track 50 interface Tunnel10 line-protocol
 delay up 20
```

Step 2: Link HSRP with the tracked object.

All data or voice subinterfaces should enable HSRP tracking.

HSRP can monitor the tracked object status. If the status is down, the HSRP priority is decremented by the configured priority. If the decrease is large enough, the HSRP standby router preempts.

```
interface [interface type] [number].[sub-interface number]
 standby 1 track 50 decrement 10
```

Example

```
track 50 interface Tunnel10 line-protocol
 delay up 20

interface GigabitEthernet0/2.64
 standby 1 track 50 decrement 10

interface GigabitEthernet0/2.69
 standby 1 track 50 decrement 10
```

PROCESS

Configuring Remote-Site DMVPN Router (Router 2)

1. Configure the WAN remote site router
2. Configure IP multicast routing
3. Configure the WAN-facing VRF
4. Connect to the Internet
5. Configure IKEv2 and IPsec
6. Configure the mGRE tunnel
7. Configure EIGRP
8. Configure IP multicast routing
9. Connect router to access layer switch
10. Configure access layer interfaces
11. Configure access layer HSRP
12. Configure transit network
13. Configure EIGRP on the transit network

Use these procedures when configuring the second router of a dual-router, dual-link design for either the hybrid design model or the dual-Internet design model.

This set of procedures includes the additional steps necessary to configure a second router as a DMVPN spoke router when the first router has already been configured with the process “Configuring Remote-Site DMVPN Spoke Router.”

The previous process, “Router 1 Modifications for Dual Router Design,” must also be completed.

The following flowchart provides details about how to complete the configuration of a remote-site DMVPN spoke router.

Figure 25 Remote-site DMVPN second router configuration flowchart

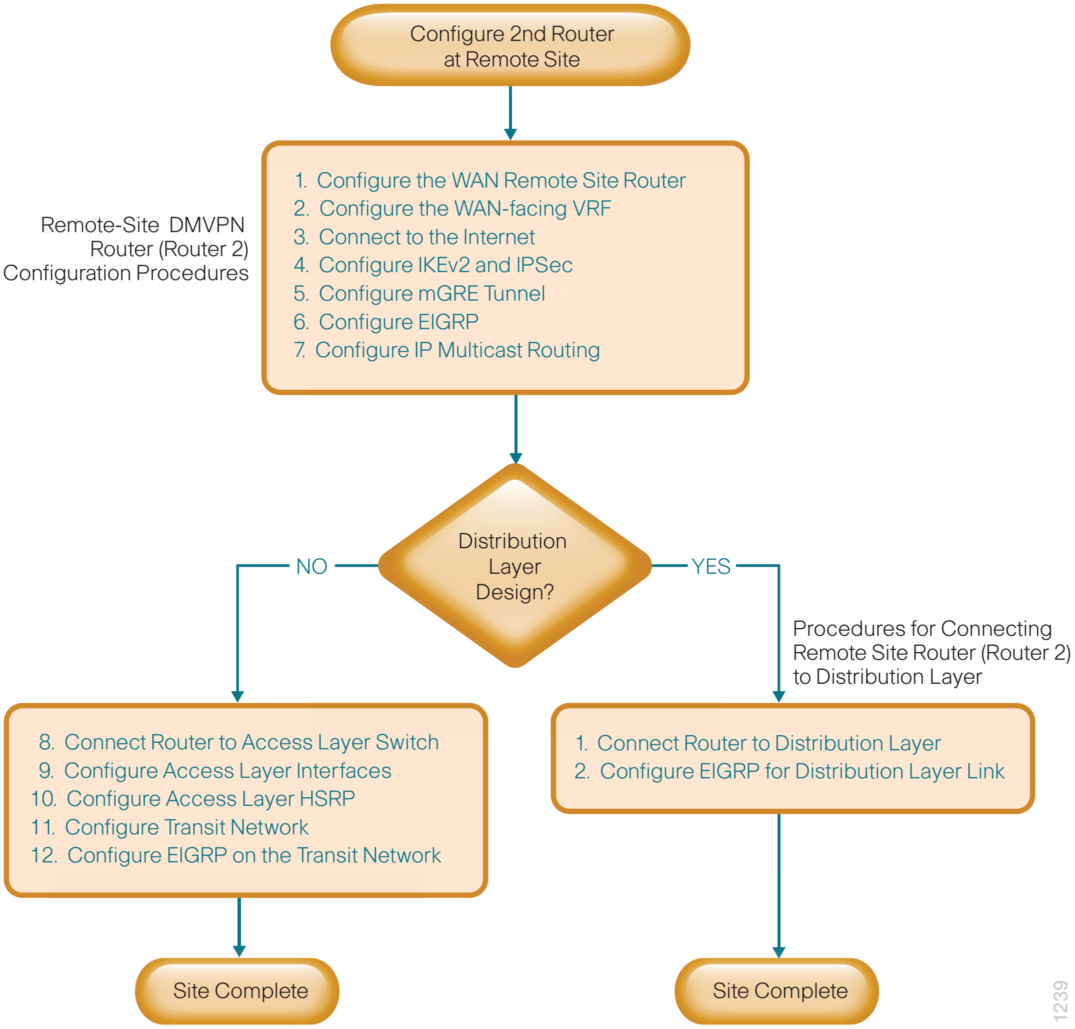
Configure 2nd Router at Remote Site
Remote-Site DMVPN Router (Router 2) Configuration Procedures
1. Configure the WAN Remote Site Router
2. Configure the WAN-facing VRF
3. Connect to the Internet
4. Configure IKEv2 and IPSec
5. Configure mGRE Tunnel
6. Configure EIGRP
7. Configure IP Multicast Routing
Distribution Layer Design?
NO
YES
Procedures for Connecting Remote Site Router (Router 2) to Distribution Layer
8. Connect Router to Access Layer Switch
9. Configure Access Layer Interfaces
10. Configure Access Layer HSRP
11. Configure Transit Network
12. Configure EIGRP on the Transit Network
1. Connect Router to Distribution Layer
2. Configure EIGRP for Distribution Layer Link
Site Complete
Site Complete
1239

Procedure 1 Configure the WAN remote site router

Within this design, there are features and services that are common across all WAN remote-site routers. These are system settings that simplify and secure the management of the solution.

To complete the base configuration for this router, follow the steps in “Configure the platform base features” in Appendix C.

Procedure 2 Configure IP multicast routing

Optional

This optional procedure includes additional steps for configuring IP Multicast on a router. Skip this procedure if you do not want to use IP Multicast in your environment.

In this design, which is based on sparse mode multicast operation, Auto RP is used to provide a simple yet scalable way to provide a highly resilient RP environment.

Step 1: Enable IP Multicast routing on the platforms in the global configuration mode.

```
ip multicast-routing
```

Step 2: Every Layer 3 switch and router must be configured to discover the IP Multicast RP with autorp. Use the **ip pim autorp listener** command to allow for discovery across sparse mode links. This configuration provides for future scaling and control of the IP Multicast environment and can change based on network needs and design.

```
ip pim autorp listener
```

Step 3: All Layer 3 interfaces in the network must be enabled for sparse mode multicast operation.

```
ip pim sparse-mode
```

Procedure 3 Configure the WAN-facing VRF

A WAN-facing VRF is created to support FVRF for DMVPN. The VRF name is arbitrary, but it is useful to select a name that describes the VRF. The VRF must be enabled for IPv4.

Table 49 VRF assignments

IWAN design model	Primary WAN VRF	Secondary WAN VRF
Hybrid	IWAN-TRANSPORT-1	IWAN-TRANSPORT-2
Dual Internet	IWAN-TRANSPORT-3	IWAN-TRANSPORT-4

This design uses VRF Lite, so the selection is only locally significant to the device. It is a best practice to use the same VRF/RD combination across multiple devices when using VRFs in a similar manner. However, this convention is not strictly required.

Step 1: Configure the secondary WAN VRF.

Example: Secondary WAN in the IWAN hybrid design model

```
vrf definition IWAN-TRANSPORT-2
 address-family ipv4
```

Procedure 4 Connect to the Internet

The remote sites using DMVPN can use either static or dynamically assigned IP addresses. Cisco tested the design with a DHCP assigned external address, which also provides a dynamically configured default route.

The DMVPN spoke router connects directly to the Internet without a separate firewall. This connection is secured in two ways. Because the Internet interface is in a separate VRF, no traffic can access the global VRF except traffic sourced through the DMVPN tunnel. This design provides implicit security. Additionally, an IP access list permits only the traffic required for an encrypted tunnel, as well as DHCP and various ICMP protocols for troubleshooting.

Step 1: Enable the interface, select VRF and enable DHCP.

The DMVPN design uses FVRF, so you must place this interface into the VRF configured in the previous procedure.

```
interface GigabitEthernet0/0
 vrf forwarding IWAN-TRANSPORT-2
 ip address dhcp
 no cdp enable
 no shutdown
```

Do not enable PIM on this interface because no multicast traffic should be requested from this interface.

Step 2: Configure and apply the access list.

The IP access list must permit the protocols specified in the following table. The access list is applied inbound on the WAN interface, so filtering is done on traffic destined to the router.

Table 50 Required DMVPN protocols

Name	Protocol	Usage
non500-isakmp	UDP 4500	IPsec via NAT-T
isakmp	UDP 500	ISAKMP
esp	IP 50	IPsec
bootpc	UDP 68	DHCP

Example

```
interface GigabitEthernet0/0
 ip access-group ACL-INET-PUBLIC in

ip access-list extended ACL-INET-PUBLIC
 permit udp any any eq non500-isakmp
 permit udp any any eq isakmp
 permit esp any any
 permit udp any any eq bootpc
```

The additional protocols listed in the following table may assist in troubleshooting, but are not explicitly required to allow DMVPN to function properly.

Table 51   Optional protocols: DMVPN spoke router

Name	Protocol	Usage
icmp echo	ICMP Type 0, Code 0	Allow remote pings
icmp echo-reply	ICMP Type 8, Code 0	Allow ping replies (from your requests)
icmp ttl-exceeded	ICMP Type 11, Code 0	Allow traceroute replies (from your requests)
icmp port-unreachable	ICMP Type 3, Code 3	Allow traceroute replies (from your requests)
UDP high ports	UDP > 1023, TTL=1	Allow remote traceroute

The additional optional entries for an access list to support ping are as follows:

```
 permit icmp any any echo
 permit icmp any any echo-reply
```

The additional optional entries for an access list to support traceroute are as follows:

```
 permit icmp any any ttl-exceeded      ! for traceroute (sourced)
 permit icmp any any port-unreachable  ! for traceroute (sourced)
 permit udp any any gt 1023 ttl eq 1   ! for traceroute (destination)
```

Procedure 5 Configure IKEv2 and IPsec

This procedure uses the parameters in the tables below. Choose the table that represents the design model that you are configuring. This procedure applies to the Secondary WAN.

Table 52 Crypto parameters: IWAN hybrid design model

Parameter	Primary WAN	Secondary WAN
vrf	IWAN-TRANSPORT-1	IWAN-TRANSPORT-2
crypto ikev2 keyring	DMVPN-KEYRING-1	DMVPN-KEYRING-2
crypto ikev2 profile	FVRF-IKEv2-IWAN-TRANSPORT-1	FVRF-IKEv2-IWAN-TRANSPORT-2
crypto ipsec profile	DMVPN-PROFILE-TRANSPORT-1	DMVPN-PROFILE-TRANSPORT-2

Table 53 Crypto parameters: IWAN dual Internet design model

Parameter	Primary WAN	Secondary WAN
vrf	IWAN-TRANSPORT-3	IWAN-TRANSPORT-4
crypto ikev2 keyring	DMVPN-KEYRING-3	DMVPN-KEYRING-4
crypto ikev2 profile	FVRF-IKEv2-IWAN-TRANSPORT-3	FVRF-IKEv2-IWAN-TRANSPORT-4
crypto ipsec profile	DMVPN-PROFILE-TRANSPORT-3	DMVPN-PROFILE-TRANSPORT-4

IPsec uses a key exchange between the routers in order to encrypt/decrypt the traffic. These keys can be exchanged using pre-shared keys or PKI certificates with a certificate authority. It is also possible to use a combination of the two, which is useful during a migration from one method to the other. Choose one of two options below as your method of key exchange.

Option 1: Configure with Pre-Shared Keys

Step 1: Configure the crypto keyring for pre-shared keys.

The crypto keyring defines a pre-shared key (or password) valid for IP sources that are reachable within a particular VRF. This key is a wildcard pre-shared key if it applies to any IP source. A wildcard key is configured using the 0.0.0.0 0.0.0.0 network/mask combination.

```
crypto ikev2 keyring [keyring name]
 peer ANY
  address 0.0.0.0 0.0.0.0
  pre-shared-key [password]
```

Example: Secondary WAN in the IWAN hybrid design model

```
crypto ikev2 keyring DMVPN-KEYRING-2
 peer ANY
  address 0.0.0.0 0.0.0.0
  pre-shared-key c1sco123
```

Step 2: Configure the IKEv2 proposal.

The user-defined IKEv2 proposal includes only the following:

- Encryption with AES cipher with a 256-bit key
- Integrity with SHA and a 512-bit digest
- Pseudo-random function with SHA and a 512-bit digest
- Diffie-Hellman group: 14

```
crypto ikev2 proposal [proposal name]
 encryption [encryption algorithm]
 integrity [integrity hash algorithm]
 group [Diffie-Hellman group]
```

Example

```
crypto ikev2 proposal AES/CBC/256
 encryption aes-cbc-256
 integrity sha512
 group 14
```

The default IKEv2 proposal is also used.

A **show crypto ikev2 proposal** displays the details of the two proposals.

```
RS12-2911-2# show crypto ikev2 proposal
IKEv2 proposal: AES/CBC/256
     Encryption : AES-CBC-256
     Integrity  : SHA512
     PRF        : SHA512
     DH Group   : DH_GROUP_2048_MODP/Group 14
IKEv2 proposal: default
    Encryption: AES-CBC-256 AES-CBC-192 AES-CBC-128
    Integrity : SHA512 SHA384 SHA256 SHA96 MD596
    PRF       : SHA512 SHA384 SHA256 SHA1 MD5
    DH Group  : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2
```

Step 3: Configure the IKEv2 profile.

The IKEv2 profile creates an association between an identity address, a VRF, and a crypto keyring. A wildcard address within a VRF is referenced with 0.0.0.0. The profile also defines what method of key sharing will be used on this router with **authentication local** and what methods will be accepted from remote locations with **authentication remote**. The **pre-share** keyword is used with the keyring defined above.

DPD is essential to facilitate fast reconvergence and for spoke registration to function properly in case a DMVPN hub is restarted. The IWAN design recommends you set the remote site DPD timer to 40 seconds with a 5 second retry.

```
crypto ikev2 profile [profile name]
 match fvrf [vrf name]
 match identity remote address [IP address]
 authentication remote pre-share
 authentication local pre-share
 keyring local [keyring name]
 dpd [interval in seconds] [retry interval] on-demand
```

Example: Secondary WAN in the IWAN hybrid design model

```
crypto ikev2 profile FVRF-IKEv2-IWAN-TRANSPORT-2
 match fvrf IWAN-TRANSPORT-2
 match identity remote address 0.0.0.0
 authentication remote pre-share
 authentication local pre-share
 keyring local DMVPN-KEYRING-2
 dpd 40 5 on-demand
```

Step 4: Define the IPsec transform set.

A transform set is an acceptable combination of security protocols, algorithms, and other settings to apply to IPsec-protected traffic. Peers agree to use a particular transform set when protecting a particular data flow.

The IPsec transform set for DMVPN uses the following:

- ESP with the 256-bit AES encryption algorithm
- ESP with the SHA (HMAC variant) authentication algorithm

Since the DMVPN hub router is behind a NAT device, the IPsec transform must be configured for transport mode.

```
crypto ipsec transform-set [transform set] esp-aes 256 esp-sha-hmac
 mode transport
```

Example

```
crypto ipsec transform-set AES256/SHA/TRANSPORT esp-aes 256 esp-sha-hmac
 mode transport
```

Step 5: Configure the IPsec profile.

The IPsec profile creates an association between an IKEv2 profile and an IPsec transform-set.

```
crypto ipsec profile [profile name]
 set transform-set [transform set]
 set ikev2-profile [ikev2 profile name]
```

Example: Secondary WAN in the IWAN hybrid design model

```
crypto ipsec profile DMVPN-PROFILE-TRANSPORT-2
 set transform-set AES256/SHA/TRANSPORT
 set ikev2-profile FVRF-IKEv2-IWAN-TRANSPORT-2
```

Step 6: Increase the IPsec anti-replay window size.

```
crypto ipsec security-association replay window-size 1024
```

Tech Tip

QoS queuing delays can cause anti-replay packet drops, so it is important to extend the window size to prevent the drops from occurring.

Increasing the anti-replay window size has no impact on throughput and security. The impact on memory is insignificant because only an extra 128 bytes per incoming IPsec SA are needed.

It is recommended that you use the 1024 window size in order to minimize future anti-replay problems.

If you do not increase the window size, the router may drop packets and you may see the following error message on the router CLI and in the log:

```
%CRYPTO-4-PKT_REPLAY_ERR: decrypt: replay check failed
```

Step 7: Skip to the next procedure.

Option 2: Configure with a Certificate Authority

If you want to use a certificate authority, you will have to configure a pre-shared key on one of the hub border routers in order to allow each remote site to establish a DMVPN tunnel to the WAN aggregation site. After the first DMVPN tunnel at a remote site is established, the router will be able to authenticate to the CA and obtain a certificate. After obtaining the certificate, you can configure the remote site to use PKI.

The **crypto pki trustpoint** is the method of specifying the parameters associated with a CA. The router must authenticate to the CA first and then enroll with the CA in order to obtain its own identity certificate.

Step 1: The fingerprint command limits the responding CA. You can find this fingerprint by using **show crypto pki server** on the IOS CA.

```
IWAN-IOS-CA# show crypto pki server
Certificate Server IWAN-IOS-CA:
    Status: enabled
    State: enabled
    Server’s configuration is locked  (enter "shut" to unlock it)
    Issuer name: CN=IWAN-IOS-CA.cisco.local L=SanJose St=CA C=US
    CA cert fingerprint: 75BEF625 9A9876CF 6F341FE5 86D4A5D8
    Granting mode is: auto
    Last certificate issued serial number (hex): 11
    CA certificate expiration timer: 10:28:57 PST Nov 11 2017
    CRL NextUpdate timer: 09:47:47 PST Dec 4 2014
    Current primary storage dir: nvram:
    Current storage dir for .crl files: nvram:
    Database Level: Complete - all issued certs written as <serialnum>.cer
```

Step 2: Configure the PKI trust point.

```
crypto pki trustpoint [name]
 enrollment url [URL of IOS CA]
 serial-number none
 fqdn [fully qualified domain name of this router]
 ip-address [loopback IP address of this router]
 fingerprint [fingerprint from IOS CA]
 revocation-check none
 rsakeypair [name] 2048 2048
```

Example: RS14-2921-2

This example is from the secondary WAN remote site router in the dual Internet design model. After the DMVPN tunnel is established with pre-shared keys, it can reach the IOS CA through the internal network at 10.6.24.11 using the default VRF.

```
crypto pki trustpoint IWAN-CA
 enrollment url http://10.6.24.11:80
 serial-number none
 fqdn RS14-2921-2.cisco.local
 ip-address 10.255.244.14
 fingerprint 75BEF6259A9876CF6F341FE586D4A5D8
 revocation-check none
 rsakeypair IWAN-CA-KEYS 2048 2048
```

Step 3: Authenticate to the CA and obtain the CA’s certificate

Exit the trustpoint configuration mode on the hub router and issue the following command to authenticate to the CA and get its certificate.

```
RS14-2921-2(config)# crypto pki authenticate IWAN-CA
Certificate has the following attributes:
       Fingerprint MD5: 1A070F43 38068E1C BE04A8FB CBAA406F
      Fingerprint SHA1: F463AAF1 54C4D994 E2732F36 3BEBED23 D410192E
Trustpoint Fingerprint: 1A070F43 38068E1C BE04A8FB CBAA406F
Certificate validated - fingerprints matched.
Trustpoint CA certificate accepted.
```

Step 4: Enroll with the CA, enter a password for key retrieval and obtain a certificate for this hub router.

```
RS14-2921-2(config)# crypto pki enroll IWAN-CA

% Start certificate enrollment ..
% Create a challenge password. You will need to verbally provide this
   password to the CA Administrator in order to revoke your certificate.
   For security reasons your password will not be saved in the configuration.
   Please make a note of it.

Password: c1sco123
Re-enter password: c1sco123

% The subject name in the certificate will include: RS14-2921-2.cisco.local
% The IP address in the certificate is 10.255.244.14

Request certificate from CA? [yes/no]: yes
% Certificate request sent to Certificate Authority
% The ‘show crypto pki certificate verbose IWAN-CA’ command will show the finger-
print.
```

Step 5: Configure the IKEv2 proposal.

The user-defined IKEv2 proposal includes only the following:

- Encryption with AES cipher with a 256-bit key
- Integrity with SHA and a 512-bit digest
- Pseudo-random function with SHA and a 512-bit digest
- Diffie-Hellman group: 14

```
crypto ikev2 proposal [proposal name]
 encryption [encryption algorithm]
 integrity [integrity hash algorithm]
 group [Diffie-Hellman group]
```

Example

```
crypto ikev2 proposal AES/CBC/256
 encryption aes-cbc-256
 integrity sha512
 group 14
```

The default IKEv2 proposal is also used.

A **show crypto ikev2 proposal** displays the details of the two proposals.

```
RS14-2921-2# show crypto ikev2 proposal
IKEv2 proposal: AES/CBC/256
     Encryption : AES-CBC-256
     Integrity  : SHA512
     PRF        : SHA512
     DH Group   : DH_GROUP_2048_MODP/Group 14
IKEv2 proposal: default
    Encryption: AES-CBC-256 AES-CBC-192 AES-CBC-128
    Integrity : SHA512 SHA384 SHA256 SHA96 MD596
    PRF       : SHA512 SHA384 SHA256 SHA1 MD5
    DH Group  : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2
```

Step 6: Configure the IKEv2 profile.

The IKEv2 profile creates an association between an identity address, a VRF, and a crypto keyring. A wildcard address within a VRF is referenced with 0.0.0.0. The **identity local address** must match the crypto pki trustpoint’s **ip-address** value from the step above in order to avoid CRYPTO-6-IKMP_NO_ID_CERT_ADDR_MATCH warning messages on the router CLI and in the log.

The profile also defines what method of key sharing will be used on this router with **authentication local** and what methods will be accepted from remote locations with **authentication remote**. The **rsa-sig** keyword is used when certificates contain the encryption key.

DPD is essential to facilitate fast reconvergence and for spoke registration to function properly in case a DMVPN hub is restarted. The IWAN design recommends you set the remote site DPD timer to 40 seconds with a 5 second retry.

```
crypto ikev2 profile [profile name]
 match fvrf [vrf name]
 match identity remote address [IP address]
 identity local address [Loopback IP address of this router]
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint [trustpoint name]
 dpd [interval in seconds] [retry interval] on-demand
```

Example: RS14-2921-2

```
crypto ikev2 profile FVRF-IKEv2-IWAN-TRANSPORT-4
 match fvrf IWAN-TRANSPORT-4
 match identity remote address 0.0.0.0
 identity local address 10.255.244.14
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint IWAN-CA
 dpd 40 5 on-demand
```

Step 7: Define the IPsec transform set.

A transform set is an acceptable combination of security protocols, algorithms, and other settings to apply to IPsec-protected traffic. Peers agree to use a particular transform set when protecting a particular data flow.

The IPsec transform set for DMVPN uses the following:

- ESP with the 256-bit AES encryption algorithm
- ESP with the SHA (HMAC variant) authentication algorithm

Because the DMVPN hub router is behind a NAT device, the IPsec transform must be configured for transport mode.

```
crypto ipsec transform-set [transform set] esp-aes 256 esp-sha-hmac
 mode transport
```

Example

```
crypto ipsec transform-set AES256/SHA/TRANSPORT esp-aes 256 esp-sha-hmac
 mode transport
```

Step 8: Configure the IPsec profile.

The IPsec profile creates an association between an IKEv2 profile and an IPsec transform-set.

```
crypto ipsec profile [profile name]
 set transform-set [transform set]
 set ikev2-profile [ikev2 profile name]
```

Example: Secondary WAN in the dual Internet design model

```
crypto ipsec profile DMVPN-PROFILE-TRANSPORT-4
 set transform-set AES256/SHA/TRANSPORT
 set ikev2-profile FVRF-IKEv2-IWAN-TRANSPORT-4
```

Step 9: Increase the IPsec anti-replay window size.

```
crypto ipsec security-association replay window-size [value]
```

Example

```
crypto ipsec security-association replay window-size 1024
```

Tech Tip

QoS queuing delays can cause anti-replay packet drops, so it is important to extend the window size to prevent the drops from occurring.

Increasing the anti-replay window size has no impact on throughput and security. The impact on memory is insignificant because only an extra 128 bytes per incoming IPsec SA is needed.

It is recommended that you use the maximum window size in order to minimize future anti-replay problems. On the Cisco ASR 1K, ISR 4K and ISR G2 router platforms, the maximum replay window size is 1024.

If you do not increase the window size, the router may drop packets and you may see the following error message on the router CLI and in the log:

```
%CRYPTO-4-PKT_REPLAY_ERR: decrypt: replay check failed
```

Procedure 6 Configure the mGRE tunnel

This procedure uses the parameters in the table below. Choose the rows that represent the design model that you are configuring. This procedure applies to the Secondary WAN.

Table 54 DMVPN tunnel parameters

IWAN design model	Tunnel VRF	Tunnel number	Tunnel network	NHRP network ID/ tunnel key
Hybrid—Primary WAN	IWAN-TRANSPORT-1	10	10.6.34.0/23	101
Hybrid—Secondary WAN	IWAN-TRANSPORT-2	11	10.6.36.0/23	102
Dual Internet—Primary WAN	IWAN-TRANSPORT-3	20	10.6.38.0/23	201
Dual Internet—Secondary WAN	IWAN-TRANSPORT-4	21	10.6.40.0/23	202

Step 1: Configure basic interface settings.

The tunnel number is arbitrary, but it is best to begin tunnel numbering at 10 or above, because other features deployed in this design may also require tunnels and they may select lower numbers by default.

You must set the bandwidth setting to match the bandwidth of the respective transport, which corresponds to the actual interface speed. Or, if you are using a subrate service, use the policed rate from the carrier. QoS and PfR require the correct bandwidth setting to operate properly.

Configure the **ip mtu** to 1400 and the **ip tcp adjust-mss** to 1360. There is a 40 byte difference, which corresponds to the combined IP and TCP header length.

The tunnel interface throughput delay setting should be set to influence the routing protocol path preference. Set the primary WAN path to 10000 usec and the secondary WAN path to 200000 usec to prefer one over the other. The delay command is entered in 10 usec units.

```
interface Tunnel11
 bandwidth 10000
 ip address 10.6.36.12 255.255.254.0
 no ip redirects
 ip mtu 1400
 ip tcp adjust-mss 1360
 delay 20000
```

Step 2: Configure the tunnel.

DMVPN uses mGRE tunnels. This type of tunnel requires a source interface only. The source interface should be the same interface you use to connect to the Internet. You should set the **tunnel vrf** command to the VRF defined previously for FVRF.

To enable encryption on this interface, you must apply the IPsec profile that you configured in the previous procedure.

```
interface Tunnel11
 tunnel source GigabitEthernet0/0
 tunnel mode gre multipoint
 tunnel key 102
 tunnel vrf IWAN-TRANSPORT-2
 tunnel protection ipsec profile DMVPN-PROFILE-TRANSPORT-2
```

Step 3: Configure NHRP.

The spoke router requires an additional configuration statement to define the NHRP server. This statement includes the NBMA definition for the DMVPN hub router tunnel endpoint. EIGRP relies on a multicast transport. Spoke routers require the NHRP multicast keyword in this statement.

When hub BRs are added for horizontal scaling or a second data center is added as a transit site, spoke routers require additional NHS statements for each BR in their environment. The configuration details are covered in subsequent sections of this guide.

The value used for the NHS is the mGRE tunnel address for the DMVPN hub router. The NBMA entry must be set to either the MPLS DMVPN hub router's actual public address or the outside NAT value of the DMVPN hub, as configured on the Cisco ASA 5500. This design uses the values shown in the following tables.

Table 55 DMVPN tunnel NHRP parameters: IWAN hybrid design model

	Transport 1	Transport 2
VRF	IWAN-TRANSPORT-1	IWAN-TRANSPORT-2
DMVPN hub public address (actual)	192.168.6.1	192.168.146.10
DMVPN hub public address (externally routable after NAT)	n/a (MPLS)	172.16.140.1
DMVPN hub tunnel IP address (NHS)	10.6.34.1	10.6.36.1
Tunnel number	10	11
NHRP network ID	101	102

Table 56 DMVPN tunnel NHRP parameters: IWAN dual Internet design model

	Transport 3	Transport 4
VRF	IWAN-TRANSPORT-3	IWAN-TRANSPORT-4
DMVPN hub public address (actual)	192.168.146.20	192.168.146.21
DMVPN hub public address (externally routable after NAT)	172.16.140.11	172.17.140.11
DMVPN tub tunnel IP address (NHS)	10.6.38.1	10.6.40.1
Tunnel number	20	21
NHRP network ID	201	202

NHRP requires all devices within a DMVPN cloud to use the same network ID and authentication key. The NHRP cache holdtime should be configured to 600 seconds.

This design supports DMVPN spoke routers that receive their external IP addresses through DHCP. It is possible for these routers to acquire different IP addresses after a reload. When the router attempts to register with the NHRP server, it may appear as a duplicate to an entry already in the cache and be rejected. The **registration no-unique** option allows you to overwrite existing cache entries. This feature is only required on NHRP clients (DMVPN spoke routers). The **if-state nhrp** option ties the tunnel line-protocol state to the reachability of the NHRP NHS, and if the NHS is unreachable the tunnel line-protocol state changes to down. This feature is used in conjunction with EOT.

```
interface Tunnel11
 ip nhrp authentication cisco123
 ip nhrp network-id 102
 ip nhrp holdtime 600
 ip nhrp nhs 10.6.36.1 nbma 172.16.140.1 multicast
 ip nhrp registration no-unique
 ip nhrp shortcut
 if-state nhrp
```

It is not necessary to disable nhrp route-watch on the second router of a dual router remote site location because there is only one WAN path in the RIB.

Procedure 7 Configure EIGRP

A single EIGRP process runs on the DMVPN spoke router. All interfaces on the router are EIGRP interfaces, but only the DMVPN tunnel interface is non-passive. The network range must include all interface IP addresses either in a single network statement or in multiple network statements. This design uses a best practice of assigning the router ID to a loopback address. All DMVPN spoke routers should run EIGRP stub routing to improve network stability and reduce resource utilization.

Step 1: Configure an EIGRP process for DMVPN by using EIGRP named mode on the spoke router.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface default
   passive-interface
  exit-af-interface
  af-interface Tunnel11
   no passive-interface
  exit-af-interface
network 10.6.36.0 0.0.1.255
network 10.7.0.0 0.0.255.255
network 10.255.0.0 0.0.255.255
eigrp router-id [IP address of Loopback0]
eigrp stub connected summary redistributed
exit-address-family
```

Step 2: Configure EIGRP values for the mGRE tunnel interface.

The EIGRP hello interval is increased to 20 seconds and the EIGRP hold time is increased to 60 seconds in order to accommodate up to 2000 remote sites on a single DMVPN cloud. Increasing the EIGRP timers also slows down the routing convergence to improve network stability and the IWAN design allows PfR to initiate the fast failover, so changing the timers is recommended for all IWAN deployments.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel11
   hello-interval 20
   hold-time 60
  exit-af-interface
exit-address-family
```

Step 3: Configure EIGRP neighbor authentication. Neighbor authentication enables the secure establishment of peering adjacencies and exchange route tables over the DMVPN tunnel interface.

```
key chain WAN-KEY
 key 1
  key-string c1sco123

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel11
   authentication mode md5
   authentication key-chain WAN-KEY
  exit-af-interface
 exit-address-family
```

Step 4: Configure EIGRP network summarization.

The remote-site LAN networks must be advertised. The IP assignment for the remote sites was designed so that all of the networks in use can be summarized within a single aggregate route. As configured below, the **summary-address** command suppresses the more specific routes. If any network within the summary is present in the route table, the summary is advertised to the DMVPN hub, which offers a measure of resiliency. If the various LAN networks cannot be summarized, then EIGRP continues to advertise the specific routes.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel10
   summary-address [summary network] [summary mask]
  exit-af-interface
 exit-address-family
```

Step 5: Tag and filter the routes.

This design uses a single EIGRP autonomous system for the hub sites, WAN, and all of the WAN remote sites. Every remote site is dual-connected for resiliency. However, due to the multiple paths that exist within this topology, you must try to avoid routing loops and to prevent remote sites from becoming transit sites if WAN failures were to occur.

In this design, there are different IP subnets for each DMVPN network, and the EIGRP tags are defined for each segment in order to help with readability and troubleshooting. When a design uses more than one data center, additional tags are required in order to identify the different DMVPN hub router locations.

The following logic is used to control the routing.

- Each DMVPN network will have an EIGRP route tag to prevent routes from being re-advertised over the other DMVPN networks.
- All prefixes that are advertised towards the WAN are uniquely tagged.
- All DMVPN learned WAN prefixes, except those that originate locally from a hub, are advertised towards the LAN and tagged.
- The design always uses DMVPN hub routers in pairs. Each DMVPN hub router blocks DMVPN WAN routes from the LAN that are tagged with the opposite hub's tags.
- Remote sites can learn routes from other remote sites when NHRP uses shortcut routing to bypass the hub locations. The spoke-to-spoke routes are blocked using tag filtering to prevent issues with PfR.

Outbound distribute-lists are used to set tags towards the WAN and LAN. The remote-site routers use the tags set towards the WAN in order to protect against becoming transit sites.

The remote-site routers use an inbound distribute-list in order to limit which routes are accepted for installation into the route table. These routers are configured to only accept routes that do not originate from the WAN sources. To accomplish this task, the remote site router must explicitly tag the DMVPN learned WAN routes.

The following tables show specific route tags in use.

Table 57 Route tag information for IWAN hybrid remote site routers

Tunnel: transport	Tunnel key	Match inbound	Block outbound
Tunnel10: DMVPN1	101 (MPLS)	101 (All routes)	101 (WAN routes)
Tunnel11: DMVPN2	102 (INET)	102 (All routes)	102 (WAN routes)

Table 58 Route tag information for IWAN dual Internet remote site routers

Tunnel: transport	Tunnel key	Match inbound	Block outbound
Tunnel20 DMVPN3	201 (INET1)	201 (All routes)	201 (WAN routes)
Tunnel21: DMVPN4	202 (INET2)	202 (All routes)	202 (WAN routes)

The following example shows a dual router design in the IWAN hybrid design model.

Example: Second Router of a Dual Router Site with DMVPN2

```
route-map DMVPN2-BR-IN Permit 10
 description Match tagged routes inbound
 match tag 102

route-map BLOCK-LEARNED deny 10
 description Block learned routes outbound
 Match tag 101 102

route-map BLOCK-LEARNED permit 20
 description Advertise all other routes outbound

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  topology base
   distribute-list route-map DMVPN2-BR-IN in Tunnel11
   distribute-list route-map BLOCK-LEARNED out Tunnel11
```

Procedure 8 Configure IP multicast routing

Optional

This optional procedure includes additional steps for configuring IP Multicast for a DMVPN tunnel on a router with IP Multicast already enabled. Skip this procedure if you do not want to use IP Multicast in your environment.

Step 1: Configure PIM on the DMVPN tunnel interface.

Enable IP PIM sparse mode on the DMVPN tunnel interface.

```
interface Tunnel11
 ip pim sparse-mode
```

Step 2: Enable PIM NBMA mode for the DMVPN tunnel.

Spoke-to-spoke DMVPN networks present a unique challenge because the spokes cannot directly exchange information with one another, even though they are on the same logical network. This inability to directly exchange information can also cause problems when running IP Multicast.

To resolve the NBMA issue, you need to implement a method where each remote PIM neighbor has its join messages tracked separately. A router in PIM NBMA mode treats each remote PIM neighbor as if it were connected to the router through a point-to-point link.

```
interface Tunnel11
 ip pim nbma-mode
```

Step 3: Configure the DR priority for the DMVPN spoke router.

Proper multicast operation across a DMVPN cloud requires that the hub router assumes the role of PIM DR. Spoke routers should never become the DR. You can prevent that by setting the DR priority to 0 for the spokes.

```
interface Tunnel11
 ip pim dr-priority 0
```

Procedure 9 Connect router to access layer switch

Optional

If you are using a remote-site distribution layer, skip to the “Connecting Remote-Site Router to Distribution Layer” process.

Reader Tip

This guide includes only the additional steps needed to complete the access layer configuration. For complete access layer configuration details, refer to the Campus LAN Layer 2 Access with Simplified Distribution Deployment Guide .

Layer 2 EtherChannels are used to interconnect the router to the access layer in the most resilient method possible, unless the access layer device is a single fixed configuration switch, Otherwise a simple Layer 2 trunk between the router and switch is used.

In the access layer design, the remote sites use collapsed routing, with 802.1Q trunk interfaces to the LAN access layer. The VLAN numbering is locally significant only.

Option 1: Layer 2 EtherChannel from router to access layer switch

Step 1: Configure port-channel interface on the router.

```
interface Port-channel2
 description EtherChannel link to RS12-A2960X
 no shutdown
```

Step 2: Configure EtherChannel member interfaces on the router.

Configure the physical interfaces to tie to the logical port-channel using the channel-group command. The number for the port-channel and channel-group must match. Not all router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet0/1
 description RS12-A2960X Gig1/0/48

interface GigabitEthernet0/2
 description RS12-A2960X Gig2/0/48

interface range GigabitEthernet0/1, GigabitEthernet0/2
  no ip address
  channel-group 2
  no shutdown
```

Step 3: Configure EtherChannel member interfaces on the access layer switch.

Connect the router EtherChannel uplinks to separate switches in the access layer switch stack.

The physical interfaces that are members of a Layer 2 EtherChannel are configured prior to configuring the logical port-channel interface. Doing the configuration in this order allows for minimal configuration and reduces errors because most of the commands entered to a port-channel interface are copied to its members interfaces and do not require manual replication.

Configure two or more physical interfaces to be members of the EtherChannel. It is recommended that they are added in multiples of two. Also, apply the egress QoS macro that was defined in the platform configuration procedure to ensure traffic is prioritized appropriately.

Not all connected router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet1/0/48
 description Link to RS12-2911-2 Gig0/1

interface GigabitEthernet2/0/48
 description Link to RS12-2911-2 Gig0/2

interface range GigabitEthernet1/0/48, GigabitEthernet2/0/48
 switchport
 channel-group 2 mode on
 logging event link-status
 logging event trunk-status
 logging event bundle-status
 load-interval 30
 macro apply EgressQoS
```

Step 4: Configure EtherChannel trunk on the access layer switch.

An 802.1Q trunk is used which allows the router to provide the Layer 3 services to all the VLANs defined on the access layer switch. The VLANs allowed on the trunk are pruned to only the VLANs that are active on the access layer switch. When using EtherChannel the interface type will be port-channel and the number must match the channel group configured in Step 3. DHCP Snooping and ARP inspection are set to trust.

```
interface Port-channel2
 description EtherChannel link to RS12-2911-2
 switchport trunk allowed vlan 64,69,99
 switchport mode trunk
 ip arp inspection trust
 spanning-tree portfast trunk
 ip dhcp snooping trust
 load-interval 30
 no shutdown
```

The Cisco Catalyst 3750 Series Switch requires the **switchport trunk encapsulation dot1q** command.

Option 2: Layer 2 trunk from router to access layer switch

Step 1: Enable the physical interface on the router.

```
interface GigabitEthernet0/2
 description RS12-A2960X Gig1/0/48
 no ip address
 no shutdown
```

Step 2: Configure the trunk on the access layer switch.

Use an 802.1Q trunk for the connection, which allows the router to provide the Layer 3 services to all the VLANs defined on the access layer switch. The VLANs allowed on the trunk are pruned to only the VLANs that are active on the access switch. DHCP Snooping and ARP inspection are set to trust.

```
interface  GigabitEthernet1/0/48
 description Link to RS12-2911-2 Gig0/2
 switchport trunk allowed vlan 64,69,99
 switchport mode trunk
 ip arp inspection trust
 spanning-tree portfast trunk
 logging event link-status
 logging event trunk-status
 ip dhcp snooping trust
 no shutdown
 load-interval 30
 macro apply EgressQoS
```

The Cisco Catalyst 3750 Series Switch requires the **switchport trunk encapsulation dot1q** command.

Procedure 10 Configure access layer interfaces

Step 1: Create subinterfaces and assign VLAN tags.

After the physical interface or port-channel have been enabled, then the appropriate data or voice subinterfaces can be mapped to the VLANs on the LAN switch. The subinterface number does not need to equate to the 802.1Q tag, but making them the same simplifies the overall configuration. The subinterface portion of the configuration should be repeated for all data or voice VLANs.

```
interface [type][number].[sub-interface number]
 encapsulation dot1Q [dot1q VLAN tag]
```

Step 2: Configure IP settings for each subinterface.

This design uses an IP addressing convention with the default gateway router assigned an IP address and IP mask combination of **N.N.N.1 255.255.255.0** where N.N.N is the IP network and 1 is the IP host.

When you are using a centralized DHCP server, your routers with LAN interfaces connected to a LAN using DHCP for end-station IP addressing must use an IP helper.

This remote-site DMVPN spoke router is the second router of a dual-router design and HSRP is configured at the access layer. The actual interface IP assignments will be configured in the following procedure.

```
interface  [type][number].[sub-interface number]
 description [usage]
 encapsulation dot1Q [dot1q VLAN tag]
 ip helper-address 10.4.48.10
 ip pim sparse-mode
```

Example: Layer 2 EtherChannel

```
interface Port-channel2
 no ip address
 no shutdown

interface Port-channel2.64
 description Data
 encapsulation dot1Q 64
 ip helper-address 10.4.48.10
 ip pim sparse-mode

interface Port-channel2.69
 description Voice
 encapsulation dot1Q 69
 ip helper-address 10.4.48.10
 ip pim sparse-mode
```

Example: Layer 2 Link

```
interface GigabitEthernet0/2
 no ip address
 no shutdown

interface GigabitEthernet0/2.64
 description Data
 encapsulation dot1Q 64
 ip helper-address 10.4.48.10
 ip pim sparse-mode

interface GigabitEthernet0/2.69
 description Voice
 encapsulation dot1Q 69
 ip helper-address 10.4.48.10
 ip pim sparse-mode
```

Procedure 11 Configure access layer HSRP

Step 1: You configure HSRP to enable a VIP that you use as a default gateway that is shared between two routers. The HSRP active router is the router connected to the primary carrier and the HSRP standby router is the router connected to the secondary carrier or backup link.

Step 2: Configure the HSRP standby router with a standby priority that is lower than the HSRP active router.

The router with the higher standby priority value is elected as the HSRP active router. The preempt option allows a router with a higher priority to become the HSRP active, without waiting for a scenario where there is no router in the HSRP active state. The relevant HSRP parameters for the router configuration are shown in the following table.

Table 59 WAN remote-site HSRP parameters (dual router)

Router	HSRP role	VIP	Real IP address	HSRP priority	PIM DR priority
Primary	Active	.1	.2	110	110
Secondary	Standby	.1	.3	105	105

The dual-router access-layer design requires a modification for resilient multicast. The PIM DR should be on the HSRP active router. The DR is normally elected based on the highest IP address and has no awareness of the HSRP configuration. In this design, the HSRP active router has a lower real IP address than the HSRP standby router, which requires a modification to the PIM configuration. The PIM DR election can be influenced by explicitly setting the DR priority on the LAN-facing subinterfaces for the routers.

Tech Tip

The HSRP priority and PIM DR priority are shown in the previous table to be the same value; however there is no requirement that these values must be identical.

Step 3: Repeat this procedure for all data or voice subinterfaces.

```
interface [interface type][number].[sub-interface number]
 ip address [LAN network 1 address] [LAN network 1 netmask]
 ip pim dr-priority 105
 standby version 2
 standby 1 ip [LAN network 1 gateway address]
 standby 1 priority 105
 standby 1 preempt
 standby 1 authentication md5 key-string c1sco123
```

Example: Layer 2 EtherChannel

```
interface PortChannel2
 no ip address
 no shutdown

interface PortChannel2.64
 description Data
 encapsulation dot1Q 64
 ip address 10.7.18.3 255.255.255.0
 ip helper-address 10.4.48.10
 ip pim dr-priority 105
 ip pim sparse-mode
 standby version 2
 standby 1 ip 10.7.18.1
 standby 1 priority 105
 standby 1 preempt
 standby 1 authentication md5 key-string c1sco123

interface PortChannel2.69
 description Voice
 encapsulation dot1Q 69
 ip address 10.7.19.3 255.255.255.0
 ip helper-address 10.4.48.10
 ip pim dr-priority 105
 ip pim sparse-mode
 standby version 2
 standby 1 ip 10.7.19.1
 standby 1 priority 105
 standby 1 preempt
 standby 1 authentication md5 key-string c1sco123
```

Example: Layer 2 Link

```
interface GigabitEthernet0/2
 no ip address
 no shutdown

interface GigabitEthernet0/2.64
 description Data
 encapsulation dot1Q 64
 ip address 10.7.18.3 255.255.255.0
 ip helper-address 10.4.48.10
 ip pim dr-priority 105
 ip pim sparse-mode
 standby version 2
 standby 1 ip 10.7.18.1
 standby 1 priority 105
 standby 1 preempt
 standby 1 authentication md5 key-string c1sco123

interface GigabitEthernet0/2.69
 description Voice
 encapsulation dot1Q 69
 ip address 10.7.19.3 255.255.255.0
 ip helper-address 10.4.48.10
 ip pim dr-priority 105
 ip pim sparse-mode
 standby version 2
 standby 1 ip 10.7.19.1
 standby 1 priority 105
 standby 1 preempt
 standby 1 authentication md5 key-string c1sco123
```

Procedure 12 Configure transit network

You use this network for router-router communication and to avoid hairpinning. The transit network should use an additional subinterface on the router interface that is already being used for data or voice.

Step 1: Configure the transit network between the two routers.

At a remote site location where there are multiple border routers, the interface throughput delay setting should be set to influence the routing protocol path preference. Set the internal LAN path to 250000 usec. The delay command is entered in 10 usec units.

There are no end stations connected to this network, so HSRP and DHCP are not required.

```
interface [interface type][number].[sub-interface number]
 encapsulation dot1Q [dot1q VLAN tag]
 ip address [transit net address] [transit net netmask]
 ip pim sparse-mode
 delay [in 10 usecs]
```

Example: Layer 2 EtherChannel

```
interface PortChannel0/2.99
 description Transit Net
 encapsulation dot1Q 99
 ip address 10.7.16.10 255.255.255.252
 ip pim sparse-mode
 delay 25000
```

Example: Layer 2 Link

```
interface GigabitEthernet0/2.99
 description Transit Net
 encapsulation dot1Q 99
 ip address 10.7.16.10 255.255.255.252
 ip pim sparse-mode
 delay 25000
```

Procedure 13 Configure EIGRP on the transit network

A single EIGRP process runs on the DMVPN spoke router, which has already been enabled during the configuration of the DMVPN tunnel. All interfaces on the router are EIGRP interfaces, but only the DMVPN tunnel interface and transit network are non-passive. The network range must include all interface IP addresses either in a single network statement or in multiple network statements.

EIGRP stub routers normally do not exchange routes with other stub routers, which is problematic in a dual router remote site. It is useful to maintain the benefits of EIGRP stub routing in this type of design. This requires the configuration of stub route leaking between the two remote-site routers for full route reachability. This is needed so that if a router loses its DMVPN link, it knows to how to reach the WAN from the other router at the branch.

In order to prevent the remote site from becoming a transit site during certain failure conditions when stub route leaking is enabled, you must also configure a distribute-list on the tunnel interfaces to control route advertisement.

Step 1: Configure the transit network subinterface as non-passive.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   af-interface Port-channel2.99
    no passive-interface
   exit-af-interface
 exit-address-family
```

Step 2: Configure EIGRP neighbor authentication. Neighbor authentication enables the secure establishment of peering adjacencies and exchange route tables over the transit network interface.

```
key chain LAN-KEY
 key 1
  key-string c1sco123

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Port-channel2.99
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
 exit-address-family
```

Step 3: Configure stub route leaking.

A simple route-map statement with no match statements matches all routes, which permits full route leaking between two routers configured as EIGRP stub.

```
route-map STUB-LEAK-ALL permit 100
 description Leak all routes

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  eigrp stub connected summary redistributed leak-map STUB-LEAK-ALL
 exit-address-family
```

Deploying an IWAN Remote-Site Distribution Layer

PROCESS

Connecting Remote-site Router to Distribution Layer

1. Connect router to distribution layer
2. Configure EIGRP for distribution layer link
3. Configure transit network for dual router design

This process helps you configure a DMVPN spoke router for an IWAN remote site and connect to a distribution layer.

This process covers the IWAN hybrid design model and the IWAN dual-Internet design model. Use this process to:

- Connect a distribution layer to a router in the single-router, dual-link design.
- Connect a distribution layer to the first router of the dual-router, dual-link design.

The distribution layer remote-site options are shown in the following figures.

Figure 26 IWAN single router remote-site: Connection to distribution layer

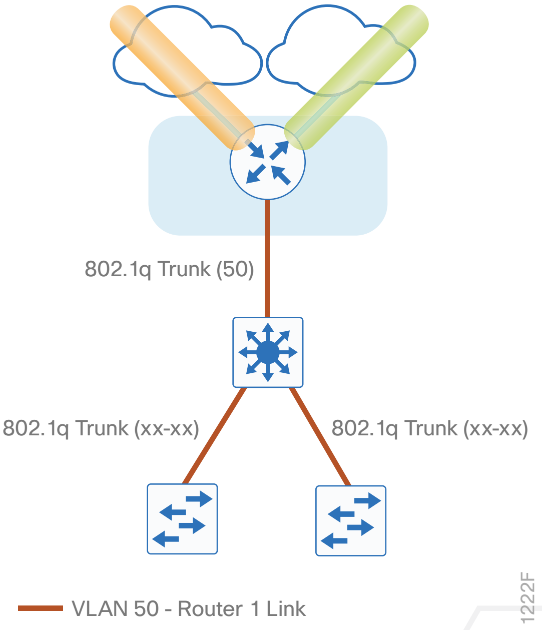

Figure 27 IWAN dual router remote-site: Connection to distribution layer

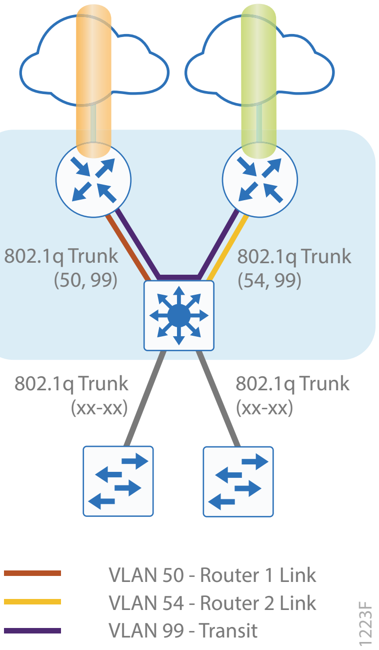

Procedure 1 Connect router to distribution layer

Reader Tip

This guide includes only the additional steps for completing the distribution layer configuration. For complete distribution layer configuration details, see the Campus LAN Layer 2 Access with Simplified Distribution Deployment Guide .

Layer 2 EtherChannels are used to interconnect the remote-site router to the distribution layer in the most resilient method possible. This connection allows for multiple VLANs to be included on the EtherChannel as necessary.

Step 1: Configure port-channel interface on the router.

```
interface Port-channel1
 description EtherChannel link to RS42-D3850
 no shutdown
```

Step 2: Configure the port channel sub-interfaces and assign IP addresses.

After you have enabled the interface, map the appropriate sub-interfaces to the VLANs on the distribution layer switch. The sub-interface number does not need to equate to the 802.1Q tag, but making them the same simplifies the overall configuration.

The sub-interface configured on the router corresponds to a VLAN interface on the distribution-layer switch.

Traffic is routed between the devices with the VLAN acting as a point-to-point link.

```
interface Port-channel1.50
 description R1 routed link to distribution layer
 encapsulation dot1Q 50
 ip address 10.7.208.1 255.255.255.252
 ip pim sparse-mode
```

Step 3: Configure EtherChannel member interfaces on the router.

Configure the physical interfaces to tie to the logical port-channel using the channel-group command. The number for the port-channel and channel-group must match. Not all router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet0/0/2
 description RS42-D3850 Gig1/1/1

interface GigabitEthernet0/0/3
 description RS42-D3850 Gig2/1/1

interface range GigabitEthernet0/0/2, GigabitEthernet0/0/3
  no ip address
  channel-group 1
  no shutdown
```

Step 4: Configure VLAN on the distribution layer switch.

```
vlan 50
 name R1-link
```

Step 5: Configure Layer 3 on the distribution layer switch.

Configure a VLAN interface, also known as a switch virtual interface (SVI), for the new VLAN added. The SVI is used for point to point IP routing between the distribution layer and the WAN router.

```
interface Vlan50
 ip address 10.5.208.2 255.255.255.252
 ip pim sparse-mode
 no shutdown
```

Step 6: Configure EtherChannel member interfaces on the distribution layer switch.

Connect the router EtherChannel uplinks to separate switches in the distribution layer switches or stack.

The physical interfaces that are members of a Layer 2 EtherChannel are configured prior to configuring the logical port-channel interface. Doing the configuration in this order allows for minimal configuration and reduces errors because most of the commands entered to a port-channel interface are copied to its members interfaces and do not require manual replication.

Configure two or more physical interfaces to be members of the EtherChannel. It is recommended that they are added in multiples of two. Also, apply the egress QoS macro that was defined in the platform configuration procedure to ensure traffic is prioritized appropriately.

Not all connected router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet1/1/1
 description Link to RS42-4451X-1 Gig0/0/2

interface GigabitEthernet2/1/1
 description Link to RS42-4451X-1 Gig0/0/3

interface range GigabitEthernet1/1/1, GigabitEthernet2/1/1
 switchport
 channel-group 1 mode on
 logging event link-status
 logging event trunk-status
 logging event bundle-status
 load-interval 30
 macro apply EgressQoS
```

Step 7: Configure EtherChannel trunk on the distribution layer switch.

An 802.1Q trunk is used which allows the router to provide the Layer 3 services to all the VLANs defined on the distribution layer switch. The VLANs allowed on the trunk are pruned to only the VLANs that are active on the distribution layer switch. When using EtherChannel the interface type will be port-channel and the number must match the channel group configured in Step 3. DHCP Snooping and ARP inspection are set to trust.

```
interface Port-channel1
 description EtherChannel link to RS42-4451X-1
 switchport trunk allowed vlan 50
 switchport mode trunk
 spanning-tree portfast trunk
 load-interval 30
 no shutdown
```

The Cisco Catalyst 3750 Series Switch requires the **switchport trunk encapsulation dot1q** command

Procedure 2 Configure EIGRP for distribution layer link

A single EIGRP process runs on the DMVPN spoke router, which has already been enabled during configuration of the DMVPN tunnel. All interfaces on the router are EIGRP interfaces, but only the DMVPN tunnel interface and the link to the distribution layer are non-passive. The network range must include all interface IP addresses either in a single network statement or in multiple network statements.

EIGRP stub routers normally do not exchange routes with other stub routers and distribution switches, which is problematic in a dual-router remote site. It is useful to maintain the benefits of EIGRP stub routing in this type of design. This requires the configuration of stub route leaking between the remote-site routers and distribution layer switch for full route reachability.

In order to prevent the remote site from becoming a transit site during certain failure conditions when stub route leaking is enabled, you must also configure a distribute-list on the tunnel interfaces to control route advertisement.

Step 1: Configure the distribution layer link sub-interface as non-passive.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   af-interface Port-channel1.50
    no passive-interface
   exit-af-interface
 exit-address-family
```

Step 2: Configure EIGRP neighbor authentication. Neighbor authentication enables the secure establishment of peering adjacencies and exchange route tables over the interface.

```
key chain LAN-KEY
 key 1
  key-string c1sco123

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Port-channel1.50
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
 exit-address-family
```

Step 3: Configure stub route leaking.

A simple route-map statement with no match statements matches all routes, which permits full route leaking between two routers configured as EIGRP stub.

```
route-map STUB-LEAK-ALL permit 100
 description Leak all routes

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  eigrp stub connected summary redistributed leak-map STUB-LEAK-ALL
 exit-address-family
```

Step 4: On the distribution layer switch VLAN interface, enable EIGRP.

EIGRP is already configured on the distribution layer switch. The VLAN interface that connects to the router must be configured for EIGRP neighbor authentication and as a non-passive EIGRP interface.

```
key chain LAN-KEY
 key 1
  key-string c1sco123

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Vlan50
   authentication mode md5
   authentication key-chain LAN-KEY
   no passive-interface
  exit-af-interface
exit-address-family
```

Procedure 3 Configure transit network for dual router design

This procedure is only for dual-router remote sites.

Step 1: Configure the transit network between the two routers.

You use this network for router-router communication and to avoid hairpinning. The transit network should use an additional sub-interface on the EtherChannel interface that is already used to connect to the distribution layer.

At the remote site location where there are multiple border routers, the interface throughput delay setting should be set to influence the routing protocol path preference. Set the internal LAN path to 250000 usec. The delay command is entered in 10 usec units.

There are no end stations connected to this network so HSRP and DHCP are not required. The transit network uses Layer 2 pass through on the distribution layer switch, so no SVI is required.

```
interface Port-channel1.99
 description Transit Net
 encapsulation dot1Q 99
 ip address 10.7.208.9 255.255.255.252
 ip pim sparse-mode
 delay 25000
```

Step 2: Enable EIGRP on the transit net interface on the router.

The transit network must be a non-passive EIGRP interface.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
af-interface Port-channel1.99
   authentication mode md5
   authentication key-chain LAN-KEY
   no passive-interface
 exit-af-interface
exit-address-family
```

Step 3: Configure transit network VLAN on the distribution layer switch.

```
vlan 99
 name Transit-net
```

Step 4: Add transit network VLAN to the existing distribution layer switch EtherChannel trunk.

```
interface Port-channel1
 switchport trunk allowed vlan add 99
```

PROCESS

Connecting Remote-Site Router to Distribution Layer (Router 2)

1. Connect router to distribution layer
2. Configure EIGRP for distribution layer link

This process helps you configure a DMVPN spoke router for an IWAN remote site and connect to a distribution layer.

This process covers the IWAN hybrid design model and the IWAN dual-Internet design model. Use this process to connect a distribution layer to the second router of the dual-router, dual-link design.

The dual-router distribution layer remote-site option is shown in the following figure.

Figure 28 *WAN remote site: Connection to distribution layer*

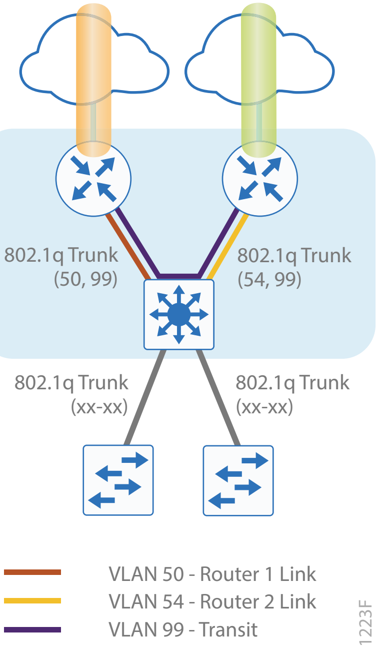

Procedure 1 Connect router to distribution layer

Reader Tip

Please refer to the Campus Wired LAN Design Guide for complete distribution layer configuration details. This guide only includes the additional steps to complete the distribution layer configuration.

Layer 2 EtherChannels are used to interconnect the remote-site router to the distribution layer in the most resilient method possible. This connection allows for multiple VLANs to be included on the EtherChannel as necessary.

Step 1: Configure port-channel interface on the router.

```
interface Port-channel2
 description EtherChannel link to RS42-D3850
 no shutdown
```

Step 2: Configure the port channel sub-interfaces and assign IP address.

After you have enabled the interface, map the appropriate sub-interfaces to the VLANs on the distribution layer switch. The sub-interface number does not need to equate to the 802.1Q tag, but making them the same simplifies the overall configuration.

The sub-interface configured on the router corresponds to a VLAN interface on the distribution-layer switch. Traffic is routed between the devices with the VLAN acting as a point-to-point link.

```
interface Port-channel2.54
 description R2 routed link to distribution layer
 encapsulation dot1Q 54
 ip address 10.7.208.5 255.255.255.252
 ip pim sparse-mode
```

Step 3: Configure the transit network interface on the router.

At a remote site location where there are multiple border routers, the interface throughput delay setting should be set to influence the routing protocol path preference. Set the internal LAN path to 250000 usec. The delay command is entered in 10 usec units.

```
interface Port-channel2.99
 description Transit Net
 encapsulation dot1Q 99
 ip address 10.7.208.10 255.255.255.252
 ip pim sparse-mode
 delay 25000
```

Step 4: Configure EtherChannel member interfaces on the router.

Configure the physical interfaces to tie to the logical port-channel using the channel-group command. The number for the port-channel and channel-group must match. Not all router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet0/0/2
 description RS42-D3850X Gig1/1/2

interface GigabitEthernet0/0/3
 description RS42-D3850X Gig2/1/2
```

```
interface range GigabitEthernet0/0/2, GigabitEthernet0/0/3
  no ip address
  channel-group 2
  no shutdown
```

Step 5: Configure VLAN on the distribution layer switch.

```
vlan 54
 name R2-link
```

Step 6: Configure Layer 3 on the distribution layer switch.

Configure a VLAN interface, also known as a SVI, for the new VLAN added. The SVI is used for point to point IP routing between the distribution layer and the WAN router.

```
interface Vlan54
 ip address 10.7.208.6 255.255.255.252
 ip pim sparse-mode
 no shutdown
```

Step 7: Configure EtherChannel member interfaces on the distribution layer switch.

Connect the router EtherChannel uplinks to separate switches in the distribution layer switches or stack.

The physical interfaces that are members of a Layer 2 EtherChannel are configured prior to configuring the logical port-channel interface. Doing the configuration in this order allows for minimal configuration and reduces errors because most of the commands entered to a port-channel interface are copied to its members interfaces and do not require manual replication.

Configure two or more physical interfaces to be members of the EtherChannel. It is recommended that they are added in multiples of two. Also, apply the egress QoS macro that was defined in the platform configuration procedure to ensure traffic is prioritized appropriately.

Not all connected router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet1/1/2
 description Link to RS42-4451X-2 Gig0/0/2
interface GigabitEthernet2/1/2
 description Link to RS42-4451X-2 Gig0/0/3

interface range GigabitEthernet1/1/2, GigabitEthernet2/1/2
 switchport
 channel-group 2 mode on
 logging event link-status
 logging event trunk-status
 logging event bundle-status
 load-interval 30
 macro apply EgressQoS
```

Step 8: Configure EtherChannel trunk on the distribution layer switch.

An 802.1Q trunk is used which allows the router to provide the Layer 3 services to all the VLANs defined on the distribution layer switch. The VLANs allowed on the trunk are pruned to only the VLANs that are active on the distribution layer switch. When using EtherChannel the interface type will be port-channel and the number must match the channel group configured in Step 3. DHCP Snooping and ARP inspection are set to trust.

```
interface Port-channel2
 description EtherChannel link to RS42-4451X-2
 switchport trunk allowed vlan 54,99
 switchport mode trunk
 spanning-tree portfast trunk
 no shutdown
```

The Cisco Catalyst 3750 Series Switch requires the **switchport trunk encapsulation dot1q** command

Procedure 2 Configure EIGRP for distribution layer link

A single EIGRP process runs on the DMVPN spoke router, which has already been enabled during DMVPN tunnel configuration. All interfaces on the router are EIGRP interfaces, but only the DMVPN tunnel interface, the link to the distribution layer, and the transit network link are non-passive. The network range must include all interface IP addresses either in a single network statement or in multiple network statements.

EIGRP stub routers normally do not exchange routes with other stub routers and distribution switches, which is problematic in a dual-router remote site. It is useful to maintain the benefits of EIGRP stub routing in this type of design. This requires the configuration of stub route leaking between the remote-site routers and distribution layer switch for full route reachability.

In order to prevent the remote site from becoming a transit site during certain failure conditions when stub route leaking is enabled, you must also configure a distribute-list on the tunnel interfaces in order to control route advertisement.

Step 1: Configure the distribution layer link subinterface as non-passive.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   af-interface Port-channel2.54
    no passive-interface
   exit-af-interface
 exit-address-family
```

Step 2: Configure the transit network link subinterface as non-passive.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   af-interface Port-channel2.99
    no passive-interface
   exit-af-interface
 exit-address-family
```

Step 3: Configure EIGRP neighbor authentication. Neighbor authentication enables the secure establishment of peering adjacencies and exchange route tables over the interface.

```
key chain LAN-KEY
 key 1
  key-string c1sco123

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Port-channel2.54
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
af-interface Port-channel2.99
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
 exit-address-family
```

Step 4: Configure stub route leaking.

A simple route-map statement with no match statements matches all routes, which permits full route leaking between two routers configured as EIGRP stub.

```
route-map STUB-LEAK-ALL permit 100
 description Leak all routes

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  eigrp stub connected summary redistributed leak-map STUB-LEAK-ALL
 exit-address-family
```

Step 5: Enable EIGRP on distribution layer switch VLAN interface.

EIGRP is already configured on the distribution layer switch. The VLAN interface that connects to the router must be configured for EIGRP neighbor authentication and as a non-passive EIGRP interface.

```
key chain LAN-KEY
 key 1
  key-string c1sco123

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Vlan54
   authentication mode md5
   authentication key-chain LAN-KEY
   no passive-interface
  exit-af-interface
exit-address-family
```

Deploying IWAN Performance Routing

Performance Routing Version 3 (PfRv3) consists of two major Cisco IOS components, an MC and a BR. The MC defines the policies and applies them to various traffic classes that traverse the BR systems. The MC can be configured to learn and control traffic classes on the network.

- The MC is the policy decision-maker. At a large site, such as a data center or campus, the MC is a standalone router. For smaller locations, the MC is typically collocated (configured) on the same platform as the BR. As a general rule, the large locations manage more network prefixes and applications than a remote site deployment so they consume more CPU and memory resources for the MC function. Therefore, Cisco recommends a dedicated router for the MC at large sites.
- The BR is in the data-forwarding path. A BR collects data from its Performance Monitor cache and smart probes, provides a degree of aggregation of this information, and influences the packet-forwarding path as directed by the MC to optimize traffic.

The remote site typically manages fewer active TCs, which are made up of prefixes and applications. In most remote site deployments, it is possible to co-locate the MC and BR on the same hardware platform. CPU and memory utilization should be monitored on MC platforms, and if utilization is high, the network manager should consider an MC platform with a higher capacity CPU and memory. The MC communicates with the border routers over an authenticated TCP socket but has no requirement for populating its own IP routing table with anything more than a route to reach the border routers.

Because PfRv3 is an intelligent path selection technology, there must be at least two external interfaces under the control of PfRv3 and at least one internal interface. IWAN is not limited to two paths. Three paths per preference logic are supported, so you can configure the path-preference as follows:

- **Path-preference**–MPLS1, MPLS2 and MPLS3
- **Fallback**–INET1, INET2 and INET3
- **Next-fallback**–INET4, INET5, INET6

Path-preference can also include fallback to the routing protocol if there is no fallback provider, or the traffic can be dropped if the primary provider is not available.

There must be at least one BR configured. If only one BR is configured, then both external interfaces are attached to the single BR. If more than one BR is configured, then the two or more external interfaces are configured across these BR platforms. External links, or exit points, are therefore owned by the BR; for a supported IWAN solution, the external links must be logical (tunnel) interfaces using a DMVPN overlay.

There are four different roles a device can play in a standard PfRv3 configuration:

- **Hub Master Controller**—The hub MC is the MC at the primary WAN aggregation site. This is the MC device where all PfRv3 policies are configured. It also acts as MC for that site and makes path optimization decision. There is only one hub MC per IWAN domain, and you cannot configure the hub BR on the same router platform.
- **Hub Border Router**—This is a BR at the hub MC site. This is the device where WAN interfaces terminate. There can be only one WAN interface on the hub device. There can be one or more hub BRs. On the Hub BRs, PfRv3 must be configured with:
    - The address of the local MC
    - The path name on external interfaces
    - The path ID on external interfaces
- **Branch Master Controller**—The Branch MC is the MC at the branch-site. There is no policy configuration on this device. It receives policy from the Hub MC. This device acts as MC for that site for making path optimization decision. The configuration includes the IP address of the hub MC.
- **Branch Border Router**—This is a BR at the branch-site. The configuration on this device enables BR functionality and includes the IP address of the site local MC. The WAN interface that terminates on the device is detected automatically.

The first design model is the IWAN hybrid, which uses a primary MPLS transport paired with Internet VPN as a secondary transport. In this design model, the MPLS WAN provides SLA class of service guarantees for key applications. The second design model is the IWAN Dual Internet, which uses a pair of Internet service providers to further reduce cost while leveraging PfR in order to mitigate network performance problems on a single Internet provider. The PfR configuration is the same for both design models.

The following diagrams show the four different device roles and where they fit into the IWAN hybrid design model.

Figure 29 IWAN hybrid design model: PfR hub location

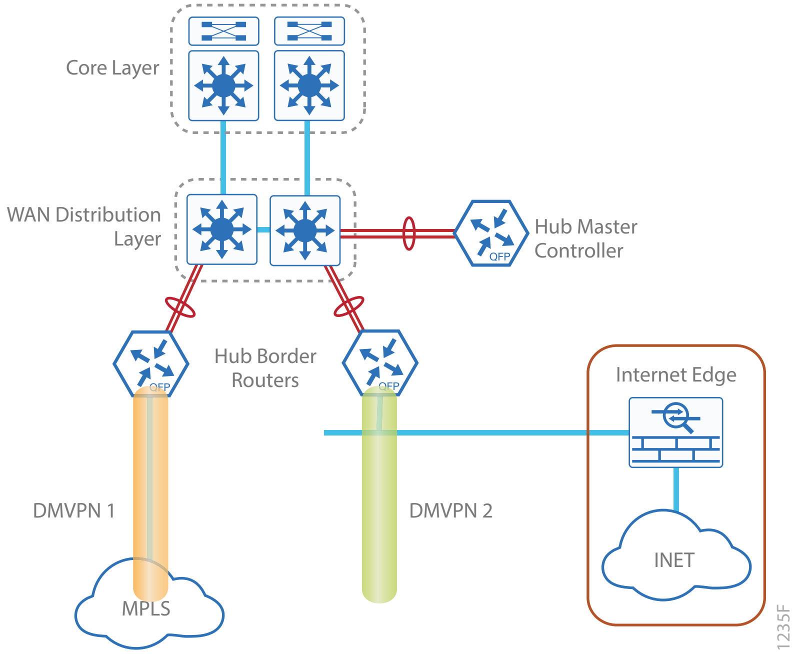

Figure 30 IWAN hybrid design model: PfR branch locations

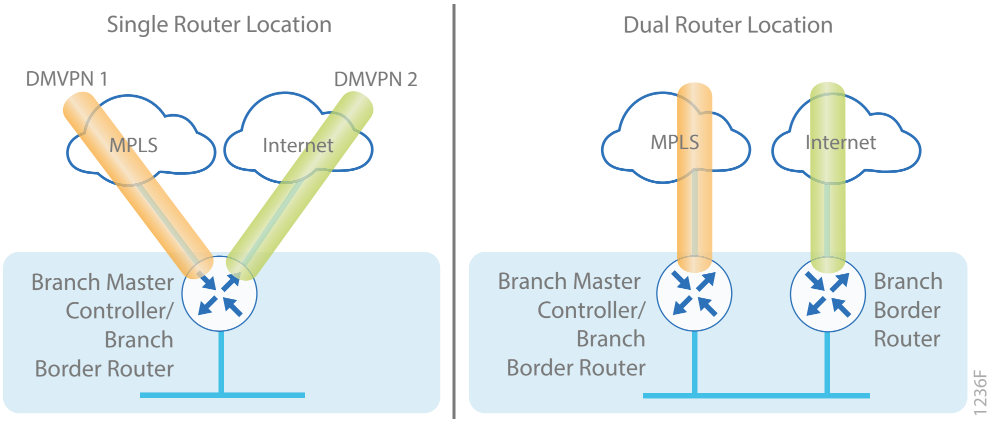

PROCESS

Configuring Hub Master Controller

1. Connect router to distribution layer
2. Configure the Hub MC platform
3. Configure connectivity to the LAN

This section describes configuring the PfR Hub MC as a new router. Only the core relevant features are included.

Table 60 Hub MC IP addresses

IWAN design model	Host name	Loopback IP address	Port-channel IP address
Hybrid	MC-HY-CSR1000V-1	10.6.32.251/32	10.6.32.151/26
Dual Internet	MC-DI-ASR1004-1	10.6.32.252/32	10.6.32.152/26

Procedure 1 Connect router to distribution layer

Reader Tip

Refer to the Campus Wired LAN Design Guide for complete distribution layer configuration details. This guide only includes the additional steps to complete the distribution layer configuration.

Step 1: If a VLAN does not already exist for the hub MC on the distribution layer switch, configure it now.

```
vlan 350
 name WAN_Service_Net
```

Step 2: If the Layer 3 SVI has not yet been configured, configure it now.

Be sure to configure a VLAN interface (SVI) for every new VLAN you add, so devices in the VLAN can communicate with the rest of the network.

```
interface Vlan350
 ip address 10.6.32.129 255.255.255.192
 no shutdown
```

Next, configure EtherChannel member interfaces.

Tech Tip

EtherChannel is a logical interface that bundles multiple physical LAN links into a single logical link.

Step 3: Connect the hub MC EtherChannel uplinks in order to separate switches in the distribution layer switches or stack, and then configure two physical interfaces to be members of the EtherChannel.

Also, apply the egress QoS macro that was defined in the platform configuration procedure. This ensures traffic is prioritized appropriately. The EtherChannel provides extra resiliency for the hub MC in case there is a link, line card or switch failure.

Tech Tip

Configure the physical interfaces that are members of a Layer 2 EtherChannel prior to configuring the logical port-channel interface. Doing the configuration in this order allows for minimal configuration and reduces errors because most of the commands entered to a port-channel interface are copied to its members interfaces and do not require manual replication.

```
interface GigabitEthernet 1/0/15
 description MC-HY-CSR1000V-1 (WAN-IWAN-C220-1) (vmnic4)

interface GigabitEthernet 2/0/15
 description MC-HY-CSR1000V-1 (WAN-IWAN-C220-1) (vmnic5)

interface range GigabitEthernet 1/0/15, GigabitEthernet 2/0/15
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 350
 switchport mode trunk
 logging event trunk-status
 load-interval 30
 macro description EgressQoS
 spanning-tree portfast trunk
 channel-group 21 mode on
```

Next, configure the EtherChannel. Access mode interfaces are used for the connection to the hub MCs.

Step 4: Assign the VLAN created at the beginning of the procedure to the interface. When using EtherChannel, the port-channel number must match the channel group configured in Step 3.

```
interface Port-channel 21
 description MC-HY-CSR1000V-1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 350
 switchport mode trunk
 logging event trunk-status
 logging event bundle-status
 spanning-tree portfast trunk
 no shutdown
```

Step 5: Allow the routing protocol to form neighbor relationships across the vlan interface.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Vlan350
   no passive-interface
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
 exit-address-family
```

Procedure 2 Configure the Hub MC platform

Within this design, there are features and services that are common across all PfR routers. In this procedure, you configure system settings that simplify and secure the management of the solution.

To complete the base configuration for this router, follow the steps in “Configure the platform base features” in Appendix C.

Procedure 3 Configure connectivity to the LAN

Step 1: Configure IP unicast routing authentication key.

```
key chain LAN-KEY
 key 1
   key-string c1sco123
```

Step 2: Configure IP unicast routing using EIGRP named mode.

EIGRP is configured facing the LAN distribution or core layer. In this design, the port-channel interface and the loopback must be EIGRP interfaces. The loopback may remain a passive interface. The network range must include both interface IP addresses, either in a single network statement or in multiple network statements. This design uses a best practice of assigning the router ID to a loopback address.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface default
   passive-interface
  exit-af-interface
 network 10.6.0.0 0.1.255.255
 eigrp router-id 10.6.32.251
 exit-address-family
```

Any links to adjacent distribution layers should be Layer 3 links or Layer 3 EtherChannels.

Step 3: Configure a Layer 3 interface.

```
interface Port-channel21
  description IW-WAN-D3750X
  ip address 10.6.32.151 255.255.255.192
  no shutdown
```

Step 4: Configure EtherChannel member interfaces.

Configure the physical interfaces to tie to the logical port-channel by using the **channel-group** command. The number for the port-channel and channel-group must match. Not all router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet1
 description IW-WAN-D3750X Gig1/0/15

interface GigabitEthernet2
 description IW-WAN-D3750X Gig2/0/15

interface range GigabitEthernet1, GigabitEthernet2
  no ip address
  cdp enable
  channel-group 21
  no shutdown
```

Step 5: Configure the EIGRP interface.

Allow EIGRP to form neighbor relationships across the interface to establish peering adjacencies and exchange route tables. In this step, you configure EIGRP authentication by using the authentication key specified in the previous procedure.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Port-channel21
   no passive-interface
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
 exit-address-family
```

PROCESS

Configuring PfR for Hub Location

1. Verify IP connectivity to remote site loopback interfaces
2. Configure prefixes for the enterprise and data center
3. Configure PfR domain in the hub MC
4. Configure PfR domain in the hub BR
5. Verify PfR domain is operational on the hub MC

All sites belong to an PfR domain where the remote site MCs are peered together. Peering has been greatly enhanced in PfRv3 which allows site information exchange and single touch provisioning.

PfRv3 has simplified policies with pre-existing templates. The policy configuration for the PfR domain is done in the hub MC and the information is distributed to all sites via MC peering. This not only simplifies provisioning substantially but also makes the policy consistent across the entire IWAN network.

PfRv3 uses Unified Monitor (also called Performance Monitor) to monitor traffic going into WAN links and traffic coming from the WAN links. It monitors performance metrics per differentiated service code point (DSCP) rather than monitoring on per-flow or per-prefix basis. When application-based policies are used, the MC will use a mapping table between the Application Name and the DSCP discovered. This reduces the number of records significantly. PfRv3 relies on performance data measured on the existing data traffic on all paths whenever it can, thereby reducing the need of synthetic traffic. Furthermore, the measurement data is not exported unless there is a violation, which further reduces control traffic and processing of those records.

PfRv3 is also VRF-aware and instances of the MC work under a VRF.

Procedure 1 Verify IP connectivity to remote site loopback interfaces

It is mandatory to use loopback interfaces for the peering traffic between the BR and MC routers. For this design, you put the loopback addresses into a specific subnet range, so they are easily identified in the routing table. The loopback address ranges for the remote sites are as follows:

Table 61 Remote-site loopback IP address ranges

IWAN design model	Tunnel type	Loopback 0 address range
Hybrid—Primary Router	MPLS	10.255.241.0/24
Hybrid—Secondary Router	INET	10.255.242.0/24
Dual Internet—Primary Router	INET1	10.255.243.0/24
Dual Internet—Secondary Router	INET2	10.255.244.0/24

Step 1: Verify that the loopback 0 interfaces on each of your remote sites are reachable from the hub MC by using the **show ip route** command.

This example shows a loopback address range of 10.255.241.0/24 for nine remote site primary routers and an address range of 10.255.242.0/24 for four remote site secondary routers.

```
MC-HY-CSR1000V-1# show ip route | include 10.255.241
D        10.255.241.11/32 [90/25610880] via 10.6.32.129, 1w2d, Port-channel21
D        10.255.241.12/32 [90/25610880] via 10.6.32.129, 1w2d, Port-channel21
D        10.255.241.21/32 [90/25610880] via 10.6.32.129, 1w2d, Port-channel21
D        10.255.241.22/32 [90/25610880] via 10.6.32.129, 1w2d, Port-channel21
D        10.255.241.31/32 [90/25610880] via 10.6.32.129, 1w2d, Port-channel21
D        10.255.241.32/32 [90/25610880] via 10.6.32.129, 1w2d, Port-channel21
D        10.255.241.41/32 [90/25610880] via 10.6.32.129, 1w2d, Port-channel21
D        10.255.241.42/32 [90/25610880] via 10.6.32.129, 1w2d, Port-channel21
D        10.255.241.51/32 [90/25610880] via 10.6.32.129, 1w3d, Port-channel21

MC-HY-CSR1000V-1# show ip route | include 10.255.242
D        10.255.242.12/32 [90/25613440] via 10.6.32.129, 1w1d, Port-channel21
D        10.255.242.22/32 [90/25613440] via 10.6.32.129, 1w2d, Port-channel21
D        10.255.242.32/32 [90/25613440] via 10.6.32.129, 1w2d, Port-channel21
D        10.255.242.42/32 [90/25613440] via 10.6.32.129, 1w2d, Port-channel21
```

Procedure 2 Configure prefixes for the enterprise and data center

Before the configuration of PfRv3 on the hub MC, you must create a prefix list for the enterprise and data center. The enterprise-prefix list covers the range of IP addresses to be controlled and optimized within this IWAN domain. Prefixes outside of the enterprise-prefix list will not be controlled by application policies, but they will be load-balanced.

The site-prefix range includes the prefixes at this specific site, which is normally a WAN aggregation or data center (DC) site. Site-prefixes are typically statically defined at WAN aggregation and DC sites and discovered automatically at remote sites.

Tech Tip

The ip prefix-list options ge and le are not supported by PfR.

Step 1: Create the enterprise prefix list.

```
ip prefix-list [prefix-list-name] seq [value] permit [prefix list]
```

Example

This example shows a contiguous block of private address space from 10.4.0.0 to 10.7.255.255, which covers all the IP addresses within this IWAN PfR domain. It does not include the router loopback address range of 10.255.240.0 to 10.255.247.255 because you do not want PfR controlling those prefixes.

```
ip prefix-list ENTERPRISE-PREFIXES seq 10 permit 10.4.0.0/14
```

Step 2: Create the primary site prefix list.

```
ip prefix-list [prefix-list-name] seq [value] permit [prefix list]
```

Example

This example shows a data center network with two class B private address blocks of 10.4.0.0 and 10.6.0.0.

```
ip prefix-list DC1-PREFIXES seq 10 permit 10.4.0.0/16
ip prefix-list DC1-PREFIXES seq 20 permit 10.6.0.0/16
```

Procedure 3 Configure PfR domain in the hub MC

Domain policies are configured on the hub MC. These policies are distributed to branch MCs by using the peering infrastructure. All sites that are in the same domain will share the same set of PfR policies. Policies can be based on DSCP or on application names.

Policies are created using preexisting templates, or they can be customized with manually defined thresholds for delay, loss and jitter.

Table 62 PfR domain pre-defined policy templates

Pre-defined template	Priority	Threshold definition
Voice	1	one-way-delay threshold 150 msec
	2	packet-loss-rate threshold 1.0 percent
	2	byte-loss-rate threshold 1.0 percent
	3	jitter threshold 30000 usec
Real-time-video	1	packet-loss-rate threshold 1.0 percent
	1	byte-loss-rate threshold 1.0 percent
	2	one-way-delay threshold 150 msec
	3	jitter threshold 20000 usec
Low-latency-data	1	one-way-delay threshold 100 msec
	2	packet-loss-rate threshold 5.0 percent
	2	byte-loss-rate threshold 5.0 percent
Bulk-data	1	one-way-delay threshold 300 msec
	2	packet-loss-rate threshold 5.0 percent
	2	byte-loss-rate threshold 5.0 percent
Best-effort	1	one-way-delay threshold 500 msec
	2	packet-loss-rate threshold 10.0 percent
	2	byte-loss-rate threshold 10.0 percent
Scavenger	1	one-way-delay threshold 500 msec
	2	packet-loss-rate threshold 50.0 percent
	2	byte-loss-rate threshold 50.0 percent

To avoid unwanted channel unreachable messages, it is recommended that you change the value of the **channel-unreachable-timer** command from its default setting of 1 second. The command is under the **advanced** setting and the value is specified in seconds.

The NMS collector IP address and port number are defined in the hub MC. The information is automatically propagated to devices in the IWAN domain. If you do not want to use a single collector for your entire network, you can specify a different IP address and port number in the IWAN domain for each device.

Step 1: Create the hub MC domain.

```
domain [name]
 vrf [name]
  master hub (create the hub MC)
   source-interface [interface]
   site-prefixes prefix-list [prefixes from previous procedure]
   password [password]
   advanced
```

```
    channel-unreachable-timer [value in seconds]
   enterprise-prefix prefix-list [prefixes from previous procedure]
   collector [IP address of NMS] port [NetFlow]
```

Example

```
domain iwan
 vrf default
  master hub
   source-interface Loopback0
   site-prefixes prefix-list DC1-PREFIXES
   password c1sco123
   advanced
    channel-unreachable-timer 4
   enterprise-prefix  prefix-list ENTERPRISE-PREFIXES
   collector 10.4.48.36 port 9991
```

Step 2: Create the hub MC policy.

```
domain [name]
 vrf [name]
  master hub (configure the hub MC with additional commands)
   load-balance (load balance the traffic not specified in a class)
   class [name] sequence [value] (repeat for each class)
    match dscp [value] policy [name] (repeat for each dscp value)
    path-preference [primary] fallback [secondary] (path names)
```

Example

The policies use the PfR predefined templates. The path preference for policies other than scavenger and default is to use MPLS unless the delay, jitter, and loss values on the path fall outside the values specified in the templates. Scavenger and default use INET and fallback to MPLS, while the rest of the traffic will be load-balanced between the two paths.

```
domain iwan
 vrf default
  master hub
   load-balance
   class VOICE sequence 10
    match dscp ef policy voice
    path-preference MPLS fallback INET
   class INTERACTIVE-VIDEO sequence 20
    match dscp cs4 policy real-time-video
    match dscp af41 policy real-time-video
    match dscp af42 policy real-time-video
    match dscp af43 policy real-time-video
    path-preference MPLS fallback INET
   class LOW_LATENCY_DATA sequence 30
    match dscp cs2 policy low-latency-data
    match dscp cs3 policy low-latency-data
    match dscp af21 policy low-latency-data
    match dscp af22 policy low-latency-data
    match dscp af23 policy low-latency-data
    path-preference MPLS fallback INET
   class BULK_DATA sequence 40
    match dscp af11 policy bulk-data
    match dscp af12 policy bulk-data
    match dscp af13 policy bulk-data
    path-preference MPLS fallback INET
   class SCAVENGER sequence 50
    match dscp cs1 policy scavenger
    path-preference INET fallback MPLS
   class DEFAULT sequence 60
    match dscp default policy best-effort
    path-preference INET fallback MPLS
```

Step 3: Verify the hub MC policy configuration by using the **show domain [name] master policy** command.

In this example, the **class default** listed last is automatically added when the **load-balance** key word is used in the hub MC policy.

```
MC-HY-CSR1000V-1#show domain iwan master policy
 No Policy publish pending
 Last publish Status : Peering Success
 Total publish errors : 0
--------------------------------------------------------------------


  class VOICE sequence 10
    path-preference MPLS fallback INET
    class type: Dscp Based
```

```
      match dscp ef policy voice
        priority 2 packet-loss-rate threshold 1.0 percent
        priority 1 one-way-delay threshold 150 msec
        priority 3 jitter threshold 30000 usec
        priority 2 byte-loss-rate threshold 1.0 percent

  class INTERACTIVE_VIDEO sequence 20
    path-preference MPLS fallback INET
    class type: Dscp Based
      match dscp cs4 policy real-time-video
        priority 1 packet-loss-rate threshold 1.0 percent
        priority 2 one-way-delay threshold 150 msec
        priority 3 jitter threshold 20000 usec
        priority 1 byte-loss-rate threshold 1.0 percent
      match dscp af41 policy real-time-video
        priority 1 packet-loss-rate threshold 1.0 percent
        priority 2 one-way-delay threshold 150 msec
        priority 3 jitter threshold 20000 usec
        priority 1 byte-loss-rate threshold 1.0 percent
      match dscp af42 policy real-time-video
        priority 1 packet-loss-rate threshold 1.0 percent
        priority 2 one-way-delay threshold 150 msec
        priority 3 jitter threshold 20000 usec
        priority 1 byte-loss-rate threshold 1.0 percent
      match dscp af43 policy real-time-video
        priority 1 packet-loss-rate threshold 1.0 percent
        priority 2 one-way-delay threshold 150 msec
        priority 3 jitter threshold 20000 usec
        priority 1 byte-loss-rate threshold 1.0 percent

  class LOW_LATENCY_DATA sequence 30
    path-preference MPLS fallback INET
    class type: Dscp Based
      match dscp cs2 policy low-latency-data
        priority 2 packet-loss-rate threshold 5.0 percent
        priority 1 one-way-delay threshold 100 msec
        priority 2 byte-loss-rate threshold 5.0 percent
      match dscp cs3 policy low-latency-data
        priority 2 packet-loss-rate threshold 5.0 percent
        priority 1 one-way-delay threshold 100 msec
        priority 2 byte-loss-rate threshold 5.0 percent
        Number of Traffic classes using this policy: 8
      match dscp af21 policy low-latency-data
        priority 2 packet-loss-rate threshold 5.0 percent
        priority 1 one-way-delay threshold 100 msec
        priority 2 byte-loss-rate threshold 5.0 percent
        Number of Traffic classes using this policy: 0
      match dscp af22 policy low-latency-data
        priority 2 packet-loss-rate threshold 5.0 percent
        priority 1 one-way-delay threshold 100 msec
        priority 2 byte-loss-rate threshold 5.0 percent
        Number of Traffic classes using this policy: 0
      match dscp af23 policy low-latency-data
        priority 2 packet-loss-rate threshold 5.0 percent
        priority 1 one-way-delay threshold 100 msec
        priority 2 byte-loss-rate threshold 5.0 percent
        Number of Traffic classes using this policy: 0

  class BULK_DATA sequence 40
    path-preference MPLS fallback INET
    class type: Dscp Based
      match dscp af11 policy bulk-data
        priority 2 packet-loss-rate threshold 5.0 percent
        priority 1 one-way-delay threshold 300 msec
        priority 2 byte-loss-rate threshold 5.0 percent
      match dscp af12 policy bulk-data
        priority 2 packet-loss-rate threshold 5.0 percent
        priority 1 one-way-delay threshold 300 msec
        priority 2 byte-loss-rate threshold 5.0 percent
      match dscp af13 policy bulk-data
        priority 2 packet-loss-rate threshold 5.0 percent
        priority 1 one-way-delay threshold 300 msec
```

```
        priority 2 byte-loss-rate threshold 5.0 percent

  class SCAVENGER sequence 50
    path-preference INET fallback MPLS
    class type: Dscp Based
      match dscp cs1 policy scavenger
        priority 2 packet-loss-rate threshold 50.0 percent
        priority 1 one-way-delay threshold 500 msec
        priority 2 byte-loss-rate threshold 50.0 percent

  class DEFAULT sequence 60
    path-preference INET fallback MPLS
    class type: Dscp Based
      match dscp default policy best-effort
        priority 2 packet-loss-rate threshold 10.0 percent
        priority 1 one-way-delay threshold 500 msec
        priority 2 byte-loss-rate threshold 10.0 percent
        Number of Traffic classes using this policy: 16

  class default
      match dscp all
        Number of Traffic classes using this policy: 42
```

Procedure 4 Configure PfR domain in the hub BR

The hub BRs are also the DMVPN hub WAN aggregation routers for the network. The PfRv3 configurations for standalone BRs are much simpler because they dynamically learn their policy information from the hub MC. The hub BR routers are also used to advertise the path names and path-ids specified in the hub MC configuration.

There is an optional feature called zero-SLA that reduces the probing to the only default class by muting the other DSCP probes. This feature is useful on Internet connections where nothing is guarenteed. Zero-SLA reduces bandwidth usage on metered interfaces like 4G LTE or other Internet connections with a monthly data cap limit.

Tech Tip

If you want to add the zero-SLA feature to an existing hub BR, you must shutdown the DMVPN tunnel interface before configuring. After the feature is added to the hub BR, bring the tunnel interface back up.

Table 63 Hub BR path and IP addresses

Host name	Path	Path ID	Loopback IP address	Zero SLA
VPN-MPLS-ASR1002X-1	MPLS	1	10.6.32.241/32	No
VPN-INET-4451X-2	INET	2	10.6.32.242/32	Yes (optional)

Step 1: Create the hub BR domain.

```
domain [name]
 vrf [name]
  border (create the BR)
   source-interface [interface]
   master [IP address of local MC]
   password [password of hub MC]
```

Example

```
domain iwan
 vrf default
  border
   source-interface Loopback0
   master 10.6.32.251
   password c1sco123
```

Step 2: Add the path names and path-ids to the tunnel interfaces of the hub BR.

```
interface Tunnel [value]
 domain [name] path [name] path-id [number] zero-sla
```

Example

This example is the primary hub BR using Tunnel 10 with MPLS as the provider.

```
interface Tunnel10
 domain iwan path MPLS path-id 1
```

(Optional) This example is the secondary hub BR using Tunnel 11 with INET as the provider and the zero-sla feature. If this is an existing configuration, you shut down the interface, add the zero SLA feature and then, bring the interface back up.

```
interface Tunnel11
 shutdown
 domain iwan path INET path-id 2 zero-sla
 no shutdown
```

Step 3: Verify the border is operational by using the **show domain [name] border status** command.

This example shows the primary hub BR of the IWAN hybrid model with MPLS as the provider. There is only one external WAN interface because the second path is on the secondary hub BR which is reachable via the Tunnel 0 interface at IP address 10.6.32.242.

```
VPN-MPLS-ASR1002X-1#show domain iwan border status
Fri Nov 20 08:10:09.866
--------------------------------------------------------------------
  **** Border Status ****
Instance Status: UP
Present status last updated: 1w0d ago
Loopback: Configured Loopback0 UP (10.6.32.241)
Master: 10.6.32.251
Master version: 2
Connection Status with Master: UP
MC connection info: CONNECTION SUCCESSFUL
Connected for: 3d20h
External Collector: 10.4.48.36  port: 9991
Route-Control: Enabled
Asymmetric Routing: Disabled
Minimum Mask length: 28
Sampling: off
Channel Unreachable Threshold Timer: 4 seconds
Minimum Packet Loss Calculation Threshold: 15 packets
Minimum Byte Loss Calculation Threshold: 1 bytes
Monitor cache usage: 100000 (20%) Auto allocated
Minimum Requirement: Met
External Wan interfaces:
     Name: Tunnel10 Interface Index: 18 SNMP Index: 13 SP: MPLS path-id: 1 Status: UP Zero-SLA: NO Path of Last Resort: Disabled


Auto Tunnel information:


   Name:Tunnel1 if_index: 20
   Virtual Template: Not Configured
   Borders reachable via this tunnel:  10.6.32.242
```

Step 4: Repeat this procedure for each hub BR by using the appropriate path name and path-id.

Procedure 5 Verify PfR domain is operational on the hub MC

The PfR path names and path-ids are automatically discovered at the remote site routers from the configuration entered into the tunnel interfaces at the hub site. The hub MC uses the path names and path-ids to determine where traffic should be sent according to its policies.

Step 1: Verify the domain is operational from the hub MC using the **show domain [name] master status** command.

This example shows the hub MC of the IWAN hybrid model in an operational state. The hub BRs are both connected and using their respective Tunnel interfaces as the exits for the hub location.

```
MC-HY-CSR1000V-1#show domain iwan master status
  *** Domain MC Status ***
Master VRF: Global
  Instance Type:    Hub
  Instance id:      0
  Operational status:  Up
  Configured status:  Up
  Loopback IP Address: 10.6.32.251
  Global Config Last Publish status: Peering Success
  Load Balancing:
   Admin Status: Enabled
   Operational Status: Up
   Enterprise top level prefixes configured: 1
   Max Calculated Utilization Variance: 0%
   Last load balance attempt: never
   Last Reason:  Variance less than 20%
   Total unbalanced bandwidth:
         External links: 0 Kbps  Internet links: 0 Kbps
  External Collector: 10.4.48.36 port: 9991
  Route Control: Enabled
  Transit Site Affinity: Enabled
  Load Sharing: Enabled
  Mitigation mode Aggressive: Disabled
  Policy threshold variance: 20
  Minimum Mask Length: 28
  Syslog TCA suppress timer: 180 seconds
  Traffic-Class Ageout Timer: 5 minutes
  Channel Unreachable Threshold Timer: 4 seconds
  Minimum Packet Loss Calculation Threshold: 15 packets
  Minimum Bytes Loss Calculation Threshold: 1 bytes
  Borders:
    IP address: 10.6.32.241
    Version: 2
    Connection status: CONNECTED (Last Updated 3d20h ago )
    Interfaces configured:
      Name: Tunnel10 | type: external | Service Provider: MPLS path-id:1 | Sta-
tus: UP | Zero-SLA: NO | Path of Last Resort: Disabled
          Number of default Channels: 0
    Tunnel if: Tunnel1
    IP address: 10.6.32.242
    Version: 2
    Connection status: CONNECTED (Last Updated 3d20h ago )
    Interfaces configured:
      Name: Tunnel11 | type: external | Service Provider: INET path-id:2 | Sta-
tus: UP | Zero-SLA: YES | Path of Last Resort: Disabled
          Number of default Channels: 0
    Tunnel if: Tunnel0
```

PROCESS

Configuring PfR for Remote Site Locations

1. Verify IP connectivity to hub MC loopback interface
2. Configure PfR in the primary remote site router
3. Configure PfR in the secondary remote site router
4. Verify PfR traffic classes are controlled

Remote sites are discovered using peering. Each remote site MC peers with the hub MC. The remote site MC advertises local site information and learns information about every other site. Prefixes specific to sites are advertised along with the site-id. The site-prefix to site-id mapping is used in monitoring and optimization. This mapping is also used for creating reports for specific sites.

WAN interfaces at each site are discovered using a special probing mechanism referred to as smart probes. This further reduces provisioning on the remote sites. The WAN interface discovery also creates mapping of the interface to a particular service provider. The mapping is used in monitoring and optimization. It can also be used to draw the WAN topology in an NMS GUI like Cisco Prime Infrastructure or LiveAction.

Procedure 1 Verify IP connectivity to hub MC loopback interface

PfRv3 requires loopback interfaces for the peering traffic between the BR and MC routers. For this design, you put the hub MC loop back interface into the subnet range of the hub location. The following table shows the loopback addresses for the hub MC.

Table 64 Hub MC loopback IP addresses

IWAN design model	Loopback 0 IP address
Hybrid—Hub MC	10.6.32.251
Dual Internet—Hub MC	10.6.32.252

Each remote site must have a route to the hub MC in the EIGRP topology table over each exit path. You can have more than two paths. You can also have two routes and Equal Cost Multiple Paths.

Step 1: Verify that there are at least two available paths to the loopback 0 interface on the hub MC from each remote site router by using the **show ip eigrp topology** command.

This example shows there are two available paths to the hub MC (10.6.32.251) using summarized routes (10.6.0.0/16) from the hub border routers in the IWAN hybrid design model. The internal tags are from the DMVPN hub configurations configured previously in this design guide. The first path is the one shown in the IP routing table because the bandwidth is higher than the feasible successor listed second.

```
RS11-2921# show ip eigrp topology 10.6.0.0 255.255.0.0
EIGRP-IPv4 VR(IWAN-EIGRP) Topology Entry for AS(400)/ID(10.255.241.11) for 
10.6.0.0/16
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 657626453, RIB 
is 5137706
  Descriptor Blocks:
  10.6.34.1 (Tunnel10), from 10.6.34.1, Send flag is 0x0
      Composite metric is (657626453/163840), route is Internal
      Vector metric:
        Minimum bandwidth is 300000 Kbit
        Total delay is 10001250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1400
        Hop count is 1
        Originating router is 10.6.32.241
        Internal tag is 101
  10.6.36.1 (Tunnel11), from 10.6.36.1, Send flag is 0x0
      Composite metric is (1314078720/163840), route is Internal
      Vector metric:
        Minimum bandwidth is 200000 Kbit
        Total delay is 20001250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1400
        Hop count is 1
        Originating router is 10.6.32.242
        Internal tag is 102
```

Procedure 2 Configure PfR in the primary remote site router

Each remote site must have a branch MC and branch BR configured. At dual-router sites it is recommended that you configure the primary router as both an MC and BR and the secondary router as only a BR.

The domain name, VRF, and password must match the hub MC configuration. Use the loopback 0 interface as the source. Configure the hub MC IP address.

Step 1: If you are not on the router console port, turn on terminal monitoring with the **terminal monitor** command from the global command line interface.

```
terminal monitor
```

Step 2: Create the branch MC domain.

```
domain [name]
 vrf [name]
  master branch (create the branch MC)
   source-interface [interface]
   Password [password]
   hub [IP address of hub MC]
```

Example

This example configures the branch MC and points to the IP address of the hub MC in the IWAN hybrid design model.

```
domain iwan
 vrf default
  master branch
   source-interface Loopback0
   password c1sco123
   hub 10.6.32.251
```

Step 3: After approximately two minutes, the console displays an EIGRP SAF message similar to the one below, which indicates the branch MC has created an adjacency with the loopback interface of the hub MC.

```
Dec  5 14:16:00.389: %DUAL-5-NBRCHANGE: EIGRP-SFv4 59501: Neighbor 10.6.32.251
(Loopback0) is up: new adjacency
```

Step 4: Verify the PfR policy from the hub MC has been propagated to the branch MC by using the **show domain [name] master policy** command.

The output from this command should look the same as the output on the hub MC.

Step 5: Enable the BR function.

```
domain [name]
 vrf [name]
  border (create the border)
   source-interface [interface]
   master [local] (local keyword means this router)
   password [password]
```

Example

This example configures the branch BR and points it to the local branch MC, which is running on the same router platform.

```
domain iwan
 vrf default
  border
   source-interface Loopback0
   master local
   password c1sco123
```

Step 6: After approximately thirty seconds, the console displays a line protocol up/down message similar to the one below, which indicates the automatically generated tunnel interface has been created.

```
Dec  5 14:31:26.317: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel1,
changed state to up
```

Step 7: Verify that the branch BR is operational by using the **show domain [name] border status** command.

This example shows the branch BR operational in the IWAN hybrid model and the external WAN interfaces are up.

```
RS11-2921# show domain iwan border status
Fri Dec 05 14:38:45.911
--------------------------------------------------------------------
  ****Border Status****
Instance Status: UP
Present status last updated: 1w4d ago
Loopback: Configured Loopback0 UP (10.255.241.11)
Master: 10.255.241.11
Connection Status with Master: UP
MC connection info: CONNECTION SUCCESSFUL
Connected for: 1w4d
Route-Control: Enabled
Minimum Mask length: 28
```

```
Sampling: off
Minimum Requirement: Met
External Wan interfaces:
     Name: Tunnel10 Interface Index: 15 SNMP Index: 12 SP:MPLS Status: UP
     Name: Tunnel11 Interface Index: 16 SNMP Index: 13 SP:INET Status: UP

Auto Tunnel information:

   Name:Tunnel0 if_index: 25
```

Step 8: Verify that the branch MC is operational by using the **show domain [name] master status** command.

Procedure 3 Configure PfR in the secondary remote site router

Use this procedure only when there is a secondary remote site router. If you have a single router at a remote location, skip this procedure.

PfRv3 requires loopback interfaces for the peering traffic between the BR and MC routers. For this design, you put the hub MC loop back interface into the subnet range of the hub location.

Each remote site must have a route to the hub MC in the EIGRP topology table over each exit path. You can have more than two paths. You can also have two routes and Equal Cost Multiple Paths.

Step 1: Verify that there are at least two available paths to the loopback 0 interface on the hub MC from each remote site router by using the **show ip eigrp topology** command.

This example shows there are two available paths to the hub MC (10.6.32.251) in the IWAN hybrid design model from the secondary router at a remote site. The internal tags are from the DMVPN hub configurations configured previously in this design guide. The first path is through the port-channel2.99 interface, which is the transit network between the primary router and the secondary router in a dual-router configuration.

```
RS12-2911-2# show ip eigrp topology 10.6.32.251 255.255.255.255
EIGRP-IPv4 VR(IWAN-EIGRP) Topology Entry for AS(400)/ID(10.255.242.12) for
10.6.32.251/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 3310960640,
RIB is 25866880
  Descriptor Blocks:
  10.7.16.9 (Port-channel2.99), from 10.7.16.9, Send flag is 0x0
      Composite metric is (3310960640/3310632960), route is Internal
      Vector metric:
        Minimum bandwidth is 20000 Kbit
        Total delay is 50021250000 picoseconds
        Reliability is 255/255
```

```
          Load is 1/255
          Minimum MTU is 1400
          Hop count is 4
          Originating router is 10.6.32.251
          Internal tag is 101
    10.6.36.1 (Tunnel11), from 10.6.36.1, Send flag is 0x0
        Composite metric is (3343400960/1720320), route is Internal
        Vector metric:
          Minimum bandwidth is 10000 Kbit
          Total delay is 50016250000 picoseconds
          Reliability is 255/255
          Load is 1/255
          Minimum MTU is 1400
          Hop count is 3
          Originating router is 10.6.32.251
          Internal tag is 102
```

Step 2: Enable the BR function.

```
domain [name]
 vrf [name]
  border (create the border)
   source-interface [interface]
   master [IP address of branch MC]
   Password [password]
```

Example

This example configures the branch BR and points it to the branch MC, which is running on the primary remote site router.

```
domain iwan
 vrf default
  border
   source-interface Loopback0
   master 10.255.241.12
   password c1sco123
```

Step 3: After approximately thirty seconds, the console displays several EIGRP messages and a line protocol up/down message similar to the ones below. These messages indicate the branch BR has neighbored with the branch MC and automatically generated the tunnel interface from the loopback of the branch BR to the loopback of the branch MC.

```
Dec  5 16:09:11.202: %DUAL-5-NBRCHANGE: EIGRP-SFv4 59501: Neighbor 10.7.18.2
(Port-channel2.64) is up: new adjacency

Dec  5 16:09:11.202: %DUAL-5-NBRCHANGE: EIGRP-SFv4 59501: Neighbor 10.7.19.2
(Port-channel2.69) is up: new adjacency

Dec  5 16:09:11.202: %DUAL-5-NBRCHANGE: EIGRP-SFv4 59501: Neighbor 10.255.241.12
(Loopback0) is up: new adjacency

Dec  5 16:09:11.690: %DUAL-5-NBRCHANGE: EIGRP-SFv4 59501: Neighbor 10.7.16.9
(Port-channel2.99) is up: new adjacency

Dec  5 16:09:12.174: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel0,
changed state to up
```

Step 4: Verify that the branch BR is operational by using the **show domain [name] border status** command.

Step 5: Repeat Procedure 1 through Procedure 3 for each remote site in your network.

Procedure 4 Verify PfR traffic classes are controlled

The final procedure is to verify that the configured and default traffic classes are controlled by the MC at the hub and branch locations.

Step 1: With traffic flowing over the WAN, verify that the PfR traffic classes are controlled in the outbound direction on the hub MC by using the **show domain [name] master traffic-classes summary** command.

This example shows the traffic classes are controlled as signified by the CN in the State column. The default class is load-balanced between the MPLS and INET paths across the network. This example is truncated due to the overall length.

```
PFR-MC-ASR1004-1# show domain iwan master traffic-classes summary


APP - APPLICATION, TC-ID - TRAFFIC-CLASS-ID, APP-ID - APPLICATION-ID

SP - SERVICE PROVIDER, PC = PRIMARY CHANNEL ID,

BC - BACKUP CHANNEL ID, BR - BORDER, EXIT - WAN INTERFACE

UC - UNCONTROLLED, PE - PICK-EXIT, CN - CONTROLLED, UK - UNKNOWN


Dst-Site-Pfx      Dst-Site-Id APP            DSCP    TC-ID APP-ID     State   SP
PC/BC       BR/EXIT


10.7.83.0/24      10.255.241.2N/A            cs3     7632  N/A        CN      INET
10131/NA    10.6.32.242/Tunnel11
```

```
10.255.241.21/32  10.255.241.2N/A           default 12063 N/A       CN      MPLS
3643/10128  10.6.32.241/Tunnel10

10.7.82.0/24      10.255.241.2N/A           default 7656  N/A       CN      INET
10130/NA    10.6.32.242/Tunnel11

10.7.18.13/32     10.255.241.1N/A           default 7646  N/A       CN      INET
10126/NA    10.6.32.242/Tunnel11

10.255.241.12/32  10.255.241.1N/A           default 7651  N/A       CN      INET
10126/NA    10.6.32.242/Tunnel11

10.7.2.0/24       10.255.241.1N/A           default 7636  N/A       CN      MPLS
3638/NA     10.6.32.241/Tunnel10

10.7.2.0/24       10.255.241.1N/A           42      7633  N/A       CN      MPLS
3650/10125  10.6.32.241/Tunnel10

10.255.242.42/32  10.255.241.4N/A           default 7653  N/A       CN      INET
10140/NA    10.6.32.242/Tunnel11

10.7.195.0/24     10.255.241.4N/A           default 15340 N/A       CN      INET
10138/NA    10.6.32.242/Tunnel11

10.7.195.0/24     10.255.241.4N/A           cs3     7640  N/A       CN      MPLS
3654/NA     10.6.32.241/Tunnel10

10.7.18.0/24      10.255.241.1N/A           default 7638  N/A       CN      INET
10126/NA    10.6.32.242/Tunnel11

10.255.242.12/32  10.255.241.1N/A           default 17896 N/A       CN      INET
10126/NA    10.6.32.242/Tunnel11

10.255.241.41/32  10.255.241.4N/A           default 7623  N/A       CN      INET
10138/NA    10.6.32.242/Tunnel11
```

Step 2: With traffic flowing over the WAN, verify that the PfR traffic classes are controlled in the outbound direction on one of the branch MC routers by using the **show domain [name] master traffic-classes dscp** command.

This example shows a video call is taking place from remote site RS11 to the HQ location. The traffic class is controlled, as signified by the Present State row. The INTERACTIVE-VIDEO, with a DSCP of AF41 (34), is in-policy and using the MPLS path. The traffic class has a valid backup channel, which means the INET path is available if the primary path falls out of policy.

```
RS11-2921# show domain iwan master traffic-classes dscp af41


 Dst-Site-Prefix: 10.4.0.0/16         DSCP: af41 [34] Traffic class id:304
  TC Learned:                    00:00:31 ago
  Present State:                 CONTROLLED
  Current Performance Status:    in-policy
  Current Service Provider:      MPLS since 00:00:01 (hold until 88 sec)
  Previous Service Provider:     Unknown
  BW Used:                       416 Kbps
  Present WAN interface:         Tunnel10 in Border 10.255.241.11
```

```
          Present Channel (primary):  312
          Backup Channel:             313
          Destination Site ID:        10.6.32.251
          Class-Sequence in use:      20
          Class Name:                 INTERACTIVE-VIDEO using policy real-time-video
          BW Updated:                 00:00:01 ago
          Reason for Route Change:    Uncontrolled to Controlled Transition
        --------------------------------------------------------------------
```

NetFlow monitoring with Cisco Prime Infrastructure and LiveAction is configured in the "Deploying IWAN Monitoring" section, later in this guide. The LiveAction example below shows the AF41 traffic flow through the MPLS path under normal conditions.

Figure 31 *LiveAction: AF41 traffic flow through the MPLS path on tunnel 10*

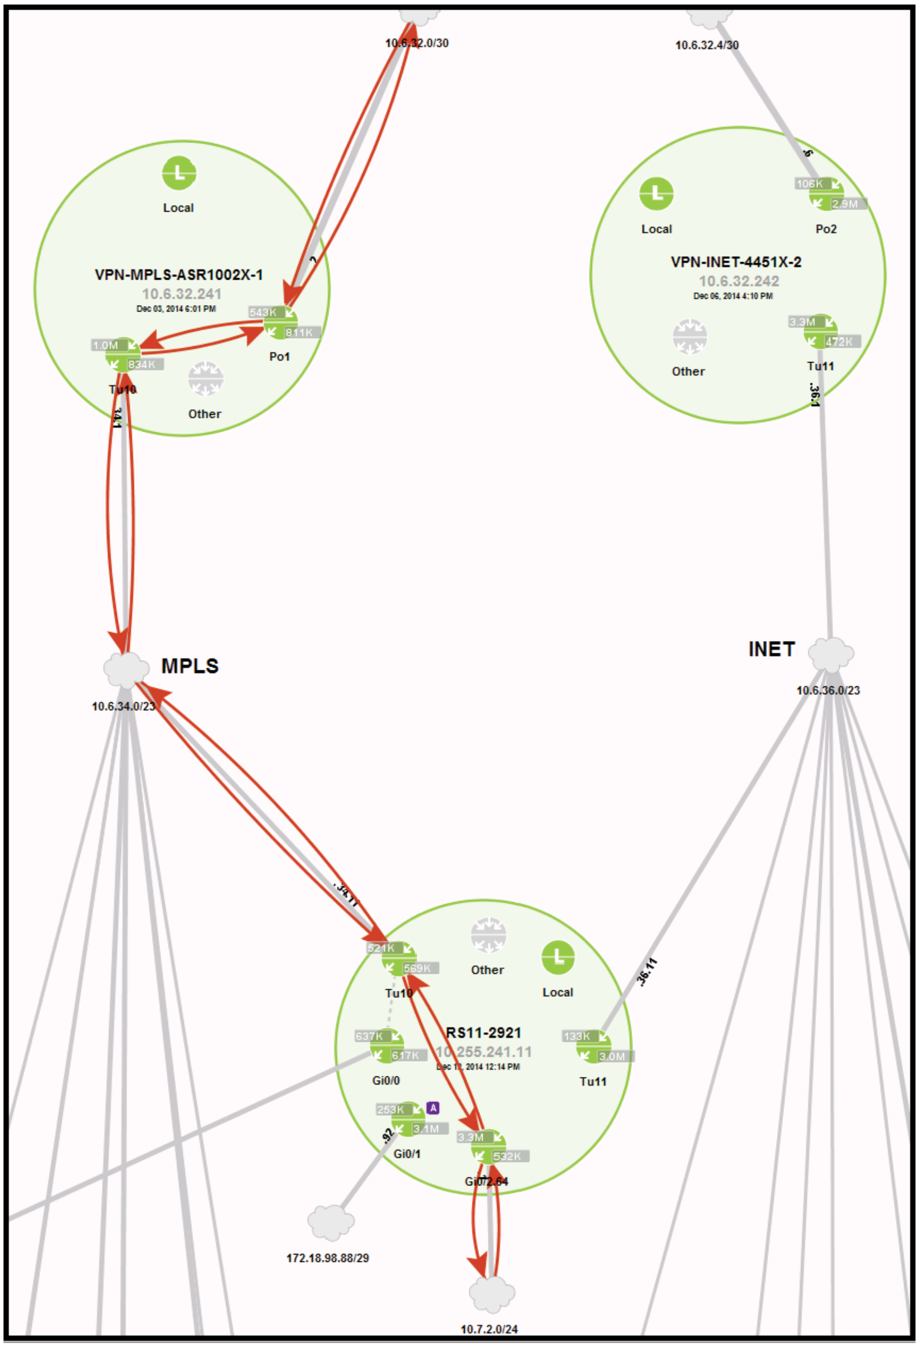

Step 3: After introducing loss into the MPLS path, verify that the protected traffic class is moved to the backup INET path by using the **show domain [name] master traffic-classes dscp** command.

This example shows the INTERACTIVE-VIDEO class, with a DSCP of AF41 (34), using the backup INET path. The branch MC has moved the traffic due to packet loss of greater than 1%. The traffic is considered in-policy because it has already been moved to the INET path where there is no loss occurring.

```
RS11-2921# show domain iwan master traffic-classes dscp af41

 Dst-Site-Prefix: 10.4.0.0/16        DSCP: af41 [34] Traffic class id:303
  TC Learned:                 00:25:40 ago
  Present State:              CONTROLLED
  Current Performance Status: in-policy
  Current Service Provider:   INET since 00:01:09
  Previous Service Provider:  INET for 180 sec
  (A fallback provider. Primary provider will be re-evaluated 00:02:53 later)
  BW Used:                    414 Kbps
  Present WAN interface:      Tunnel11 in Border 10.255.241.11
  Present Channel (primary):  311
  Backup Channel:             310
  Destination Site ID:        10.6.32.251
  Class-Sequence in use:      10
  Class Name:                 INTERACTIVE-VIDEO using policy real-time-video
  BW Updated:                 00:00:10 ago
  Reason for Route Change:    Loss
--------------------------------------------------------------------
```

NetFlow monitoring with Cisco Prime Infrastructure and LiveAction is configured in the “Deploying IWAN Monitoring” section, later in this guide. The LiveAction example below shows the AF41 traffic flow through the INET path after loss has been introduced on the MPLS path.

Figure 32 *LiveAction: AF41 traffic flow through the INET path on tunnel 11 after loss*

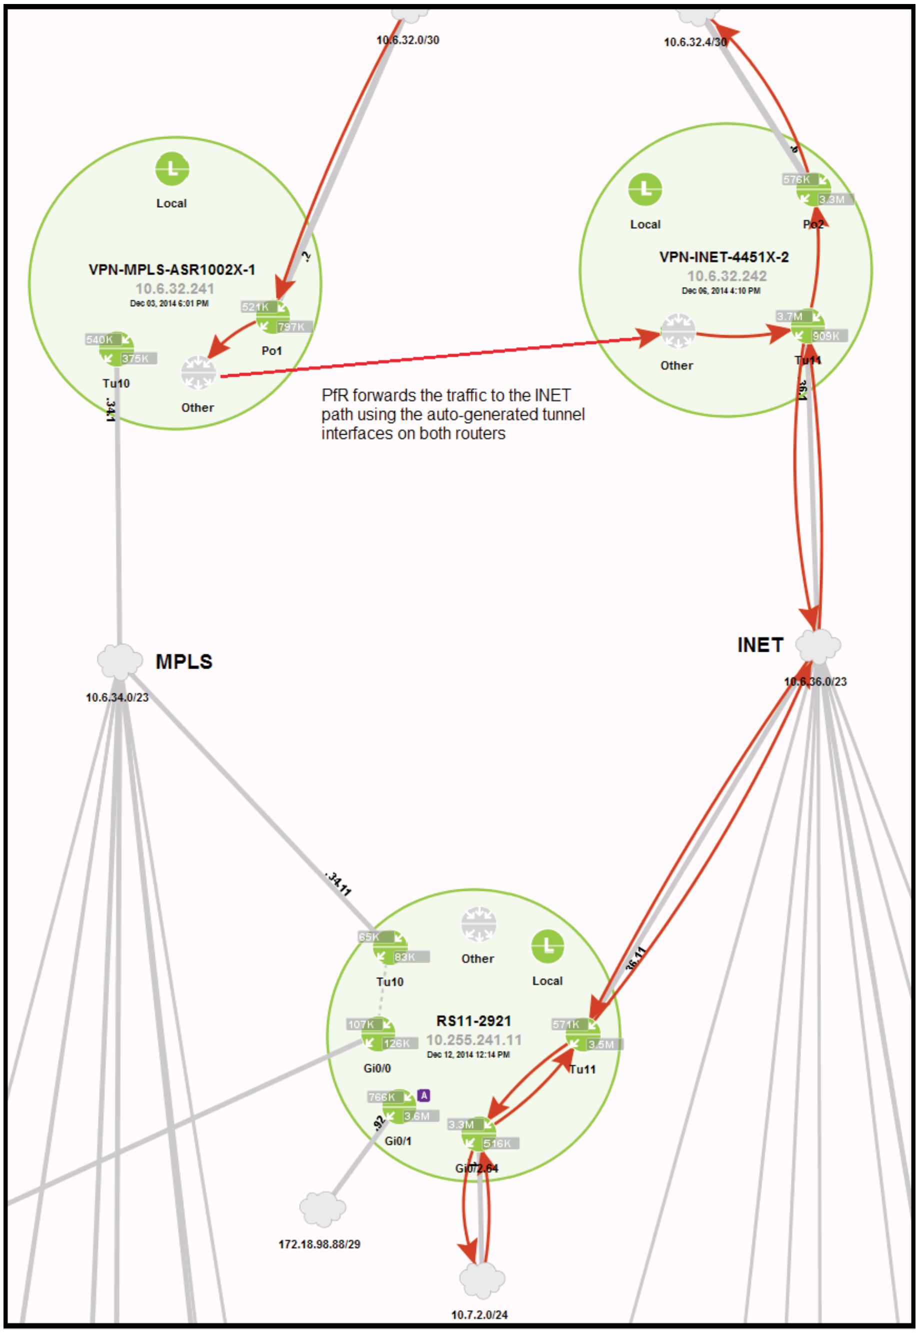

PfR forwards the traffic to the INET path using the auto-generated tunnel interfaces on both routers

PROCESS

Configuring Hub Master Controller High Availability

1. Copy the configuration from existing router to the new router
2. Configure the hub MC HA platform
3. Configure connectivity to the LAN
4. Test the failover from the primary hub MC

Use this optional process if you want to deploy a second hub MC for high availability (HA) using IP Anycast. Skip this process if you do not want to add HA to your hub MC. This concept works with any of the IWAN design models.

For this process, you configure a second hub MC with the same base configuration as the first one. You have to make a few minor changes to allow it to take over when the first hub MC goes offline. The two hub MCs must be kept in sync manually, but the failover will occur automatically within a few minutes depending on the size of your IWAN implementation.

The following diagram shows the hub MC HA and where it fits into the IWAN dual Internet design model.

Figure 33 IWAN dual Internet design model—Hub MC high availability

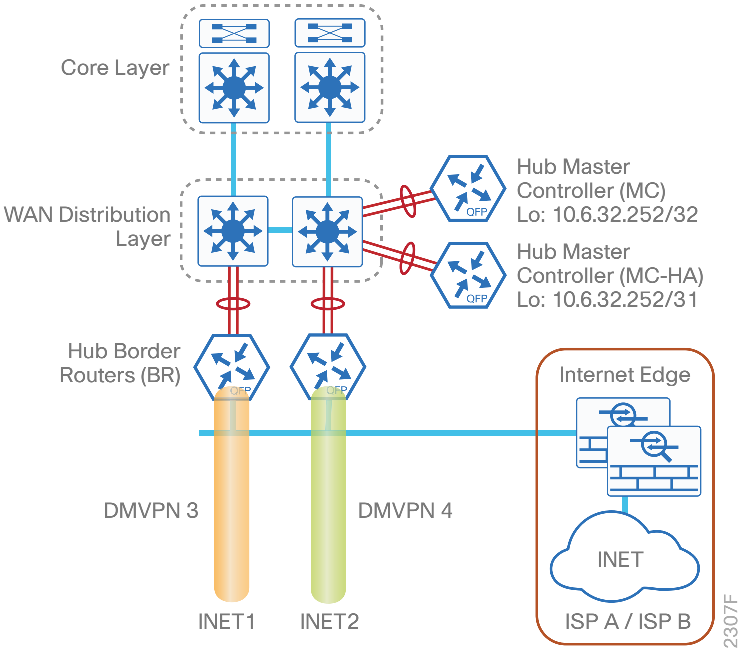

The table below shows the two loopback IP addresses for the pair of hub MCs are the same, except for the network mask. The second hub MC uses a /31 mask, which makes it a less desirable choice by the adjacent router's routing table unless the first hub MC is no longer reachable. The port channel IP addresses are unique.

Table 65 Hub MC IP addresses

IWAN design model	Host name	Loopback IP address	Port-channel IP address
Dual Internet	MC-DI-ASR1004-1	10.6.32.252/32	10.6.32.152/26
Dual Internet	MC-DI-ASR1004-2	10.6.32.252/31	10.6.32.153/26

The other change involves the use of the loopback interface for the various device processes running on the router. To allow access to the processes while both routers are active, you must choose an alternate interface that has a different IP address.

Follow the process "Configuring Hub Master Controller" and the first three procedures of the process "Configuring PfR for Hub Location" using the base PfR information from the first hub MC. Make the required changes from the procedures below in order to enable hub MC HA in the IWAN domain.

Procedure 1 Copy the configuration from existing router to the new router

Optional

If the hardware for the second hub MC is identical to the first, you can use this optional procedure to copy the configuration file from one router to the other as a starting point and then, follow the procedures below. Skip this procedure if you do not want to copy the configuration from an existing router.

Step 1: Copy the running configuration from an existing router to your FTP server.

```
MC-DI-ASR1004-1# copy running-config ftp://cisco:cisco@10.4.48.27
Address or name of remote host [10.4.48.27]?
Destination filename [mc-di-asr1004-1-confg]?
Writing mc-di-asr1004-1-confg !
6175 bytes copied in 0.700 secs (8821 bytes/sec)
```

Step 2: From the console of the new hub MC, copy and paste the configuration into the router before making the changes below.

Procedure 2 Configure the hub MC HA platform

In this procedure, you configure system settings that are unique to the new hub MC.

Step 1: Configure an in-band management interface.

The loopback interface is a logical interface that is always reachable as long as the device is powered on and any IP interface is reachable to the network. However, when you try to access the hub MC HA using the loopback address, the network will direct you to the first hub MC by design. For this reason, you will need access the router using the port-channel interface.

The loopback address is commonly a host address with a 32-bit address mask. In the case of the hub MC HA, you will use the same IP address as the hub MC and a 31-bit address mask.

```
interface Loopback 0
  ip address 10.6.32.252 255.255.255.254
```

Bind the device processes for SNMP, SSH, PIM, TACACS+ and NTP to an interface that is not the loopback, to allow access when both routers are active:

```
snmp-server trap-source Port-channel23
ip ssh source-interface Port-channel23
ip pim register-source Port-channel23
ip tacacs source-interface Port-channel23
ntp source Port-channel23
```

Step 2: Configure IP unicast routing using EIGRP named mode.

EIGRP is configured facing the LAN distribution or core layer. In this design, the port-channel interface and the loopback must be EIGRP interfaces. The loopback may remain a passive interface. The network range must include both interface IP addresses, either in a single network statement or in multiple network statements.

This design uses a best practice of assigning the router ID to a loopback address, but in the case if the hub MC HA, you will need to pick a unique number that does not conflict with other loopback addresses in your network.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface default
   passive-interface
  exit-af-interface
 network 10.6.0.0 0.1.255.255
 eigrp router-id 10.6.32.253
 exit-address-family
```

Procedure 3 Configure connectivity to the LAN

Any links to adjacent distribution layers should be Layer 3 links or Layer 3 EtherChannels. Choose a unique port-channel interface from the LAN switch perspective and an IP address that is different from the first hub MC.

Step 1: Configure a Layer 3 interface.

```
interface Port-channel23
  description IW-WAN-D3750X
  ip address 10.6.32.153 255.255.255.192
  no shutdown
```

Step 2: Configure EtherChannel member interfaces.

Configure the physical interfaces to tie to the logical port-channel by using the **channel-group** command. The number for the port-channel and channel-group must match. Not all router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet0/0/0
 description IW-WAN-D3750X Gig1/0/14

interface GigabitEthernet0/0/1
 description IW-WAN-D3750X Gig2/0/14

interface range GigabitEthernet0/0/0, GigabitEthernet0/0/1
  no ip address
  cdp enable
  channel-group 23
  no shutdown
```

Step 3: Configure the EIGRP interface.

Allow EIGRP to form neighbor relationships across the interface to establish peering adjacencies and exchange route tables. In this step, you configure EIGRP authentication by using the authentication key specified in the previous procedure.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Port-channel23
   no passive-interface
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
 exit-address-family
```

Procedure 4 Test the failover from the primary hub MC

Optional

Use this optional procedure if you want to test the failover to the second hub MC. Skip this procedure if you do not want to test the HA functionality of your hub MC.

During a primary hub MC failure, the remote site will register with the hub MC HA as soon as the branch MC sends the next set of smart probes. The branch MC will continue to use the existing PfR policies until the switchover occurs. If you follow the procedures outlined above, the hub MC HA policy will be identical to the primary hub MC policy.

Step 1: To monitor the progress, log into the second hub MC HA from the console port or using SSH.

Step 2: If you plan to use SSH, turn on console monitoring with **terminal monitor**.

```
MC-DI-ASR1004-2#terminal monitor
```

Step 3: From the console port of primary hub MC, turn off the port-channel interface to the LAN.

```
MC-DI-ASR1004-1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
MC-DI-ASR1004-1(config)#interface Port-channel22
MC-DI-ASR1004-1(config-if)#shut
```

Step 4: From the second hub MC HA, you will see the following messages when the hub BRs and branch MCs register to the backup MC. Depending on the size of the IWAN domain, this step can take several minutes to complete.

```
MC-DI-ASR1004-2#
Nov 19 13:25:26.375: %DUAL-5-NBRCHANGE: EIGRP-IPv4 400: Neighbor 10.6.32.152
(Port-channel23) is down: holding time expired
10.255.243.43 (Loopback0) is up: new adjacency
Nov 19 13:26:37.629: %DUAL-5-NBRCHANGE: EIGRP-SFv4 59501: Neighbor 10.6.32.244
(Loopback0) is up: new adjacency
Nov 19 13:27:00.748: %DUAL-5-NBRCHANGE: EIGRP-SFv4 59501: Neighbor 10.255.243.13
(Loopback0) is up: new adjacency
Nov 19 13:27:04.580: %DUAL-5-NBRCHANGE: EIGRP-SFv4 59501: Neighbor 10.6.32.243
(Loopback0) is up: new adjacency
```

```
Nov 19 13:27:20.402: %DUAL-5-NBRCHANGE: EIGRP-SFv4 59501: Neighbor 10.255.243.44
(Loopback0) is up: new adjacency

Nov 19 13:27:23.259: %DUAL-5-NBRCHANGE: EIGRP-SFv4 59501: Neighbor 10.255.243.14
(Loopback0) is up: new adjacency
```

Step 5: After the messages stop, confirm that the second hub MC is acting as the hub MC with **show domain [domain name] master status**.

```
MC-DI-ASR1004-2#show domain iwan2 master status
  *** Domain MC Status ***
 Master VRF: Global
  Instance Type:    Hub
  Instance id:      0
  Operational status:  Up
  Configured status:  Up
  Loopback IP Address: 10.6.32.252
  Global Config Last Publish status: Peering Success
  Load Balancing:
   Admin Status: Enabled
   Operational Status: Up
   Enterprise top level prefixes configured: 1
   Max Calculated Utilization Variance: 0%
   Last load balance attempt: never
   Last Reason:  Variance less than 20%
   Total unbalanced bandwidth:
         External links: 0 Kbps  Internet links: 0 Kbps
  External Collector: 10.4.48.36 port: 9991
  Route Control: Enabled
  Transit Site Affinity: Enabled
  Load Sharing: Enabled
  Mitigation mode Aggressive: Disabled
  Policy threshold variance: 20
  Minimum Mask Length: 28
  Syslog TCA suppress timer: 180 seconds
  Traffic-Class Age out Timer: 5 minutes
  Channel Unreachable Threshold Timer: 4 seconds
  Minimum Packet Loss Calculation Threshold: 15 packets
  Minimum Bytes Loss Calculation Threshold: 1 bytes
  Borders:
    IP address: 10.6.32.243
    Version: 2
    Connection status: CONNECTED (Last Updated 00:00:54 ago )
    Interfaces configured:
      Name: Tunnel20 | type: external | Service Provider: INET1 path-id:1 | 
Status: UP | Zero-SLA: NO | Path of Last Resort: Disabled
          Number of default Channels: 0
    Tunnel if: Tunnel0
    IP address: 10.6.32.244
    Version: 2
    Connection status: CONNECTED (Last Updated 00:00:52 ago )
    Interfaces configured:
      Name: Tunnel21 | type: external | Service Provider: INET2 path-id:2 | 
Status: UP | Zero-SLA: NO | Path of Last Resort: Disabled
          Number of default Channels: 0
    Tunnel if: Tunnel0
```

After you have verified that the second hub MC is operational, log into the primary hub MC to bring it back online.

Step 6: From the console port of the primary hub MC, turn on the port-channel interface to the LAN.

```
MC-DI-ASR1004-1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
MC-DI-ASR1004-1(config)#interface Port-channel22
MC-DI-ASR1004-1(config-if)#no shut
```

PROCESS

Configuring Hub Border Router Scalability

1. Copy the configuration from existing router to the new router
2. Configure the hub BR platform
3. Configure connectivity to the LAN
4. Connect to the Internet
5. Configure the mGRE tunnel
6. Configure EIGRP
7. Configure network address translation on the firewall
8. Configure PfR domain in the hub BR
9. Configure remote sites for additional hub BRs

Use this optional process if you want to deploy additional hub BRs at the same location for horizontal scaling. Skip this process if you do not want to horizontally scale your hub BRs. This concept works with any of the IWAN design models.

This type of configuration offers the following benefits:

- Distribute traffic across multiple hub BRs on a single DMVPN to utilize all WAN and router capacity
- Convergence across hub BRs should only occur when all exits in a hub BR fail or reach their maximum bandwidth limits
- If the current exit to a remote site fails, converge to an alternate exit on the same (DMVPN1) network or converge to the alternate (DMVPN2) network

The following diagram shows two additional hub BRs and where they fit into the IWAN dual Internet design model.

Figure 34 *IWAN dual Internet design model—Hub BR scalability*

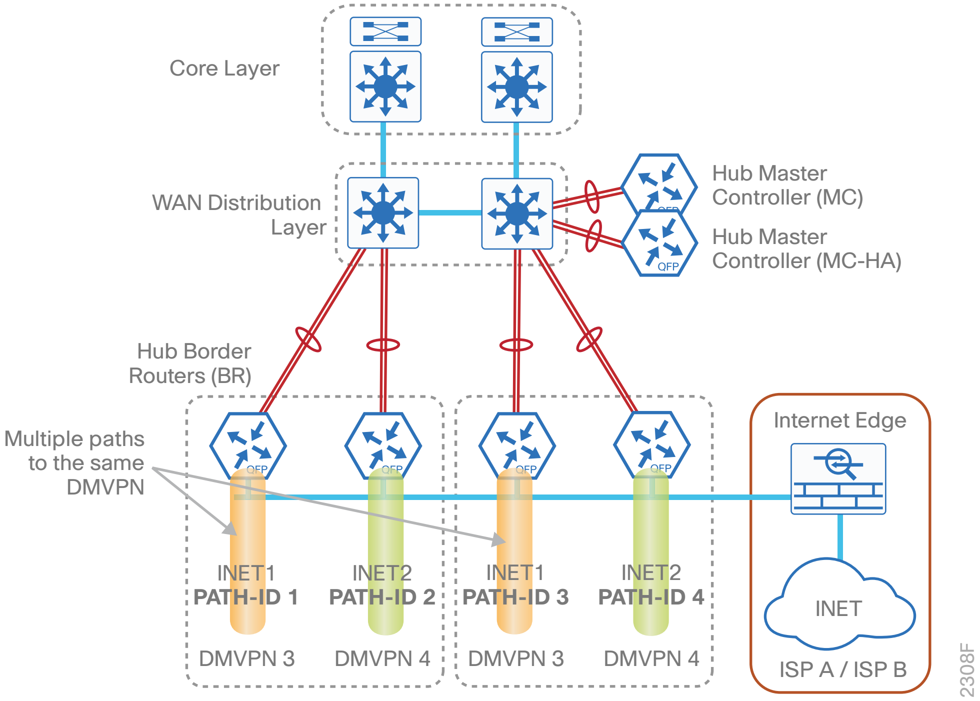

For this process, you configure two additional hub BRs with base configurations similar to the existing hub BRs.' You have to make changes to the base configurations and the remote site routers in order to take advantage of the new hub BRs

The additional routers have unique path information, IP addresses, and port-channel assignments, but the rest of the configurations are the same.

Table 66 *Hub BR path and IP addresses*

Host name	Path	Path ID	Loopback IP address	Port-channel IP address	Internet DMZ IP address
VPN-INET-ASR1002X-3	INET1	1	10.6.32.243/32	10.6.32.18/30	192.168.146.20/24
VPN-INET-ASR1002X-4	INET2	2	10.6.32.244/32	10.6.32.22/30	192.168.146.21/24
VPN-INET-ASR1002X-5	INET1	3	10.6.32.245/32	10.6.32.26/30	192.168.146.22/24
VPN-INET-ASR1002X-6	INET2	4	10.6.32.246/32	10.6.32.30/30	192.168.146.23/24

Follow the process “Configuring DMVPN Hub Router” using the base PfR information from the first two hub BRs. Make the required changes from the procedures below to horizontally scale your IWAN domain.

Procedure 1 Copy the configuration from existing router to the new router

Optional

If the hardware for the corresponding hub BR is identical to the first, you can use this optional procedure to copy the configuration file from one router to the other as a starting point and then, follow the procedures below. Skip this procedure if you do not want to copy the configuration from an existing router.

Step 1: Copy the running configuration from an existing router to your FTP server.

```
VPN-INET-ASR1002X-3# copy running-config ftp://cisco:cisco@10.4.48.27
Address or name of remote host [10.4.48.27]?
Destination filename [vpn-inet-asr1002x-3-confg]?
Writing vpn-inet-asr1002x-3-confg !
13228 bytes copied in 0.7500 secs (9921 bytes/sec)
```

Step 2: From the console of the new hub BR, copy and paste the configuration into the router before making the changes below.

Procedure 2 Configure the hub BR platform

In this procedure, you configure system settings that are unique to the new hub BR.

Step 1: Configure the device host name to make it easy to identify the device.

```
hostname VPN-INET-ASR1002X-5
```

Step 2: Configure an in-band management interface.

The loopback interface is a logical interface that is always reachable as long as the device is powered on and any IP interface is reachable to the network.

The loopback address is commonly a host address with a 32-bit address mask.

```
interface Loopback 0
  ip address 10.6.32.245 255.255.255.255
```

Step 3: Configure IP unicast routing using EIGRP named mode.

In this design, the port-channel interface and the loopback must be EIGRP interfaces. The loopback may remain a passive interface. The network range must include both interface IP addresses, either in a single network statement or in multiple network statements.

This design uses a best practice of assigning the router ID to a loopback address.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface default
   passive-interface
  exit-af-interface
 network 10.6.0.0 0.1.255.255
 eigrp router-id 10.6.32.245
 exit-address-family
```

Procedure 3 Configure connectivity to the LAN

Any links to adjacent distribution layers should be Layer 3 links or Layer 3 EtherChannels. Choose a unique port-channel interface from the LAN switch perspective and an IP address that is different from the other hub BRs.

Step 1: Configure a Layer 3 interface.

At the hub location where there are multiple border routers, the interface throughput delay setting should be set to influence the routing protocol path preference. Set the internal LAN path to 250000 usec. The delay command is entered in 10 usec units.

```
interface Port-channel5
  description IW-WAN-D3750X
  ip address 10.6.32.26 255.255.255.252
  ip pim sparse-mode
  delay 25000
  no shutdown
```

Step 2: Configure EtherChannel member interfaces.

Configure the physical interfaces to tie to the logical port-channel by using the **channel-group** command. The number for the port-channel and channel-group must match. Not all router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet0/0/0
 description IW-WAN-D3750X Gig1/0/5

interface GigabitEthernet0/0/1
 description IW-WAN-D3750X Gig2/0/5

interface range GigabitEthernet0/0/0, GigabitEthernet0/0/1
  no ip address
  cdp enable
  channel-group 5
  no shutdown
```

Step 3: Configure the EIGRP interface.

Allow EIGRP to form neighbor relationships across the interface to establish peering adjacencies and exchange route tables. In this step, you configure EIGRP authentication by using the authentication key specified in the previous procedure.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Port-channel5
   no passive-interface
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
 exit-address-family
```

Procedure 4 Connect to the Internet

The DMVPN hub is connected through a Cisco ASA 5500 using a DMZ interface specifically created and configured for a VPN termination router.

The IP address that you use for the Internet-facing interface of the DMVPN hub router must be an Internet-routable address. There are two possible methods for accomplishing this task:

- Assign a routable IP address directly to the router.
- Assign a non-routable RFC-1918 address directly to the router and use a static NAT on the Cisco ASA 5500 in order to translate the router IP address to a routable IP address.

This design assumes that the Cisco ASA 5500 is configured for static NAT for the DMVPN hub router.

Step 1: Enable the interface, select the VRF, and assign the IP address.

```
interface GigabitEthernet0/0/3
 vrf forwarding IWAN-TRANSPORT-3
 ip address 192.168.146.22 255.255.255.0
 no shutdown
```

Procedure 5 Configure the mGRE tunnel

The parameters in the table below are used in this procedure. Choose the row that represents the hub BR that you are configuring. This procedure applies to the scale hub BR in the IWAN dual Internet design model.

Table 67 DMVPN tunnel parameters

Hostname	Tunnel type	Tunnel number	Tunnel IP address
VPN-INET-ASR1002X-5	INET1	20	10.6.38.2/23
VPN-INET-ASR1002X-6	INET2	21	10.6.40.2/23

Step 1: Configure the basic interface settings.

The tunnel number is arbitrary, but it is best to begin tunnel numbering at 10 or above, because other features deployed in this design may also require tunnels and they may select lower numbers by default.

```
interface Tunnel20
 ip address 10.6.38.2 255.255.254.0
```

Step 2: Configure NHRP.

Hub BRs require an additional configuration statement in order to create an EIGRP neighbor adjacency with the other hub BR. This statement includes the NBMA definition for the DMVPN hub router tunnel endpoint. EIGRP relies on a multicast transport. Hub BRs require the NHRP multicast keyword in this statement.

The value used for the NHS is the mGRE tunnel address for the DMVPN hub router. The NBMA entry must be set to the hub router's DMZ IP address because both of the hub routers are behind the firewall. This design uses the values shown in the following table.

Table 68 NHRP parameters

Hostname	Tunnel type	Tunnel number	Tunnel IP address	DMZ IP address
VPN-INET-ASR1002X-3	INET1	20	10.6.38.1	192.168.146.20
VPN-INET-ASR1002X-4	INET2	21	10.6.40.1	192.168.146.21
VPN-INET-ASR1002X-5	INET1	20	10.6.38.2	192.168.146.22
VPN-INET-ASR1002X-6	INET2	21	10.6.40.2	192.168.146.23

The two corresponding hub BRs must point at each other to allow an EIGRP neighbor adjacency to be formed.

Example: VPN-INET-ASR1002X-3

```
interface Tunnel20
 ip nhrp nhs 10.6.38.2 nbma 192.168.146.22 multicast
```

Example: VPN-INET-ASR1002X-5

```
interface Tunnel20
 ip nhrp nhs 10.6.38.1 nbma 192.168.146.20 multicast
```

Procedure 6 Configure EIGRP

Step 1: Tag and filter the routes.

This design uses a single EIGRP autonomous system for the WAN and all of the WAN remote sites. Every remote site is dual-connected for resiliency. However, due to the multiple paths that exist within this topology, you must try to avoid routing loops and to prevent remote sites from becoming transit sites if WAN failures were to occur.

In this design, there are different IP subnets for each DMVPN network, and the EIGRP tags are clearly defined to help with readability and troubleshooting. When a design uses more than one data center, additional tags are required in order to identify the different DMVPN hub router locations.

The following logic is used to control the routing.

- Each DMVPN network will have an EIGRP route tag to prevent routes from being re-advertised over the other DMVPN networks at the remote sites.
- All prefixes that are advertised towards the WAN are uniquely tagged.
- All DMVPN learned WAN prefixes, except those that originate locally from a hub, are advertised towards the LAN and tagged.

Outbound distribute-lists are used to set tags on the DMVPN hub routers towards the WAN and LAN. The remote-site routers use the tags set towards the WAN in order to protect against becoming transit sites.

The following tables show specific route tags in use.

Table 69 Route tag information for optional DMVPN IWAN dual Internet hub routers

DMVPN hub	DMVPN prefix (tag)	Tag tunnel	Tag LAN
VPN-INET-ASR1002X-5	201 (INET1)	201 (All routes)	201 (WAN routes)
VPN-INET-ASR1002X-6	202 (INET2)	202 (All routes)	202 (WAN routes)

The following examples show both of the new hub BRs in the IWAN dual Internet design model.

Example: VPN-INET-ASR1002X-5

```
route-map SET-TAG-ALL permit 10
 description Tag all routes advertised through the tunnel
 set tag 201

! All INET1 tunnel interfaces are in this IP address range
ip access-list standard DMVPN-3-SPOKES
 permit 10.6.38.0 0.0.1.255

route-map SET-TAG-DMVPN-3 permit 10
 description Tag routes sourced from DMVPN-3
 match ip route-source DMVPN-3-SPOKES
 set tag 201

route-map SET-TAG-DMVPN-3 permit 100
 description Advertise all other routes with no tag

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   topology base
   distribute-list route-map SET-TAG-DMVPN-3 out Port-channel5
   distribute-list route-map SET-TAG-ALL out Tunnel20
```

Example: VPN-INET-ASR1002X-6

```
route-map SET-TAG-ALL permit 10
 description Tag all routes advertised through the tunnel
 set tag 202

! All INET2 tunnel interfaces are in this IP address range
ip access-list standard DMVPN-4-SPOKES
 permit 10.6.40.0 0.0.1.255

route-map SET-TAG-DMVPN-4 permit 10
 description Tag routes sourced from DMVPN-4
 match ip route-source DMVPN-4-SPOKES
 set tag 202

route-map SET-TAG-DMVPN-4 permit 100
 description Advertise all other routes with no tag

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   topology base
   distribute-list route-map SET-TAG-DMVPN-4 out Port-channel6
   distribute-list route-map SET-TAG-ALL out Tunnel21
```

Procedure 7 Configure network address translation on the firewall

You have to add the new hub BRs to your existing firewall configuration for network address translation.

The DMZ network uses private network (RFC 1918) addressing that is not Internet-routable, so the firewall must translate the DMZ address of the DMVPN hub router to an outside public address.

The example DMZ address to public IP address mapping is shown in the following table.

Table 70 *DMVPN NAT address mapping*

Hostname	DMVPN hub router DMZ address	DMVPN hub router public address (externally routable after NAT)
VPN-INET-ASR1002X-5	192.168.146.22	172.16.140.12 (ISP-A)
VPN-INET-ASR1002X-6	192.168.146.23	172.17.140.12 (ISP-B)

First, to simplify the configuration of the security policy, you create the External DMZ network objects that are used in the firewall policies.

Table 71 *External DMZ firewall network objects*

Network object name	Object type	IP address	Description
outside-dmvpn-5-ISPa	Host	172.16.140.12	DMVPN hub router 5 on ISP A (outside)
outside-dmvpn-6-ISPb	Host	172.17.140.12	DMVPN hub router 6 on ISP B (outside)

Step 1: Navigate to **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups**.

Step 2: Click **Add** > **Network Object**.

The Add Network Object dialog box appears.

Step 3: In the **Name** box, enter the name. (Example: outside-dmvpn-5-ISPa)

Step 4: In the **Type** list, choose **Host** or **Network**. (Example: Host)

Step 5: In the **IP Address** box, enter the address. (Example: 172.16.140.12)

Step 6: In the **Description** box, enter a useful description, and then click **OK**. (Example: DMVPN hub router 5 on ISP A)

Step 7: Repeat Step 2 through Step 6 for each object listed in the table above. If an object already exists, then skip to the next object listed in the table.

Step 8: After adding all of the objects listed, on the Network Objects/Groups pane, click **Apply**.

Next, you add a network object for the private DMZ address of the DMVPN hub router.

Table 72 Private DMZ firewall network objects

Network object name	Object type	IP address	Description
dmz-dmvpn-5	Host	192.168.146.22	DMVPN hub router 5 on vpn-dmz
dmz-dmvpn-6	Host	192.168.146.23	DMVPN hub router 6 on vpn-dmz

Step 9: Navigate to **Configuration > Firewall > Objects > Network Objects/Groups**.

Step 10: Click **Add > Network Object**.

The Add Network Object dialog box appears.

Step 11: In the **Name** box, enter the name. (Example: dmz-dmvpn-5)

Step 12: In the **Type** list, choose **Host** or **Network**. (Example: Host)

Step 13: In the **IP Address** box, enter the address. (Example: 192.168.146.22)

Step 14: In the **Description** box, enter a useful description, and then click **OK**. (Example: DMVPN hub router 5 on vpn-dmz)

Step 15: Click the two down arrows. The NAT pane expands.

Step 16: Select **Add Automatic Address Translation Rules**.

Step 17: In the **Translated Address** list, choose the network object created previously. (Example: outside-dmvpn-5-ISPa)

Step 18: Select **Use one-to-one address translation**, and then click **OK**.

Step 19: Repeat Step 10 through Step 18 for each object listed in the table above. If an object already exists, then skip to the next object listed in the table.

Step 20: After adding all of the objects listed, on the Network Objects/Groups pane, click **Apply**.

Procedure 8 Configure PfR domain in the hub BR

The additional hub BRs are also the DMVPN hub WAN aggregation routers for the network. The PfRv3 configurations for standalone BRs are much simpler because they dynamically learn their policy information from the hub MC. The hub BR routers are also used to advertise the path names and path-ids specified in the hub MC configuration.

Step 1: Create the hub BR domain.

```
domain [name]
 vrf [name]
  border (create the BR)
   source-interface [interface]
   master [IP address of local MC]
   password [password of hub MC]
```

Example

```
domain iwan2
 vrf default
  border
   source-interface Loopback0
   master 10.6.32.252
   password c1sco123
```

Step 2: Add the path names and path-ids to the tunnel interfaces of the hub BR.

```
interface Tunnel [value]
 domain [name] path [name] path-id [number]
```

Example

This example is the additional hub BR using Tunnel 20 with INET1 as the provider.

```
interface Tunnel20
 domain iwan2 path INET1 path-id 3
```

This example is the additional hub BR using Tunnel 21 with INET2 as the provider.

```
interface Tunnel21
 domain iwan2 path INET2 path-id 4
```

Step 3: Verify the border is operational by using the **show domain [name] border status** command.

Step 4: Repeat this procedure for each additional hub BR by using the appropriate path name and path-id.

Procedure 9 Configure remote sites for additional hub BRs

An additional NHRP command has to be added to the tunnel interfaces of remote site BRs for them to begin using the new hub BRs.

Table 73 NHRP parameters

Hostname	Tunnel number	Tunnel IP address	Public IP address
VPN-INET-ASR1002X-5	20	10.6.38.2	172.16.140.12 (ISP A)
VPN-INET-ASR1002X-6	21	10.6.40.2	172.17.140.12 (ISP B)

Step 1: Configure NHRP.

The DMVPN hub router is the NHRP server for all of the spokes. Remote routers use NHRP in order to determine the tunnel destinations for peers attached to the mGRE tunnel.

The spoke router requires an additional configuration statement in order to define the NHRP server. This statement includes the NBMA definition for the DMVPN hub router tunnel endpoint. EIGRP relies on a multicast transport. Spoke routers require the NHRP multicast keyword in this statement.

The value used for the NHS is the mGRE tunnel address for the DMVPN hub router. The NBMA entry must be set to either the MPLS DMVPN hub router's actual public address or the outside NAT value of the DMVPN hub, as configured on the Cisco ASA 5500. This design uses the values shown in the table above.

Example: Single Router Dual INET—RS13-2911

```
interface Tunnel20
 ip nhrp nhs 10.6.38.2 nbma 172.16.140.12 multicast

interface Tunnel21
 ip nhrp nhs 10.6.40.2 nbma 172.17.140.12 multicast
```

Step 2: Confirm that the new hub BRs are reachable with **show ip eigrp neighbors**.

```
RS13-2911#show ip eigrp neighbors
EIGRP-IPv4 VR(IWAN-EIGRP) Address-Family Neighbors for AS(400)
H   Address         Interface      Hold Uptime    SRTT    RTO  Q   Seq
                                   (sec)          (ms)        Cnt  Num
2   10.6.38.2       Tu20             42 1d01h       1    100  0   574
3   10.6.38.1       Tu20             58 1d01h       1    100  0   631
1   10.6.40.2       Tu21             59 1d02h       1    100  0   646
0   10.6.40.1       Tu21             55 1d02h       1    100  0   804
```

Step 3: Repeat this procedure for each remote site that will use the new hub BRs.

Deploying a Second Data Center Location

Use this optional section if you want to deploy a second data center location as a transit site for geographic redundancy and scalability. Skip this section if you do not want to add a transit site to your network. This concept works with any of the IWAN design models.

This type of configuration offers the following benefits:

- Data centers are reachable across the WAN core for each transit site.
- Remote sites can access any data center across either hub.
- Data centers can reach any remote site across any of the transit sites.
- Multiple hub BRs per DMVPN per site may be required for horizontal scaling, as noted in the previous process.

This design introduces the concept of a transit master controller and transit BRs.

- **Transit Master Controller**—The Transit MC is the MC at the transit-site. There is no policy configuration on this device. It receives policy from the Hub MC. This device acts as MC for that site for making path optimization decision. The configuration includes the IP address of the hub MC.
- **Transit Border Router**—This is a BR at the transit MC site. This is the device where WAN interfaces terminate. There can only be one WAN interface on the device. There can be one or more transit BRs. On the transit BRs, PfRv3 must be configured with:
    - The address of the transit MC.
    - The path name on external interfaces.
    - The path ID on external interfaces.

The following diagram shows the transit MC with two additional transit BRs and where they fit into the IWAN hybrid design model.

Figure 35 *IWAN hybrid design model—Second data center as a transit site*

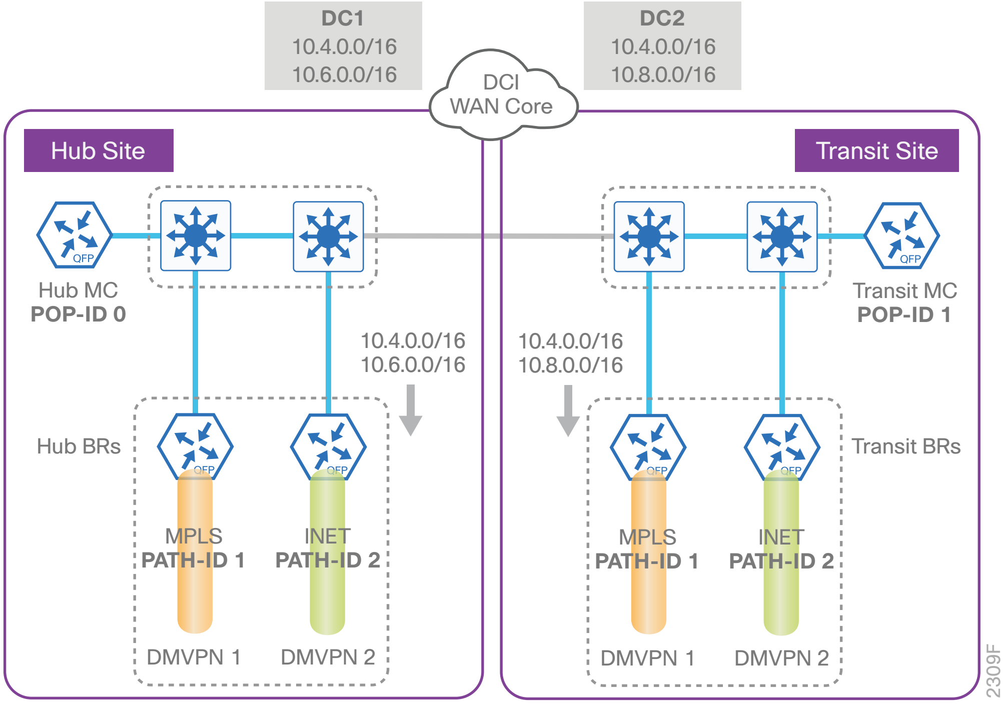

With the IOS release used for this guide, data center affinity is enabled by default. It is applicable for both path preference and load balancing. There is no CLI change required and PfR will use the primary data center as its preference for all traffic.

If the MPLS path is primary and INET path is secondary in your design, the path preference will be as follows:

- Path #1 to 10.4.0.0/16 is MPLS path to DC#1
- Path #2 to 10.4.0.0/16 is INET path to DC#1
- Path #3 to 10.4.0.0/16 is MPLS path to DC#2
- Path #4 to 10.4.0.0/16 is INET path to DC#2

If you want the path preference to be the MPLS path as primary and INET path as fallback across data centers, there is a domain **transit-site-affinity** command to disable data center affinity.

```
domain iwan
 vrf default
  master hub
   advanced
    no transit-site-affinity
```

If **no transit-site-affinity** is enabled, the failover order for the example given above would be as follows:

- Path #1 to 10.4.0.0/16 is MPLS path to DC#1
- Path #2 to 10.4.0.0/16 is MPLS path to DC#2
- Path #3 to 10.4.0.0/16 is INET path to DC#1
- Path #4 to 10.4.0.0/16 is INET path to DC#2

PROCESS

Configuring Transit Border Routers

1. Copy the configuration from existing router to the new router
2. Configure the transit BR platform
3. Configure connectivity to the LAN
4. Connect to the MPLS WAN or Internet
5. Configure the mGRE tunnel
6. Configure EIGRP
7. Configure network address translation on the firewall

For this process, you configure two transit site BRs with similar base configurations as the existing hub BRs. You have to make changes to the base configurations and the remote site routers to take advantage of the new transit site location.

The transit site BR routers have unique IP addresses and port-channel assignments, but the rest of the configuration items are the same.

Table 74 Hub and transit BR path and IP addresses

Host name	Path	Path ID	Loopback IP address	Port-channel IP address	MPLS/Internet DMZ IP address
VPN-MPLS-ASR1002X-1	MPLS	1	10.6.32.241/32	10.6.32.2/30	192.168.6.1/24
VPN-INET-4451X-2	INET	2	10.6.32.242/32	10.6.32.6/30	192.168.146.10/24
VPN-MPLS-ASR1002X-T1	MPLS	1	10.8.32.241/32	10.8.32.2/30	192.168.6.41/24
VPN-INET-ASR1002X-T2	INET	2	10.8.32.242/32	10.8.32.6/30	192.168.146.11/24

Follow the process “Configuring DMVPN Hub Router,” using the base PfR information from the first two hub BRs. Make the required changes from the procedures below to add a transit site to your IWAN domain.

Procedure 1 Copy the configuration from existing router to the new router

Optional

If the hardware for the corresponding transit BR is identical to the hub BR, you can use this optional procedure to copy the configuration file from one router to the other as a starting point and then, follow the procedures below. Skip this procedure if you do not want to copy the configuration from an existing router.

Step 1: Copy the running configuration from an existing router to your FTP server.

```
VPN-MPLS-ASR1002X-1# copy running-config ftp://cisco:cisco@10.4.48.27
Address or name of remote host [10.4.48.27]?
Destination filename [vpn-mpls-asr1002x-1-confg]?
Writing vpn-mpls-asr1002x-1-confg !
15884 bytes copied in 0.800 secs (12707 bytes/sec)
```

Step 2: From the console of the new transit BR, copy and paste the configuration into the router before making the changes below.

Procedure 2 Configure the transit BR platform

In this procedure, you configure system settings that are unique to the transit BR.

Step 1: Configure the device host name to make it easy to identify the device.

```
hostname VPN-MPLS-ASR1002X-T1
```

Step 2: Configure an in-band management interface.

The loopback interface is a logical interface that is always reachable as long as the device is powered on and any IP interface is reachable to the network.

The loopback address is commonly a host address with a 32-bit address mask.

```
interface Loopback 0
  ip address 10.8.32.241 255.255.255.255
```

Step 3: Configure IP unicast routing using EIGRP named mode.

In this design, the tunnel, port-channel and loopback must be EIGRP interfaces. The loopback may remain a passive interface. The network range must include all interface IP addresses, either in a single network statement or in multiple network statements.

This design uses a best practice of assigning the router ID to a loopback address.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface default
   passive-interface
  exit-af-interface
 network 10.6.0.0 0.1.255.255
 network 10.8.0.0 0.1.255.255
 eigrp router-id 10.8.32.241
 exit-address-family
```

Procedure 3 Configure connectivity to the LAN

Any links to adjacent distribution layers should be Layer 3 links or Layer 3 EtherChannels. Choose a unique port-channel interface from the LAN switch perspective.

Step 1: Configure a Layer 3 interface.

At the transit location where there are multiple border routers, the interface throughput delay setting should be set to influence the routing protocol path preference. Set the internal LAN path to 250000 usec. The delay command is entered in 10 usec units.

```
interface Port-channel1
  description IWAN-D3750X-T
  ip address 10.8.32.2 255.255.255.252
  ip pim sparse-mode
  delay 25000
  no shutdown
```

Step 2: Configure EtherChannel member interfaces.

Configure the physical interfaces to tie to the logical port-channel by using the **channel-group** command. The number for the port-channel and channel-group must match. Not all router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet0/0/0
 description IWAN-D3750X-T Gig1/0/1

interface GigabitEthernet0/0/1
 description IWAN-D3750X-T Gig2/0/1

interface range GigabitEthernet0/0/0, GigabitEthernet0/0/1
```

```
          no ip address
          cdp enable
          channel-group 1
          no shutdown
```

Step 3: Configure the EIGRP interface.

Allow EIGRP to form neighbor relationships across the interface to establish peering adjacencies and exchange route tables. In this step, you configure EIGRP authentication by using the authentication key specified in the previous procedure.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Port-channel1
   no passive-interface
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
 exit-address-family
```

Procedure 4 Connect to the MPLS WAN or Internet

Each IWAN DMVPN hub requires a connection to the WAN transport, which for the hybrid model is either MPLS or Internet.

If you are using MPLS in this design, the DMVPN hub is connected to the service provider’s MPLS PE router. The IP addressing used between IWAN CE and MPLS PE routers must be negotiated with your MPLS carrier.

If you are using the Internet in this design, the DMVPN hub is connected through a Cisco ASA 5500 using a DMZ interface specifically created and configured for a VPN termination router.

The IP address that you use for the Internet-facing interface of the DMVPN hub router must be an Internet-routable address. There are two possible methods for accomplishing this task:

- Assign a routable IP address directly to the router.
- Assign a non-routable RFC-1918 address directly to the router and use a static NAT on the Cisco ASA 5500 to translate the router IP address to a routable IP address.

This design assumes that the Cisco ASA 5500 is configured for static NAT for the DMVPN hub router.

Option 1: MPLS WAN physical WAN interface

The DMVPN design is using FVRF, so you must place the WAN interface into the VRF configured in the previous procedure.

Step 1: Enable the interface, select the VRF, and assign the IP address.

```
interface GigabitEthernet0/0/3
 vrf forwarding IWAN-TRANSPORT-1
 ip address 192.168.6.41 255.255.255.252
 no shutdown
```

Step 2: Configure the VRF-specific default routing.

The VRF created for FVRF must have its own default route to the MPLS. This default route points to the MPLS PE router's IP address and is used by DMVPN for tunnel establishment.

```
ip route vrf IWAN-TRANSPORT-1 0.0.0.0 0.0.0.0 192.168.6.42
```

Option 2: Internet WAN physical WAN interface

The DMVPN design is using FVRF, so you must place the WAN interface into the VRF configured in Procedure 5, "Configure the WAN-facing VRF."

Step 1: Enable the interface, select the VRF, and assign the IP address.

```
interface GigabitEthernet0/0/3
 vrf forwarding IWAN-TRANSPORT-2
 ip address 192.168.146.11 255.255.255.0
 no shutdown
```

Step 2: Configure the VRF-specific default routing.

The VRF created for FVRF must have its own default route to the Internet. This default route points to the Cisco ASA 5500's DMZ interface IP address.

```
ip route vrf IWAN-TRANSPORT-2 0.0.0.0 0.0.0.0 192.168.146.1
```

Procedure 5 Configure the mGRE tunnel

The parameters in the table below are used in this procedure. Choose the row that represents the transit site BR that you are configuring. This procedure applies to the transit site BR in the IWAN hybrid design model.

Table 75 DMVPN tunnel parameters

Hostname	Tunnel type	Tunnel number	Tunnel IP address
VPN-MPLS-ASR1002X-T1	MPLS	10	10.6.34.2/23
VPN-INET-ASR1002X-T2	INET	11	10.6.36.2/23

Step 1: Configure the basic interface settings.

The tunnel number is arbitrary, but it is best to begin tunnel numbering at 10 or above, because other features deployed in this design may also require tunnels and they may select lower numbers by default.

```
interface Tunnel10
 ip address 10.6.34.2 255.255.254.0
```

Procedure 6 Configure EIGRP

Step 1: Configure EIGRP network summarization.

The IP assignments for the entire network were designed so they can be summarized within a few aggregate routes. As configured below, the **summary-address** command suppresses the more specific routes. If any network within the summary is present in the route table, the summary is advertised to the remote sites, which offers a measure of resiliency. If the various networks cannot be summarized, then EIGRP continues to advertise the specific routes.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Tunnel10
   summary-address 10.6.0.0 255.255.0.0
   summary-address 10.7.0.0 255.255.0.0
   summary-address 10.8.0.0 255.255.0.0
   summary-address 10.255.240.0 255.255.248.0
```

Step 2: Tag and filter the routes.

This design uses a single EIGRP autonomous system for the WAN and all of the WAN remote sites. Every remote site is dual-connected for resiliency. However, due to the multiple paths that exist within this topology, you must try to avoid routing loops and to prevent remote sites from becoming transit sites if WAN failures were to occur.

In this design, there are different IP subnets for each DMVPN network, and the EIGRP tags are clearly defined to help with readability and troubleshooting. When a design uses more than one data center, additional tags are required in order to identify the different DMVPN hub router locations.

The following logic is used to control the routing.

- Each DMVPN network will have an EIGRP route tag to prevent routes from being re-advertised over the other DMVPN networks at the remote sites.
- All prefixes that are advertised towards the WAN are uniquely tagged.
- All DMVPN learned WAN prefixes, except those that originate locally from a hub, are advertised towards the LAN and tagged.

Outbound distribute-lists are used to set tags on the DMVPN hub routers towards the WAN and LAN. The remote-site routers use the tags set towards the WAN in order to protect against becoming transit sites.

The DMVPN hub routers use an inbound distribute-list in order to limit which routes are accepted for installation into the route table. These routers are configured to only accept routes that do not originate from the other data center's MPLS and DMVPN WAN sources. To accomplish this task, during the route redistribution process, the DMVPN hub router must explicitly tag the DMVPN learned WAN routes. The following tables show specific route tags in use.

Table 76 *Route tag information for hub and transit site routers*

DMVPN hub	DMVPN tunnel key	Tag tunnel	Block tunnel	Tag LAN
VPN-MPLS-ASR1002X-1	101 (MPLS)	101 (All routes)	103 (DC2 tagged routes)	101 (WAN routes)
VPN-INET-4451X-2	102 (INET)	102 (All routes)	104 (DC2 tagged routes)	102 (WAN routes)
VPN-MPLS-ASR1002X-T1	101 (MPLS)	103 (All routes)	101 (DC1 tagged routes)	103 (WAN routes)
VPN-INET-ASR1002X-T2	102 (INET)	104 (All routes)	102 (DC1 tagged routes)	104 (WAN routes)

The following examples show the hub and transit border routers in the IWAN hybrid design model.

Example: MPLS Hub Border Router—VPN-MPLS-ASR1002X-1

```
route-map SET-TAG-ALL permit 10
 description Tag all routes advertised through the tunnel
 set tag 101

! All MPLS tunnel interfaces are in this IP address range
ip access-list standard DMVPN-1-SPOKES
 permit 10.6.34.0 0.0.1.255

route-map SET-TAG-DMVPN-1 permit 10
 description Tag all incoming routes advertised through LAN interface
 match ip route-source DMVPN-1-SPOKES
 set tag 101
route-map SET-TAG-DMVPN-1 permit 100
 description Advertise all other routes with no tag

route-map BLOCK-DC2-DMVPN-1 deny 10
 description Block Summary route from other border
 match tag 103
route-map BLOCK-DC2-DMVPN-1 permit 100
 description Advertise all other routes

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   topology base
   distribute-list route-map SET-TAG-DMVPN-1 out Port-channel1
   distribute-list route-map SET-TAG-ALL out Tunnel10
   distribute-list route-map BLOCK-DC2-DMVPN-1 in Tunnel10
```

Example: INET Hub Border Router—VPN-INET-ISR4451X-2

```
route-map SET-TAG-ALL permit 10
 description tag all routes advertised through the tunnel
 set tag 102

! All INET tunnel interfaces are in this IP address range
ip access-list standard DMVPN-2-SPOKES
 permit 10.6.36.0 0.0.1.255

route-map SET-TAG-DMVPN-2 permit 10
 description Tag routes sourced from DMVPN-2
 match ip route-source DMVPN-2-SPOKES
 set tag 102
route-map SET-TAG-DMVPN-2 permit 100
 description Advertise all other routes with no tag

route-map BLOCK-DC2-DMVPN-2 deny 10
 description Block Summary route from other border
 match tag 104
route-map BLOCK-DC2-DMVPN-2 permit 100
 description Advertise all other routes

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   topology base
   distribute-list route-map SET-TAG-DMVPN-2 out Port-channel2
   distribute-list route-map SET-TAG-ALL out Tunnel11
   distribute-list route-map BLOCK-DC2-DMVPN-2 in Tunnel11
```

Example: MPLS Transit Border Router—VPN-MPLS-ASR1002X-T1

```
route-map SET-TAG-ALL permit 10
 description Tag all routes advertised through the tunnel
 set tag 103

! All MPLS tunnel interfaces are in this IP address range
ip access-list standard DMVPN-1-SPOKES
 permit 10.6.34.0 0.0.1.255

route-map SET-TAG-DMVPN-1 permit 10
 description Tag all incoming routes advertised through LAN interface
 match ip route-source DMVPN-1-SPOKES
 set tag 103
route-map SET-TAG-DMVPN-1 permit 100
 description Advertise all other routes with no tag

route-map BLOCK-DC1-DMVPN-1 deny 10
 description Block Summary route from other border
 match tag 101
route-map BLOCK-DC1-DMVPN-1 permit 100
 description Advertise all other routes

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   topology base
   distribute-list route-map SET-TAG-DMVPN-1 out Port-channel1
   distribute-list route-map SET-TAG-ALL out Tunnel10
   distribute-list route-map BLOCK-DC1-DMVPN-1 in Tunnel10
```

Example: INET Transit Border Router—VPN-INET-ASR1002X-T2

```
route-map SET-TAG-ALL permit 10
 description tag all routes advertised through the tunnel
 set tag 104

! All INET tunnel interfaces are in this IP address range
ip access-list standard DMVPN-2-SPOKES
 permit 10.6.36.0 0.0.1.255

route-map SET-TAG-DMVPN-2 permit 10
 description Tag routes sourced from DMVPN-2
 match ip route-source DMVPN-2-SPOKES
 set tag 104
route-map SET-TAG-DMVPN-2 permit 100
 description Advertise all other routes with no tag

route-map BLOCK-DC1-DMVPN-2 deny 10
 description Block Summary route from other border
 match tag 102
route-map BLOCK-DC1-DMVPN-2 permit 100
 description Advertise all other routes

router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
   topology base
   distribute-list route-map SET-TAG-DMVPN-2 out Port-channel2
   distribute-list route-map SET-TAG-ALL out Tunnel11
   distribute-list route-map BLOCK-DC1-DMVPN-2 in Tunnel11
```

Procedure 7 Configure network address translation on the firewall

You have to add the transit site Internet BR to your firewall configuration for network address translation.

The DMZ network uses private network (RFC 1918) addressing that is not Internet-routable, so the firewall must translate the DMZ address of the DMVPN hub router to an outside public address.

The example DMZ address to public IP address mapping is shown in the following table.

Table 77 DMVPN NAT address mapping

Hostname	DMVPN hub router DMZ address	DMVPN hub router public address (externally routable after NAT)
VPN-INET- ASR1002X-T2	192.168.146.11	172.16.140.2 (ISP-A)

First, to simplify the configuration of the security policy, you create the External DMZ network objects that are used in the firewall policies.

Table 78 External DMZ firewall network objects

Network object name	Object type	IP address	Description
outside-dmvpn-T2-ISPa	Host	172.16.140.2	DMVPN hub router T2 on ISP A (outside)

Step 1: Navigate to **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups**.

Step 2: Click **Add** > **Network Object**.

The Add Network Object dialog box appears.

Step 3: In the **Name** box, enter the name. (Example: outside-dmvpn-T2-ISPa)

Step 4: In the **Type** list, choose **Host** or **Network**. (Example: Host)

Step 5: In the **IP Address** box, enter the address. (Example: 172.16.140.2)

Step 6: In the **Description** box, enter a useful description, and then click **OK**. (Example: DMVPN hub router T2 on ISP A)

Step 7: Repeat Step 2 through Step 6 for each object listed in the above table. If an object already exists, then skip to the next object listed in the table.

Step 8: After adding all of the objects listed, on the Network Objects/Groups pane, click **Apply**.

Next, you add a network object for the private DMZ address of the DMVPN hub router.

Table 79 Private DMZ firewall network objects

Network object name	Object type	IP address	Description
dmz-dmvpn-T2	Host	192.168.146.11	DMVPN hub router T2 on vpn-dmz

Step 9: Navigate to **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups**.

Step 10: Click **Add** > **Network Object**.

The Add Network Object dialog box appears.

Step 11: In the **Name** box, enter the name. (Example: dmz-dmvpn-T2)

Step 12: In the **Type** list, choose **Host** or **Network**. (Example: Host)

Step 13: In the **IP Address** box, enter the address. (Example: 192.168.146.11)

Step 14: In the **Description** box, enter a useful description, and then click **OK**. (Example: DMVPN hub router T2 on vpn-dmz)

Step 15: Click the two down arrows. The NAT pane expands.

Step 16: Select **Add Automatic Address Translation Rules**.

Step 17: In the **Translated Address** list, choose the network object created previously. (Example: outside-dmvpn-T2-ISPa)

Step 18: Select **Use one-to-one address translation**, and then click **OK**.

Step 19: Repeat Step 10 through Step 18 for each object listed in the table above. If an object already exists, then skip to the next object listed in the table.

Step 20: After adding all of the objects listed, on the Network Objects/Groups pane, click **Apply**.

PROCESS

Configuring Transit Master Controller

1. Copy the configuration from existing router to the new router
2. Configure the transit MC platform
3. Configure connectivity to the LAN

For this process, you configure a transit MC with a similar base configuration as the existing hub MC. You have to make changes to the base configuration and the remote site routers in order to take advantage of the new transit site location.

The additional MC router has a unique pop-id, IP addresses and port-channel assignments, and a much simpler PfR MC configuration, but the rest of the configuration is the same. The hub MC has a default pop-id of 0 and transit MCs pop-id start at 1.

Table 80 *Hub and transit site MC IP addresses*

Host name	Pop ID	Loopback IP address	Port-channel IP address
MC-HY-CSR1000v-1	0	10.6.32.251/32	10.6.32.151/25
MC-HY-ASR1002X-T1	1	10.8.32.251/32	10.8.32.151/25

Follow the process “Configuring Hub Master Controller” using the base PfR information from the hub MC. Make the required changes from the procedures below in order to add a transit site to your IWAN domain.

Procedure 1 Copy the configuration from existing router to the new router

Optional

If the hardware for the transit MC is identical to the hub MC, you can use this optional procedure to copy the configuration file from one router to the other as a starting point and then, follow the procedures below. Skip this procedure if you do not want to copy the configuration from an existing router.

Step 1: Copy the running configuration from an existing router to your FTP server.

```
MC-HY-CSR1000v-1# copy running-config ftp://cisco:cisco@10.4.48.27
Address or name of remote host [10.4.48.27]?
Destination filename [mc-hy-csr100v-1-confg]?
Writing mc-hy-csr100v-1-confg !
7856 bytes copied in 0.800 secs (9820 bytes/sec)
```

Step 2: From the console of the new transit MC, copy and paste the configuration into the router before making the changes below.

Procedure 2 Configure the transit MC platform

In this procedure, you configure system settings that are unique to the transit MC.

Step 1: Configure the device host name to make it easy to identify the device.

```
hostname MC-HY-ASR1002X-T1
```

Step 2: Configure an in-band management interface.

The loopback interface is a logical interface that is always reachable as long as the device is powered on and any IP interface is reachable to the network.

The loopback address is commonly a host address with a 32-bit address mask.

```
interface Loopback 0
  ip address 10.8.32.151 255.255.255.255
```

Step 3: Configure IP unicast routing using EIGRP named mode.

EIGRP is configured facing the LAN distribution or core layer. In this design, the port-channel interface and the loopback must be EIGRP interfaces. The loopback may remain a passive interface. The network range must include both interface IP addresses, either in a single network statement or in multiple network statements.

This design uses a best practice of assigning the router ID to a loopback address.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface default
   passive-interface
  exit-af-interface
 network 10.8.0.0 0.1.255.255
 eigrp router-id 10.8.32.151
 exit-address-family
```

Procedure 3 Configure connectivity to the LAN

Any links to adjacent distribution layers should be Layer 3 links or Layer 3 EtherChannels.

Step 1: Configure a Layer 3 interface.

```
interface Port-channel21
  description IW-WAN-D3750X-T
  ip address 10.8.32.151 255.255.255.192
  no shutdown
```

Step 2: Configure EtherChannel member interfaces.

Configure the physical interfaces to tie to the logical port-channel by using the **channel-group** command. The number for the port-channel and channel-group must match. Not all router platforms can support LACP to negotiate with the switch, so EtherChannel is configured statically.

```
interface GigabitEthernet0/0/0
 description IW-WAN-D3750X-T Gig1/0/3

interface GigabitEthernet0/0/1
 description IW-WAN-D3750X-T Gig2/0/3

interface range GigabitEthernet0/0/0, GigabitEthernet0/0/1
  no ip address
  cdp enable
  channel-group 21
  no shutdown
```

Step 3: Configure the EIGRP interface.

Allow EIGRP to form neighbor relationships across the interface to establish peering adjacencies and exchange route tables. In this step, you configure EIGRP authentication by using the authentication key specified in the previous procedure.

```
router eigrp IWAN-EIGRP
 address-family ipv4 unicast autonomous-system 400
  af-interface Port-channel21
   no passive-interface
   authentication mode md5
   authentication key-chain LAN-KEY
  exit-af-interface
 exit-address-family
```

PROCESS

Configuring PfR for Transit Location

1. Verify IP connectivity to remote site loopback interfaces
2. Configure prefixes for the data center
3. Configure PfR domain in the transit MC
4. Configure PfR domain in the transit BR
5. Verify PfR domain is operational on the transit MC
6. Configure remote sites for transit site BRs

After the transit BRs and MC are configured, you will configure PfR for the transit site location.

Procedure 1 Verify IP connectivity to remote site loopback interfaces

It is mandatory to use loopback interfaces for the peering traffic between the BR and MC routers. For this design, you put the loopback addresses into a specific subnet range, so they are easily identified in the routing table. The loopback address ranges for the remote sites are as follows:

Table 81 Remote-site loopback IP address ranges

IWAN design model	Tunnel type	Loopback 0 address range
Hybrid–Primary Router	MPLS	10.255.241.0/24
Hybrid–Secondary Router	INET	10.255.242.0/24
Dual Internet–Primary Router	INET1	10.255.243.0/24
Dual Internet–Secondary Router	INET2	10.255.244.0/24

Step 1: Verify that the loopback 0 interfaces on each of your remote sites are reachable from the transit MC by using the **show ip route** command.

This example shows a loopback address range of 10.255.241.0/24 for nine remote site primary routers and an address range of 10.255.242.0/24 for four remote site secondary routers.

```
MC-HY-ASR1002X-T1# show ip route | include 10.255.241
D        10.255.241.11/32 [90/25610880] via 10.8.32.129, 1w2d, Port-channel21
D        10.255.241.12/32 [90/25610880] via 10.8.32.129, 1w2d, Port-channel21
D        10.255.241.21/32 [90/25610880] via 10.8.32.129, 1w2d, Port-channel21
D        10.255.241.22/32 [90/25610880] via 10.8.32.129, 1w2d, Port-channel21
D        10.255.241.31/32 [90/25610880] via 10.8.32.129, 1w2d, Port-channel21
D        10.255.241.32/32 [90/25610880] via 10.8.32.129, 1w2d, Port-channel21
D        10.255.241.41/32 [90/25610880] via 10.8.32.129, 1w2d, Port-channel21
D        10.255.241.42/32 [90/25610880] via 10.8.32.129, 1w2d, Port-channel21
D        10.255.241.51/32 [90/25610880] via 10.8.32.129, 1w3d, Port-channel21

MC-HY-ASR1002X-T1#show ip route | include 10.255.242
D        10.255.242.12/32 [90/25613440] via 10.8.32.129, 1w1d, Port-channel21
D        10.255.242.22/32 [90/25613440] via 10.8.32.129, 1w2d, Port-channel21
D        10.255.242.32/32 [90/25613440] via 10.8.32.129, 1w2d, Port-channel21
D        10.255.242.42/32 [90/25613440] via 10.8.32.129, 1w2d, Port-channel21
```

Procedure 2 Configure prefixes for the data center

Before the configuration of PfRv3 on the transit MC, you must create prefix lists for the data center. The enterprise-prefix list is only configured on the hub MC and you will not configure one on the transit MC.

The site-prefix range for the transit site includes the prefixes at this specific site, which is normally a WAN aggregation or data center site. Site-prefixes are typically statically defined at WAN aggregation and DC sites and discovered automatically at remote sites.

Tech Tip

The ip prefix-list options ge and le are not supported by PfR.

Step 1: Create the transit site prefix list.

```
ip prefix-list [prefix-list-name] seq [value] permit [prefix list]
```

Example

This example shows a data center network with two class B private address blocks of 10.4.0.0 and 10.8.0.0.

```
ip prefix-list DC2-PREFIXES seq 10 permit 10.4.0.0/16
ip prefix-list DC2-PREFIXES seq 20 permit 10.8.0.0/16
```

Procedure 3 Configure PfR domain in the transit MC

Domain policies are configured on the hub MC. These policies are distributed to branch MCs and the transit MC by using the peering infrastructure. All sites that are in the same domain will share the same set of PfR policies. The transit MC must peer to the hub MC to get the policy information.

Step 1: Create the transit MC domain.

```
domain [name]
 vrf [name]
  master transit [number]
   source-interface [interface]
   site-prefixes prefix-list [prefixes from previous procedure]
   password [password of hub MC]
   hub [IP address of hub MC]
```

Example

```
domain iwan
 vrf default
  master transit 1
   source-interface Loopback0
   site-prefixes prefix-list DC2-PREFIXES
   password c1sco123
   hub 10.6.32.251
```

Step 2: Verify the hub MC policy configuration is available by using the **show domain [name] master policy** command.

The output from this command should look the same as the output on the hub MC.

Procedure 4 Configure PfR domain in the transit BR

The transit BRs are also the DMVPN hub WAN aggregation routers for the transit site network. The PfRv3 configurations for standalone BRs are much simpler because they dynamically learn their policy information from the transit MC. The transit BR routers are also used to advertise the path names and path-ids specified in the hub MC configuration.

There is an optional feature called *zero-SLA* that reduces the probing to only the default class by muting the other DSCP probes. This feature is useful on Internet connections where nothing is guaranteed. Zero-SLA reduces bandwidth usage on metered interfaces such as 4G LTE or other Internet connections with a monthly data cap limit.

Tech Tip

If you want to add the zero-SLA feature to an existing hub BR, you must shut down the DMVPN tunnel interface before configuring. After the feature is added to the hub BR, bring the tunnel interface back up.

Table 82 Transit BR path and IP addresses

Host name	Path	Path ID	Loopback IP address	Zero SLA
VPN-MPLS-ASR1002X-T1	MPLS	1	10.8.32.241/32	No
VPN-INET-ASR1002X-T2	INET	2	10.8.32.242/32	Yes (optional)

Step 1: Create the transit BR domain.

```
domain [name]
 vrf [name]
  border (create the BR)
   source-interface [interface]
   master [IP address of transit MC]
   password [password of hub MC]
```

Example

```
domain iwan
 vrf default
  border
   source-interface Loopback0
   master 10.8.32.251
   password c1sco123
```

Step 2: Add the path names and path-ids to the tunnel interfaces of the transit BR.

```
interface Tunnel [value]
 domain [name] path [name] path-id [number] zero-sla
```

Example

This example is the primary transit BR using Tunnel 10 with MPLS as the provider.

```
interface Tunnel10
 domain iwan path MPLS path-id 1
```

Step 3: (Optional) This example is the secondary hub BR using Tunnel 11 with INET as the provider and the zero-sla feature. If this is an existing configuration, you shut down the interface, add the zero SLA feature and then, bring the interface back up.

```
interface Tunnel11
 shutdown
 domain iwan path INET path-id 2 zero-sla
 no shutdown
```

Step 4: Verify the border is operational by using the **show domain [name] border status** command.

Step 5: Repeat this procedure for each transit BR by using the appropriate path name and path-id.

Procedure 5 Verify PfR domain is operational on the transit MC

The PfR path names and path-ids are automatically discovered at the remote site routers from the configuration entered into the tunnel interfaces at the hub and transit sites. The hub MC uses the path names and path-ids to determine where traffic should be sent according to its policies.

Step 1: Verify the domain is operational from the transit MC using the **show domain [name] master status** command.

This example shows the transit MC of the IWAN hybrid model in an operational state. The transit BRs are both connected and using their respective Tunnel interfaces as the exits for the transit location.

```
MC-HY-ASR1002X-T1#show domain iwan master status
  *** Domain MC Status ***
 Master VRF: Global
  Instance Type:    Transit
  POP ID:      1
  Instance id:      0
  Operational status:  Up
  Configured status:  Up
  Loopback IP Address: 10.8.32.251
  Load Balancing:
   Operational Status: Up
   Max Calculated Utilization Variance: 0%
```

```
    Last load balance attempt: never
    Last Reason:  Variance less than 20%
    Total unbalanced bandwidth:
          External links: 0 Kbps  Internet links: 0 Kbps
   External Collector: 10.4.48.36 port: 9991
   Route Control: Enabled
   Transit Site Affinity: Enabled
   Load Sharing: Enabled
   Mitigation mode Aggressive: Disabled
   Policy threshold variance: 20
   Minimum Mask Length: 28
   Syslog TCA suppress timer: 180 seconds
   Traffic-Class Ageout Timer: 5 minutes
   Minimum Packet Loss Calculation Threshold: 15 packets
   Minimum Bytes Loss Calculation Threshold: 1 bytes
   Minimum Requirement: Met

   Borders:
     IP address: 10.8.32.241
     Version: 2
     Connection status: CONNECTED (Last Updated 1w2d ago )
     Interfaces configured:
       Name: Tunnel10 | type: external | Service Provider: MPLS path-id:1 | Sta-
tus: UP | Zero-SLA: NO | Path of Last Resort: Disabled
           Number of default Channels: 0
     Tunnel if: Tunnel0

     IP address: 10.8.32.242
     Version: 2
     Connection status: CONNECTED (Last Updated 1w2d ago )
     Interfaces configured:
       Name: Tunnel11 | type: external | Service Provider: INET path-id:2 | Sta-
tus: UP | Zero-SLA: YES | Path of Last Resort: Disabled
           Number of default Channels: 0
     Tunnel if: Tunnel1
```

Procedure 6 Configure remote sites for transit site BRs

An additional NHRP command has to be added to the tunnel interfaces of remote site BRs for them to begin using the transit BRs.

Table 83 NHRP parameters

Hostname	Tunnel type	Tunnel number	Tunnel IP address	MPLS/public IP address
VPN-MPLS-ASR1002X-T1	MPLS	10	10.6.34.2	192.168.6.41
VPN-INET-ASR1002X-T2	INET	11	10.6.36.2	172.16.140.2 (ISP A)

Step 1: Configure NHRP.

The DMVPN hub router is the NHRP server for all of the spokes. Remote routers use NHRP in order to determine the tunnel destinations for peers attached to the mGRE tunnel.

The spoke router requires an additional configuration statement in order to define the NHRP server. This statement includes the NBMA definition for the DMVPN hub router tunnel endpoint. EIGRP relies on a multicast transport. Spoke routers require the NHRP multicast keyword in this statement.

The value used for the NHS is the mGRE tunnel address for the DMVPN hub router. The NBMA entry must be set to either the MPLS DMVPN hub router's actual public address or the outside NAT value of the DMVPN hub, as configured on the Cisco ASA 5500. This design uses the values shown in the table above.

Example: Single Router Hybrid—RS11-2921

```
interface Tunnel10
 ip nhrp nhs 10.6.34.2 nbma 192.168.6.41 multicast

interface Tunnel11
 ip nhrp nhs 10.6.36.2 nbma 172.16.140.2 multicast
```

Step 2: Confirm the hub and transit BRs are reachable with **show ip eigrp neighbors**.

```
RS11-2921#show ip eigrp neighbors
EIGRP-IPv4 VR(IWAN-EIGRP) Address-Family Neighbors for AS(400)
H   Address      Interface     Hold Uptime   SRTT  RTO  Q  Seq
                               (sec)         (ms)       Cnt Num
3   10.6.36.1    Tu11            55 1w3d        1   100  0  7806
2   10.6.34.1    Tu10            55 5w5d        1   100  0  17528
0   10.6.34.2    Tu10            57 5w5d        1   100  0  8851
1   10.6.36.2    Tu11            56 5w5d        1   100  0  16134
```

Step 3: Repeat this procedure for each remote site that will use the transit BRs.

Deploying IWAN Quality of Service

QoS has already proven itself as the enabling technology for the convergence of voice, video and data networks. As business needs evolve, so do demands on QoS technologies. The need to protect voice, video and critical data with QoS mechanisms is extremely important on the WAN because access speeds are much lower than the LAN networks that feed them.

PROCESS

Configuring QoS for DMVPN Routers

1. Create the class maps in order to classify traffic
2. Create policy map with queuing policy

When configuring WAN-edge QoS, you are defining how traffic egresses your network. It is critical that the classification, marking, and bandwidth allocations align to the service provider, offering to ensure consistent QoS treatment end to end.

The Per-Tunnel QoS for DMVPN feature allows the configuration of a QoS policy on a DMVPN hub on a per-tunnel (spoke) basis. The QoS policy on a tunnel instance allows you to shape the tunnel traffic to individual spokes (parent policy) and to differentiate between traffic classes within the tunnel for appropriate treatment (child policy).

You can also mark the header of the GRE tunneled packets by using the QoS policy map classes. There are two methods for marking the DSCP of the tunnel headers in order to influence per-hop treatment within the service provider network. One method applies the policy to a virtual tunnel interface and the second method applies the policy to a physical interface.

The following table shows an example of how to mark the tunnel headers when using a 12- or 8-class model in the enterprise, while combining the traffic classes into a smaller 6-, 5- or 4-class model in the service provider network.

The tunnel markings must match the service provider offering, so you will have to adjust the table below according to your specific service level agreement.

Figure 36 *QoS class model mapping: Tunnel mappings must match provider*

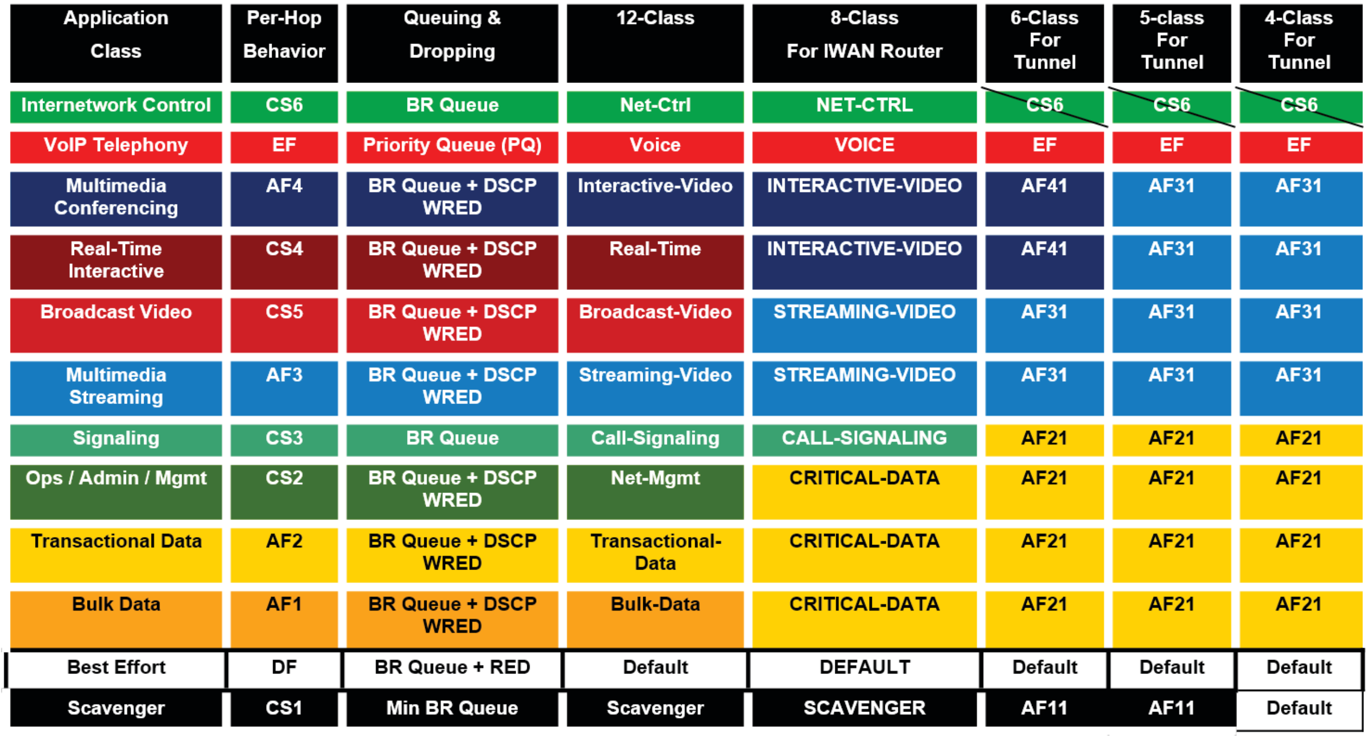

Application Class	Per-Hop Behavior	Queuing & Dropping	12-Class	8-Class For IWAN Router	6-Class For Tunnel	5-class For Tunnel	4-Class For Tunnel
Internetwork Control	CS6	BR Queue	Net-Ctrl	NET-CTRL	~~CS6~~	~~CS6~~	~~CS6~~
VoIP Telephony	EF	Priority Queue (PQ)	Voice	VOICE	EF	EF	EF
Multimedia Conferencing	AF4	BR Queue + DSCP WRED	Interactive-Video	INTERACTIVE-VIDEO	AF41	AF31	AF31
Real-Time Interactive	CS4	BR Queue + DSCP WRED	Real-Time	INTERACTIVE-VIDEO	AF41	AF31	AF31
Broadcast Video	CS5	BR Queue + DSCP WRED	Broadcast-Video	STREAMING-VIDEO	AF31	AF31	AF31
Multimedia Streaming	AF3	BR Queue + DSCP WRED	Streaming-Video	STREAMING-VIDEO	AF31	AF31	AF31
Signaling	CS3	BR Queue	Call-Signaling	CALL-SIGNALING	AF21	AF21	AF21
Ops / Admin / Mgmt	CS2	BR Queue + DSCP WRED	Net-Mgmt	CRITICAL-DATA	AF21	AF21	AF21
Transactional Data	AF2	BR Queue + DSCP WRED	Transactional-Data	CRITICAL-DATA	AF21	AF21	AF21
Bulk Data	AF1	BR Queue + DSCP WRED	Bulk-Data	CRITICAL-DATA	AF21	AF21	AF21
Best Effort	DF	BR Queue + RED	Default	DEFAULT	Default	Default	Default
Scavenger	CS1	Min BR Queue	Scavenger	SCAVENGER	AF11	AF11	Default

Tech Tip

Because the provider will generally have a network control queue that they do not declare as part of their customer-facing model, the traffic in the NET-CTRL queue will be marked as CS6 on the tunnel header. Because the network elements use this traffic to ensure stability under congestion and when a device is oversubscribed, CS6 traffic should be preserved.

Traffic is regulated from the central site (hub) routers to the remote-site routers on a per-tunnel (spoke) basis. The hub site is unable to send more traffic than a single remote-site can handle, and this ensures that high bandwidth hub sites do not overrun lower bandwidth remote-sites.

The following two procedures apply to all DMVPN WAN hub and spoke routers.

Procedure 1 Create the class maps in order to classify traffic

This number of classes in an egress policy should be chosen based on interface speed, available queues, and device capabilities. This design guide uses an 8-class model in the enterprise and the examples will have to be modified for other models, as required.

Use the **class-map** command to define a traffic class and identify traffic to associate with the class name. You use these class names when configuring policy maps that define actions you want to take against the traffic type. The **class-map** command sets the match logic. In this case, the match-any keyword indicates that the maps match any of the specified criteria. This keyword is followed by the name you want to assign to the class of service. After you have configured the **class-map** command, you define specific values, such as DSCP and protocols to match with the match command.

Step 1: Create the class maps for DSCP matching.

You do not need to explicitly configure the default class.

```
class-map match-any [class-map name]
 match dscp [dcsp value] [optional additional dscp value(s)]
```

Table 84 *Class of service for 8-class model*

Class Name	Traffic type	DSCP values
VOICE	Voice traffic	ef
INTERACTIVE-VIDEO	Interactive video (video conferencing)	cs4, af4
STREAMING-VIDEO	Video streaming	cs5, af3
NET-CTRL	Routing protocols	cs6
CALL-SIGNALING	Voice and video call signaling	cs3
CRITICAL-DATA	Network management, highly interactive	cs2, af1, af2
default	Best effort	all others
SCAVENGER	Scavenger	cs1

Step 2: Repeat this step to create a class-map for each of the seven non-default WAN classes of service listed in the table above.

Example: Class Maps for 8-class QoS model

```
class-map match-any VOICE
 match dscp ef
class-map match-any INTERACTIVE-VIDEO
 match dscp cs4 af41 af42 af43
class-map match-any STREAMING-VIDEO
```

```
 match dscp cs5 af31 af32 af33
class-map match-any NET-CTRL
 match dscp cs6
class-map match-any CALL-SIGNALING
 match dscp cs3
class-map match-any CRITICAL-DATA
 match dscp cs2 af11 af12 af13 af21 af22 af23
class-map match-any SCAVENGER
 match dscp cs1
```

Tech Tip

You do not need to configure a Best-Effort class-map. This is implicitly included within class-default as shown in Procedure 2.

Procedure 2 Create policy map with queuing policy

The WAN policy map references the class names you created in the previous procedures and defines the queuing behavior, along with the minimum guaranteed bandwidth allocated to each class. Using a policy-map accomplishes this specification. Then, each class within the policy map invokes an egress queue and assigns a percentage of bandwidth. One additional default class defines the minimum allowed bandwidth available for best effort traffic.

There are two methods for marking the tunnel headers. The method you use depends on whether you apply the policy to a virtual tunnel interface or a physical interface. Use the **set dscp tunnel** command when applying the policy-map to a tunnel interface to change the outer header, while leaving the original 12-class marking on the inner header unchanged. Use the **set dscp** command when applying the policy-map to a physical interface because in the order of operations, the outer header has already been imposed and the inner header will not be modified.

Tech Tip

The local router policy maps define eight classes, while most service providers offer only six classes of service. The NET-CTRL class is defined to ensure the correct classification, marking, and queuing of network-critical traffic on egress to the WAN. After the traffic has been transmitted to the service provider, it is put into a queue that is not part of the customer-facing class model.

Step 1: Create the policy map.

```
policy-map [policy-map-name]
```

Step 2: Apply the previously created class-map.

```
class [class-name]
```

Step 3: (Optional) Define what proportion of available bandwidth should be reserved for this class of traffic under congestion.

```
bandwidth remaining percent [percentage]
```

Step 4: (Optional) Define the congestion avoidance mechanism.

```
random-detect [type]
```

Step 5: (Optional) For priority queues, define the priority level for the class.

```
priority level [value]
```

Step 6: (Optional) For priority queues, define the amount of bandwidth that may be consumed by priority traffic.

```
police cir [percentage]
```

Step 7: (Optional) For QOS policies that will be attached to tunnel interfaces (hub router configuration), mark the DSCP in the tunnel header.

```
set dscp tunnel [dscp value]
```

Step 8: (Optional) For QOS policies that will be attached to physical interface (remote-site router configuration), mark the DSCP in the tunnel header.

```
set dscp [dscp value]
```

Table 85 Bandwidth, congestion avoidance and outbound tunnel values

Class of service for 8-class model	Bandwidth %	Congestion avoidance	Tunnel DSCP values for 6-class model
VOICE	10 (PQ)	–	ef
INTERACTIVE-VIDEO	30 remaining	DSCP based WRED	af41
STREAMING-VIDEO	10 remaining	DSCP based WRED	af31
NET-CTRL	5 remaining	–	cs6 (pass through)
CALL-SIGNALING	4 remaining	–	af21
CRITICAL-DATA	25 remaining	DSCP based WRED	af21
default	25 remaining	random	(pass through)
SCAVENGER	1 remaining	–	af11

Step 9: Repeat Step 2 through Step 8 for each class in the previous table, including the default class.

Tech Tip

The default class does not set the value, which allows traffic to pass through with markings that do not fit into the values specified by the match statements used above.

Example: WAN aggregation policy map for 6-class service provider offering

This example uses the **set dscp tunnel** command in each class because the policy is applied to the tunnel interfaces on the WAN aggregation routers.

```
policy-map WAN
 class INTERACTIVE-VIDEO
  bandwidth remaining percent 30
  random-detect dscp-based
  set dscp tunnel af41
 class STREAMING-VIDEO
  bandwidth remaining percent 10
  random-detect dscp-based
  set dscp tunnel af31
 class NET-CTRL
  bandwidth remaining percent 5
  set dscp tunnel cs6
 class CALL-SIGNALING
  bandwidth remaining percent 4
  set dscp tunnel af21
  class CRITICAL-DATA
  bandwidth remaining percent 25
  random-detect dscp-based
  set dscp tunnel af21
 class SCAVENGER
  bandwidth remaining percent 1
  set dscp tunnel af11
 class VOICE
  priority level 1
  police cir percent 10
  set dscp tunnel ef
 class class-default
  bandwidth remaining percent 25
  random-detect
```

Example: Remote site policy map for 6-class service provider offering

This example uses the **set dscp** command in each class, because the policy is applied to the physical interfaces on the remote site routers.

```
policy-map WAN
class INTERACTIVE-VIDEO
  bandwidth remaining percent 30
  random-detect dscp-based
  set dscp af41
class STREAMING-VIDEO
  bandwidth remaining percent 10
  random-detect dscp-based
  set dscp af31
 class NET-CTRL
  bandwidth remaining percent 5
  set dscp cs6
 class CALL-SIGNALING
  bandwidth remaining percent 4
  set dscp af21
  class CRITICAL-DATA
  bandwidth remaining percent 25
  random-detect dscp-based
  set dscp af21
 class SCAVENGER
  bandwidth remaining percent 1
  set dscp af11
 class VOICE
  priority level 1
  police cir percent 10
  set dscp ef
 class class-default
  bandwidth remaining percent 25
  random-detect
```

Tech Tip

Although these bandwidth assignments represent a good baseline, it is important to consider your actual traffic requirements per class and adjust the bandwidth settings accordingly.

PROCESS

Applying DMVPN QoS Policy to DMVPN Hub Routers

1. Configure shaping policy for hub router
2. Configure per-tunnel QoS policies for DMVPN hub router
3. Configure per-tunnel QoS NHRP policies on DMVPN hub router

This process applies only to DMVPN WAN Aggregation routers.

Procedure 1 Configure shaping policy for hub router

With WAN interfaces using Ethernet as an access technology, the demarcation point between the enterprise and service provider may no longer have a physical-interface bandwidth constraint. Instead, a specified amount of access bandwidth is contracted with the service provider. To ensure the offered load to the service provider does not exceed the contracted rate that results in the carrier discarding traffic, you need to configure shaping on the physical interface. When you configure the **shape average** command, ensure that the value matches the contracted bandwidth rate from your service provider.

Tech Tip

QoS on a physical interface is limited only to the class default shaper. Other QoS configurations on the physical interface are not supported.

You must apply the class default shaper policy map on the main interface before applying the tunnel policy map.

The class default shaper policy map must contain only the class class-default and shape commands.

Create the policy map and configure the shaper for the default class.

As a best practice, embed the transport number within the name of the policy map.

```
policy-map [policy-map-name]
  class class-default
    shape average [bandwidth (kbps)]
```

Step 1: Apply the shaper to the WAN interface.

You must apply the service policy needs to be applied in the outbound direction.

```
interface [interface type] [number]
    service-policy output [policy-map-name]
```

Example: Physical interface

This example shows a router with a 100 Mbps service rate configured on a 1 Gbps physical interface.

```
policy-map POLICY-TRANSPORT-1-SHAPE-ONLY
 class class-default
  shape average 100000000

interface GigabitEthernet0/0/3
 bandwidth 100000
 service-policy output POLICY-TRANSPORT-1-SHAPE-ONLY
```

Procedure 2 Configure per-tunnel QoS policies for DMVPN hub router

The QoS policy on a tunnel instance allows you to shape the tunnel traffic to individual spokes and to differentiate between traffic classes within the tunnel for appropriate treatment.

The QoS policy on the tunnel instance is defined and applied only to the DMVPN hub routers at the central site. The remote-site router signals the QoS group policy information to the hub router with a command in the NHRP configuration, which greatly reduces QoS configuration and complexity. The hub router applies the signaled policy in the egress direction for each remote site.

The **bandwidth remaining ratio** command is used to provide each site with their fair share of the remaining bandwidth when the outbound interface is experiencing congestion. If you do not use this command, the lower bandwidth sites will get all of their assigned bandwidth, while the higher bandwidth sites will get less than their fair share.

In the example below, divide the shape average bandwidth by 1 Mbps to come up with the value for the ratio. If you have sites with less than 5 Mbps of shape average bandwidth, you should divide the shape average for all of your sites by 100 Kbps to ensure they all get a reasonable ratio greater than 1.

Tech Tip

With Per-Tunnel QoS for DMVPN, the queuing and shaping is performed at the outbound physical interface for the GRE/IPsec tunnel packets. This means that the GRE header, the IPsec header and the layer2 (for the physical interface) header are included in the packet-size calculations for shaping and bandwidth queuing of packets under QoS.

The values in the table are examples; make sure to adjust these values for your specific needs and remote-site bandwidth provisioned with your ISP.

Table 86 Per-tunnel QoS policies

Policy name	Class	Bandwidth bps	Bandwidth remaining ratio
RS-GROUP-300MBPS-POLICY	class-default	**300000000**	**300**
RS-GROUP-200MBPS-POLICY	class-default	**200000000**	**200**
RS-GROUP-100MBPS-POLICY	class-default	**100000000**	**100**
RS-GROUP-50MBPS-POLICY	class-default	**50000000**	**50**
RS-GROUP-30MBPS-POLICY	class-default	**30000000**	**30**
RS-GROUP-20MBPS-POLICY	class-default	**20000000**	**20**
RS-GROUP-10MBPS-POLICY	class-default	**10000000**	**10**
RS-GROUP-4G-POLICY	class-default	**8000000**	**8**

Step 1: Create a policy.

```
policy-map [policy-map-name]
```

Step 2: Define a shaper and bandwidth remaining ratio for the default-class and apply the **WAN** QoS queuing child service policy created in Procedure 2, "Create policy map with queuing policy"

```
policy-map [policy-map-name]
   class class-default
    shape average [bandwidth (kbps)]
    bandwidth remaining ratio [shape average bandwidth/1 Mbps]
      service-policy [policy-map name]
```

Step 3: For each remote-site type, repeat steps 1 and 2.

Example: Hub router

```
policy-map RS-GROUP-300MBPS-POLICY
 class class-default
  shape average 300000000
  bandwidth remaining ratio 300
   service-policy WAN
policy-map RS-GROUP-200MBPS-POLICY
 class class-default
  shape average 200000000
  bandwidth remaining ratio 200
   service-policy WAN
policy-map RS-GROUP-100MBPS-POLICY
 class class-default
```

```
  shape average 100000000
  bandwidth remaining ratio 100
   service-policy WAN
policy-map RS-GROUP-50MBPS-POLICY
 class class-default
  shape average 50000000
  bandwidth remaining ratio 50
   service-policy WAN
policy-map RS-GROUP-30MBPS-POLICY
 class class-default
  bandwidth remaining ratio 30
  shape average 30000000
   service-policy WAN
policy-map RS-GROUP-20MBPS-POLICY
 class class-default
  shape average 20000000
  bandwidth remaining ratio 20
   service-policy WAN
policy-map RS-GROUP-10MBPS-POLICY
 class class-default
  shape average 10000000
  bandwidth remaining ratio 10
   service-policy WAN
policy-map RS-GROUP-4G-POLICY
 class class-default
  shape average 8000000
  bandwidth remaining ratio 8
   service-policy WAN
```

Procedure 3 Configure per-tunnel QoS NHRP policies on DMVPN hub router

The QoS policy that the hub uses for a particular endpoint or spoke is selected by the NHRP group in which the spoke is configured.

Prerequisites and important caveats:

- DMVPN must be fully configured and operational before you can configure an NHRP group on a spoke or map the NHRP group to a QoS policy on a hub.
- Although you may configure multiple spokes as part of the same NHRP group, the tunnel traffic for each spoke is measured individually for shaping and policing.
- Only output NHRP policies are supported. These apply to per-site traffic egressing the router towards the WAN.

Step 1: Create NHRP group policy name mapping and apply the policies configured in the previous procedure to the DMVPN tunnel interface on the hub router.

```
interface tunnel[number]
 ip nhrp map group [NHRP GROUP Policy Name] service-policy output [policy-map name]
```

Example: Hub router

```
interface tunnel10
 ip nhrp map group RS-GROUP-300MBPS service-policy output RS-GROUP-300MBPS-POLICY
 ip nhrp map group RS-GROUP-200MBPS service-policy output RS-GROUP-200MBPS-POLICY
 ip nhrp map group RS-GROUP-100MBPS service-policy output RS-GROUP-100MBPS-POLICY
 ip nhrp map group RS-GROUP-50MBPS service-policy output RS-GROUP-50MBPS-POLICY
 ip nhrp map group RS-GROUP-30MBPS service-policy output RS-GROUP-30MBPS-POLICY
 ip nhrp map group RS-GROUP-20MBPS service-policy output RS-GROUP-20MBPS-POLICY
 ip nhrp map group RS-GROUP-10MBPS service-policy output RS-GROUP-10MBPS-POLICY
 ip nhrp map group RS-GROUP-4G service-policy output RS-GROUP-4G-POLICY
```

PROCESS

Applying QoS Configurations to Remote Site Routers

1. Configure per-tunnel QoS NHRP policy on remote-site routers
2. Configure physical interface QoS policy on remote-site routers
3. Apply QoS policy to the physical interface on remote-site routers
4. Verify QoS policy on physical interfaces of remote site router
5. Verify DMVPN per-tunnel QoS from hub routers

This process completes the remote-site QoS configuration and applies to all DMVPN spoke routers.

Procedure 1 Configure per-tunnel QoS NHRP policy on remote-site routers

This procedure configures the remote-site router to reference the QoS policy configured on the hub site routers.

Step 1: Apply the NHRP group policy to the DMVPN tunnel interface on the corresponding remote-site router. Use the NHRP group name as defined on the hub router in Procedure 2, "Configure per tunnel QoS policies for DMVPN hub router," above.

```
interface Tunnel[value]
 ip nhrp group [NHRP GROUP Policy Name]
```

Example: Remote-site router with dual-link

This example shows a remote-site using a 20 Mbps policy and a 10 Mbps policy.

```
interface Tunnel10
  ip nhrp group RS-GROUP-20MBPS

interface Tunnel11
  ip nhrp group RS-GROUP-10MBPS
```

Procedure 2 Configure physical interface QoS policy on remote-site routers

Repeat this procedure in order to support remote-site routers that have multiple WAN connections attached to different interfaces.

With WAN interfaces using Ethernet as an access technology, the demarcation point between the enterprise and service provider may no longer have a physical-interface bandwidth constraint. Instead, a specified amount of access bandwidth is contracted with the service provider. To ensure the offered load to the service provider does not exceed the contracted rate that results in the carrier discarding traffic, configure shaping on the physical interface. This shaping is accomplished with a QoS service policy. You configure a QoS service policy on the outside Ethernet interface, and this parent policy includes a shaper that then references a second or subordinate (child) policy that enables queuing within the shaped rate. This is called a hierarchical Class-Based Weighted Fair Queuing configuration. When you configure the **shape average** command, ensure that the value matches the contracted bandwidth rate from your service provider.

Step 1: Create the parent policy map.

As a best practice, embed the transport number within the name of the parent policy map.

```
policy-map [policy-map-name]
```

Step 2: Configure the shaper.

```
class [class-name]
  shape  [average | peak] [bandwidth (kbps)]
```

Step 3: Apply the child service policy as defined in Procedure 2, "Create policy map with queuing policy" above.

```
service-policy WAN
```

Example: Remote-site router with dual-link

This example shows a router with a 20-Mbps rate on interface GigabitEthernet0/0 for transport 1 and a 10-Mbps rate on interface GigabitEthernet0/1 for transport 2.

```
policy-map POLICY-TRANSPORT-1
 class class-default
  shape average 20000000
 service-policy WAN

policy-map POLICY-TRANSPORT-2
 class class-default
  shape average 10000000
 service-policy WAN
```

Procedure 3 Apply QoS policy to the physical interface on remote-site routers

Repeat this procedure in order to support remote-site routers that have multiple WAN connections attached to different interfaces.

To invoke shaping and queuing on a physical interface, you must apply the parent policy that you configured in the previous procedure.

Step 1: Select the WAN interface.

```
interface [interface type] [number]
```

Step 2: Apply the WAN QoS policy.

The service policy needs to be applied in the outbound direction.

```
service-policy output [policy-map-name]
```

Example: Remote-site router with dual-link

```
interface GigabitEthernet0/0
 service-policy output POLICY-TRANSPORT-1

interface GigabitEthernet0/1
 service-policy output POLICY-TRANSPORT-2
```

Procedure 4 Verify QoS policy on physical interfaces of remote site router

After all of the physical interfaces on a router are configured, you can verify each one before moving to the next remote site.

Step 1: Verify the QoS output policy on each interface is correct by using the **show policy-map interface** command.

This example is truncated due to the overall length.

```
RS11-2921# show policy-map interface GigabitEthernet 0/0
GigabitEthernet0/0

  Service-policy output: POLICY-TRANSPORT-1

    Class-map: class-default (match-any)
      66941984 packets, 14951533762 bytes
      5 minute offered rate 83000 bps, drop rate 0000 bps
      Match: any
      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 66941987/15813338284
      shape (average) cir 30000000, bc 120000, be 120000
      target shape rate 30000000

      Service-policy : WAN
```

```
        queue stats for all priority classes:
          Queueing
          priority level 1
          queue limit 64 packets
          (queue depth/total drops/no-buffer drops) 0/0/0
          (pkts output/bytes output) 572961/124080310

        Class-map: INTERACTIVE-VIDEO (match-any)
          530608 packets, 205927572 bytes
          5 minute offered rate 0000 bps, drop rate 0000 bps
          Match:  dscp cs4 (32) af41 (34) af42 (36) af43 (38)
            530608 packets, 205927572 bytes
            5 minute rate 0 bps
          Queueing
          queue limit 64 packets
          (queue depth/total drops/no-buffer drops) 0/0/0
          (pkts output/bytes output) 530608/209102464
          bandwidth remaining 30%
            Exp-weight-constant: 9 (1/512)
            Mean queue depth: 0 packets
            dscp       Transmitted         Random drop      Tail drop
Minimum        Maximum     Mark
                    pkts/bytes            pkts/bytes       pkts/bytes
thresh         thresh     prob

            af41      530608/209102464       0/0              0/0
32            40  1/10
          QoS Set
            dscp af41
              Packets marked 530608
```

Step 2: Repeat the previous step for each interface configured with QoS.

Procedure 5 Verify DMVPN per-tunnel QoS from hub routers

After the all of the DMVPN routers are configured for Per-Tunnel QoS, you can verify the configurations from the hub router.

Step 1: Verify the Per-Tunnel QoS output policy to each remote-site is correct by using the **show dmvpn detail** command.

This example is truncated due to the overall length.

```
VPN-MPLS-ASR1002X-1# show dmvpn detail
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        T1 - Route Installed, T2 - Nexthop-override
        C - CTS Capable
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================


Interface Tunnel10 is up/up, Addr. is 10.6.34.1, VRF ""
   Tunnel Src./Dest. addr: 192.168.6.1/MGRE, Tunnel VRF "IWAN-TRANSPORT-1"
   Protocol/Transport: "multi-GRE/IP", Protect "DMVPN-PROFILE-TRANSPORT-1"
   Interface State Control: Disabled
   nhrp event-publisher : Disabled
Type:Hub, Total NBMA Peers (v4/v6): 9


# Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb    Target Network
----- --------------- --------------- ----- -------- ----- -----------------
    1 192.168.6.5          10.6.34.11    UP     1w0d     D      10.6.34.11/32
NHRP group: RS-GROUP-30MBPS
 Output QoS service-policy applied: RS-GROUP-30MBPS-POLICY
    1 192.168.6.9          10.6.34.12    UP     1w0d     D      10.6.34.12/32
NHRP group: RS-GROUP-20MBPS
 Output QoS service-policy applied: RS-GROUP-20MBPS-POLICY
    1 192.168.6.13         10.6.34.21    UP     1w0d     D      10.6.34.21/32
```

```
        NHRP group: RS-GROUP-30MBPS
         Output QoS service-policy applied: RS-GROUP-30MBPS-POLICY
            1 192.168.6.17         10.6.34.22    UP     1w0d     D      10.6.34.22/32
        NHRP group: RS-GROUP-20MBPS
         Output QoS service-policy applied: RS-GROUP-20MBPS-POLICY
            1 192.168.6.21         10.6.34.31    UP     1w0d     D      10.6.34.31/32
        NHRP group: RS-GROUP-200MBPS
         Output QoS service-policy applied: RS-GROUP-200MBPS-POLICY
            1 192.168.6.25         10.6.34.32    UP     1w0d     D      10.6.34.32/32
        NHRP group: RS-GROUP-300MBPS
         Output QoS service-policy applied: RS-GROUP-300MBPS-POLICY
            1 192.168.6.29         10.6.34.41    UP     1w0d     D      10.6.34.41/32
        NHRP group: RS-GROUP-30MBPS
         Output QoS service-policy applied: RS-GROUP-30MBPS-POLICY
            1 192.168.6.33         10.6.34.42    UP     1w0d     D      10.6.34.42/32
        NHRP group: RS-GROUP-300MBPS
         Output QoS service-policy applied: RS-GROUP-300MBPS-POLICY
```

Step 2: Repeat the previous step for each DMVPN hub router.

Deploying IWAN Monitoring

NetFlow operates by creating a NetFlow cache entry that contains information for all active flows on a NetFlow-enabled device. NetFlow builds its cache by processing the first packet of a flow through the standard switching path. It maintains a flow record within the NetFlow cache for all active flows. Each flow record in the NetFlow cache contains key fields, as well as additional non-key fields, that can be used later for exporting data to a collection device. Each flow record is created by identifying packets with similar flow characteristics and counting or tracking the packets and bytes per flow.

Flexible NetFlow (FNF) allows you to customize and focus on specific network information. To define a flow, you can use a subset or superset of the traditional seven key fields. FNF also has multiple additional fields (both key and non-key). This permits an organization to target more specific information so that the total amount of information and the number of flows being exported is reduced, allowing enhanced scalability and aggregation.

PROCESS

Configuring Flexible NetFlow for IWAN Monitoring

1. Create flexible NetFlow flow record
2. Create flow exporter
3. Create a flow monitor
4. Apply flow monitor to router interfaces
5. Configure the platform base features

These procedures include best practice recommendations for which key fields and non-key fields need to be collected in order to allow for effective IWAN monitoring.

Additional details regarding the deployment of NetFlow with NBAR2 and the usage of a broad range of NetFlow collector/analyzers are covered in the Application Monitoring Using NetFlow Technology Design Guide.

Procedure 1 Create flexible NetFlow flow record

Flexible NetFlow requires the explicit configuration of a flow record that consists of both key fields and non-key fields. This procedure provides guidance on how to configure a user-defined flow record that includes all of the Traditional NetFlow (TNF) fields (key and non-key) as well as additional FNF fields (key and non-key). The resulting flow record includes the full subset of TNF fields used in classic NetFlow deployments.

The examples in this guide are from Cisco Prime Infrastructure and LiveAction. Different NetFlow collector applications support different export version formats and you should align your flow record with the type of network management platform used by your organization.

Step 1: Specify key fields. This determines unique flow. Be sure to include a separate match statement for each key field.

```
flow record [record name]
 description [record description]
 match [key field type] [key field value]
```

Table 87 Recommended FNF key fields for IWAN

Key field type	Key field value
flow	direction
interface	input
ipv4	tos protocol source address destination address
transport	source port destination port

Step 2: Specify non-key fields to be collected for each unique flow. Be sure to include a separate **collect** statement for each non-key field.

```
flow record [record name]
 collect [non-key field type] [non-key field value]
```

Table 88 Recommended FNF non-key fields for IWAN

Non-key field type	Non-key field value
application	name
flow	sampler
routing	source as destination as next-hop address ipv4
ipv4	source prefix source mask destination mask dscp id
transport	tcp flags
interface	output
counter	bytes packets
timestamp	sys-uptime first sys-uptime last

Example

```
flow record Record-FNF-IWAN
 description Flexible NetFlow for IWAN Monitoring
 match flow direction
 match interface input
 match ipv4 destination address
 match ipv4 protocol
 match ipv4 source address
 match ipv4 tos
 match transport destination-port
 match transport source-port
 collect application name
 collect counter bytes
 collect counter packets
 collect flow sampler
 collect interface output
 collect ipv4 destination mask
 collect ipv4 dscp
 collect ipv4 id
 collect ipv4 source mask
 collect ipv4 source prefix
 collect routing destination as
 collect routing next-hop address ipv4
 collect routing source as
 collect timestamp sys-uptime first
 collect timestamp sys-uptime last
 collect transport tcp flags
```

Procedure 2 Create flow exporter

The NetFlow data that is stored in the cache of the network device can be more effectively analyzed when exported to an external collector.

Creating a flow exporter is only required when exporting data to an external collector. If data is analyzed only on the network device, you can skip this procedure.

Reader Tip

Most external collectors use SNMP to retrieve the interface table from the network device. Ensure that you have completed the relevant SNMP procedures for your platform.

Different NetFlow collector applications support different export version formats (v5, v9, IPFIX) and expect to receive the exported data on a particular UDP or TCP port (ports 2055, 9991, 9995, 9996 are popular). The NetFlow RFC 3954 does not specify a specific port for collectors to receive NetFlow data. In this deployment, the collector applications used for testing use the parameters designated in the following table.

Table 89 *NetFlow collector parameters*

Vendor	Application	Version	Export capability	Netflow destination port
Cisco	Prime Infrastructure	3.0.2	Flexible NetFlow v9	UDP 9991
LiveAction	LiveAction	4.1.2	Flexible NetFlow v9	UDP 2055

Step 1: Configure a basic flow exporter by using Netflow v9.

```
flow exporter [exporter name]
 description [exporter description]
 destination [NetFlow collector IP address]
 source Loopback0
 transport [UDP or TCP] [port number]
 export-protocol netflow
```

Step 2: For FNF records, export the interface table for FNF. The **option interface-table** command enables the periodic sending of an options table. This provides interface names through the NetFlow export.

```
flow exporter [exporter name]
 option interface-table
 template data timeout 600
```

Step 3: If you are using an NBAR flow record, export the NBAR application table. The **option application-table** command enables the periodic sending of an options table that allows the collector to map the NBAR application IDs provided in the flow records to application names.

```
flow exporter [exporter name]
 option application-table
```

Step 4: If you are using an NBAR flow record, export the NBAR application attributes. The **option application-attributes** command causes the periodic sending of NBAR application attributes to the collector.

```
flow exporter [exporter name]
 option application-attributes
```

Step 5: If you are using the Cisco ISR-G2 series routers, enable **output-features**. Otherwise, NetFlow traffic that originates from a WAN remote-site router will not be encrypted or tagged using QoS.

```
flow exporter [exporter name]
 output-features
```

Example: LiveAction

```
flow exporter Export-FNF-Monitor-1
 description FNFv9 NBAR2 with LiveAction
 destination 10.4.48.178
 source Loopback0
 output-features    ! this command is not required on IOS-XE routers
 transport udp 2055
 template data timeout 600
 option interface-table
 option application-table
 option application-attributes
```

Example: Prime Infrastructure

```
flow exporter Export-FNF-Monitor-2
 description FNFv9 NBAR2 with Prime
 destination 10.4.48.36
 source Loopback0
 output-features    ! this command is not required on IOS-XE routers
 transport udp 9991
 template data timeout 600
 option interface-table
 option application-table
 option application-attributes
```

Step 6: Verify the NetFlow exporter configuration using the **show flow exporter** command.

```
RS41-2921# show flow exporter Export-FNF-Monitor-2

Flow Exporter Export-FNF-Monitor-2:
  Description:              FNFv9 NBAR2 with Prime
  Export protocol:          NetFlow Version 9
  Transport Configuration:
    Destination IP address: 10.4.48.36
    Source IP address:      10.255.241.41
    Source Interface:       Loopback0
    Transport Protocol:     UDP
    Destination Port:       9991
    Source Port:            64254
    DSCP:                   0x0
    TTL:                    255
    Output Features:        Used
  Options Configuration:
    interface-table (timeout 600 seconds)
    application-table (timeout 600 seconds)
    application-attributes (timeout 600 seconds)
```

Procedure 3 Create a flow monitor

The network device must be configured to monitor the flows through the device on a per-interface basis. The flow monitor must include a flow record and optionally one or more flow exporters if data is to be collected and analyzed. After the flow monitor is created, it is applied to device interfaces. The flow monitor stores flow information in a cache, and the timer values for this cache are modified within the flow monitor configuration. It is recommended that you set the timeout active timer to 60 seconds, which exports flow data on existing long-lived flows.

Step 1: Create the flow monitor, and then set the cache timers.

```
flow monitor [monitor name]
 description [monitor description]
 cache timeout active 60
 cache timeout inactive 10
```

Step 2: Associate the flow record to the flow monitor. You can use either a custom or a built-in flow record.

```
flow monitor [monitor name]
 record [record name]
```

Step 3: If you are using an external NetFlow collector, associate the exporters to the flow monitor. If you are using multiple exporters, add additional lines.

```
flow monitor [monitor name]
 exporter [exporter name]
```

Example: Prime Infrastructure and LiveAction

```
flow monitor Monitor-FNF-IWAN
 description IWAN Traffic Analysis
 record Record-FNF-IWAN
 exporter Export-FNF-Monitor-1
 exporter Export-FNF-Monitor-2
 cache timeout active 60
 cache timeout inactive 10
```

Step 4: Verify the flow monitor configuration by using the **show flow monitor** command.

```
RS41-2921#show flow monitor
Flow Monitor Monitor-FNF-IWAN:
  Description:       IWAN Traffic Analysis
  Flow Record:       Record-FNF-IWAN
  Flow Exporter:     Export-FNF-Monitor-1
                     Export-FNF-Monitor-2
  Cache:
    Type:                 normal
    Status:               not allocated
    Size:                 4096 entries/0 bytes
    Inactive Timeout:     10 secs
    Active Timeout:       60 secs
    Update Timeout:       1800 secs
    Synchronized Timeout: 600 secs
    Status:            allocated
    Size:              4096 entries/376856 bytes
    Inactive Timeout:  15 secs
    Active Timeout:    60 secs
    Update Timeout:    1800 secs
```

Procedure 4 Apply flow monitor to router interfaces

A best practice for NetFlow in an IWAN deployment is to monitor all inbound and outbound traffic on the DMVPN tunnel interfaces.

Step 1: Apply the flow monitor to the tunnel interface(s).

```
interface [name]
 ip flow monitor [monitor name] input
 ip flow monitor [monitor name] output
```

Example: Single-router dual-link remote site

```
interface Tunnel10
 ip flow monitor Monitor-FNF-IWAN input
 ip flow monitor Monitor-FNF-IWAN output

interface Tunnel11
 ip flow monitor Monitor-FNF-IWAN input
 ip flow monitor Monitor-FNF-IWAN output
```

Step 2: Verify the proper interfaces are configured for NetFlow monitoring using the **show flow interface** command.

```
RS41-2921# show flow interface
Interface Tunnel10
  FNF:  monitor:          Monitor-FNF-IWAN
        direction:        Input
        traffic(ip):      on
  FNF:  monitor:          Monitor-FNF-IWAN
        direction:        Output
        traffic(ip):      on
Interface Tunnel11
  FNF:  monitor:          Monitor-FNF-IWAN
        direction:        Input
        traffic(ip):      on
  FNF:  monitor:          Monitor-FNF-IWAN
        direction:        Output
        traffic(ip):      on
```

Step 3: At dual-router sites with a distribution layer, also apply the flow monitor to the interfaces that connect to the distribution layer switch. This ensures that you capture all possible traffic flows.

Example: First router of a dual-router dual-link remote site

```
interface Port-channel1.50
 ip flow monitor Monitor-FNF-IWAN input
 ip flow monitor Monitor-FNF-IWAN output
```

Example: Second router of a dual-router dual-link remote site

```
interface Port-channel2.54
 ip flow monitor Monitor-FNF-IWAN input
 ip flow monitor Monitor-FNF-IWAN output
```

Step 4: Verify the dscp used in the network by displaying the NetFlow cache on the WAN aggregation routers. Use the **show flow monitor** command.

```
show flow monitor Monitor-FNF-IWAN cache format table
```

The Cisco Prime example below shows the site-to-site topology for RS12, which is a dual router IWAN hybrid model remote site. The remote site graphic shows the logical separation of the MC and BR functions, even though they are physically in the same router.

Figure 37 *Cisco Prime: Site to site topology*

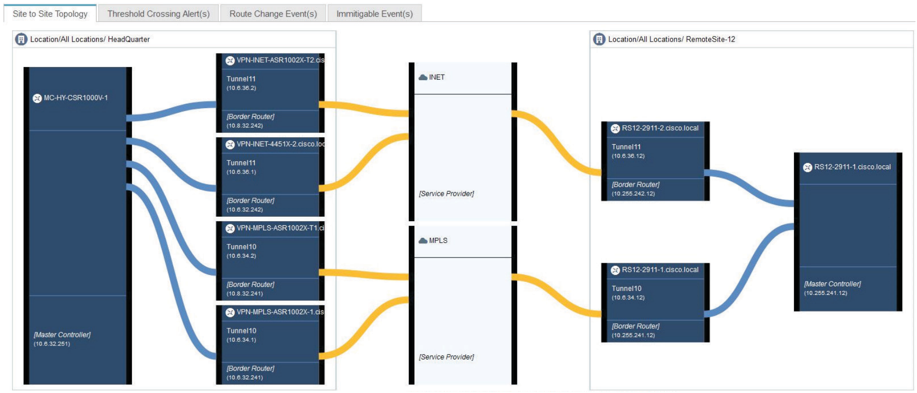

Appendix A: Product List

WAN AGGREGATION

Place In Network	Product Description	Part Number	SW Version	Feature Set
WAN-aggregation Router	Aggregation Services 1002X Router	ASR1002X-5G-VPNK9	IOS-XE 03.16.01a.S	Advanced Enterprise
	Aggregation Services 1001X Router	ASR1001X-5G-VPN	IOS XE 03.16.01a.S	Advanced Enterprise
	Cisco ISR 4451-X Security Bundle with SEC License	ISR4451-X-SEC/K9	IOS XE 03.16.01a.S	securityk9

WAN REMOTE SITE

Place In Network	Product Desccription	Part Number	SW Version	Feature Set
Modular WAN Remote-site Router	Cisco ISR 4451 AX Bundle with APP and SEC License	ISR4451-X-AX/K9	IOS-XE 03.16.01a.S	securityk9, appxk9
	Cisco ISR 4431 AX Bundle with APP and SEC License	ISR4431-AX/K9	IOS XE 03.16.01a.S	securityk9, appxk9
	Cisco ISR 4351 AX Bundle with APP and SEC License	ISR4351-AX/K9	IOS XE 03.16.01a.S	securityk9, appxk9
	Cisco ISR 4331 AX Bundle with APP and SEC License	ISR4331-AX/K9	IOS XE 03.16.01a.S	securityk9, appxk9
	Cisco ISR 4321 AX Bundle with APP and SEC License	ISR4321-AX/K9	IOS XE 03.16.01a.S	securityk9, appxk9
	Cisco ISR 3945 AX Bundle with APP and SEC License	C3945-AX/K9	15.5(3)M1	securityk9, datak9, uck9
	Cisco ISR 3925 AX Bundle with APP and SEC License	C3925-AX/K9	15.5(3)M1	securityk9, datak9, uck9
	Unified Communications Paper PAK for Cisco 3900 Series	SL-39-UC-K9		
	Cisco ISR 2951 AX Bundle with APP and SEC License	C2951-AX/K9	15.5(3)M1	securityk9, datak9, uck9
	Cisco ISR 2921 AX Bundle with APP and SEC License	C2921-AX/K9	15.5(3)M1	securityk9, datak9, uck9
	Cisco ISR 2911 AX Bundle with APP and SEC License	C2911-AX/K9	15.5(3)M1	securityk9, datak9, uck9
	Unified Communications Paper PAK for Cisco 2900 Series	SL-29-UC-K9		
	Cisco ISR 1941 AX Bundle with APP and SEC License	C1941-AX/K9	15.5(3)M1	securityk9, datak9

INTERNET EDGE

Place In Network	Product Description	Part Number	SW Version	Feature Set
Firewall	Cisco ASA 5545-X	ASA5545-K9	ASA 9.1(6)	
	Cisco ASA 5525-X	ASA5525-K9	ASA 9.1(6)	
	Cisco ASA 5515-X	ASA5515-K9	ASA 9.1(6)	
	Cisco ASA 5512-X	ASA5512-K9	ASA 9.1(6)	
	Cisco ASA 5512-X Security Plus license	ASA5512-SEC-PL		
	Firewall Management	ASDM	7.5(2)	

INTERNET EDGE LAN

Place In Network	Product Description	Part Number	SW Version	Feature Set
DMZ Switch	Cisco Catalyst 2960-X Series 24 10/100/1000 PoE and 2 SFP+ Uplink	WS-C2960X-24PS	15.2(3)E1	LAN Base
	Cisco Catalyst 2960-X FlexStack-Plus Hot-Swappable Stacking Module	C2960X-STACK		

LAN ACCESS LAYER

Place In Network	Product Description	Part Number	SW Version	Feature Set
Modular Access Layer Switch	Cisco Catalyst 4500E Series 4507R+E 7-slot Chassis with 48Gbps per slot	WS-C4507R+E	3.7.1E(15.2.3E1)	IP Base
	Cisco Catalyst 4500E Supervisor Engine 8-E, Unified Access, 928Gbps	WS-X45-SUP8-E	3.7.1E(15.2.3E1)	IP Base
	Cisco Catalyst 4500E 12-port 10GbE SFP+ Fiber Module	WS-X4712-SFP+E		
	Cisco Catalyst 4500E 48-Port 802.3at PoE+ 10/100/1000 (RJ-45)	WS-X4748-RJ45V+E		
	Cisco Catalyst 4500E Series 4507R+E 7-slot Chassis with 48Gbps per slot	WS-C4507R+E	3.7.1E(15.2.3E1)	IP Base
	Cisco Catalyst 4500E Supervisor Engine 7L-E, 520Gbps	WS-X45-SUP7L-E	3.7.1E(15.2.3E1)	IP Base
	Cisco Catalyst 4500E 48 Ethernet 10/100/1000 (RJ45) PoE+,UPoE ports	WS-X4748-UPOE+E		
	Cisco Catalyst 4500E 48 Ethernet 10/100/1000 (RJ45) PoE+ ports	WS-X4648-RJ45V+E		

Place In Network	Product Description	Part Number	SW Version	Feature Set
Stackable Access Layer Switch	Cisco Catalyst 3850 Series Stackable 48 Ethernet 10/100/1000 PoE+ ports	WS-C3850-48F	3.7.1E(15.2.3E1)	IP Base
	Cisco Catalyst 3850 Series Stackable 24 Ethernet 10/100/1000 PoE+ Ports	WS-C3850-24P	3.7.1E(15.2.3E1)	IP Base
	Cisco Catalyst 3850 Series 2 x 10GE Network Module	C3850-NM-2-10G		
	Cisco Catalyst 3850 Series 4 x 1GE Network Module	C3850-NM-4-1G		
	Cisco Catalyst 3650 Series 24 Ethernet 10/100/1000 PoE+ and 2x10GE or 4x1GE Uplink	WS-C3650-24PD	3.7.1E(15.2.3E1)	IP Base
	Cisco Catalyst 3650 Series 24 Ethernet 10/100/1000 PoE+ and 4x1GE Uplink	WS-C3650-24PS	3.7.1E(15.2.3E1)	IP Base
	Cisco Catalyst 3650 Series Stack Module	C3650-STACK		
	Cisco Catalyst 2960-X Series 24 10/100/1000 Ethernet and 2 SFP+ Uplink	WS-C2960X-24PD	15.2(3)E1	LAN Base
Standalone Access Layer Switch	Cisco Catalyst 3650 Series 24 Ethernet 10/100/1000 PoE+ and 4x1GE Uplink	WS-C3650-24PS	3.7.1E(15.2.3E1)	IP Base

LAN DISTRIBUTION LAYER

Place In Network	Product Description	Part Number	SW Version	Feature Set
Modular Distribution Layer Virtual Switch Pair	Cisco Catalyst 6500 VSS Supervisor 2T with 2 ports 10GbE and PFC4	VS-S2T-10G	15.2(1)SY1	IP Services
	Cisco Catalyst 6800 Series 6807-XL 7-Slot Modular Chassis	C6807-XL	15.2(1)SY1	IP Services
	Cisco Catalyst 6500 4-port 40GbE/16-port 10GbE Fiber Module w/DFC4	WS-X6904-40G-2T		
	Cisco Catalyst 6500 4-port 10GbE SFP+ adapter for WX-X6904-40G module	CVR-CFP-4SFP10G		
	Cisco Catalyst 6500 CEF720 48 port 10/100/1000mb Ethernet	WS-X6748-GE-TX		
	Cisco Catalyst 6500 Distributed Forwarding Card 4	WS-F6K-DFC4-A		
	Cisco Catalyst 6500 Series 6506-E 6-Slot Chassis	WS-C6506-E	15.2(1)SY1	IP services
	Cisco Catalyst 6500 VSS Supervisor 2T with 2 ports 10GbE and PFC4	VS-S2T-10G	15.2(1)SY1	IP services
	Cisco Catalyst 6500 4-port 40GbE/16-port 10GbE Fiber Module w/DFC4	WS-X6904-40G-2T		
	Cisco Catalyst 6500 4-port 10GbE SFP+ adapter for WX-X6904-40G module	CVR-CFP-4SFP10G		
	Cisco Catalyst 6500 48-port GigE Mod (SFP)	WS-X6748-SFP		
	Cisco Catalyst 6500 Distributed Forwarding Card 4	WS-F6K-DFC4-A		
	Cisco Catalyst 6500 24-port GigE Mod (SFP)	WS-X6724-SFP		
	Cisco Catalyst 6500 Distributed Forwarding Card 4	WS-F6K-DFC4-A		

Place in Network	Product Description	Part Number	SW Version	Feature Set
Extensible Fixed Distribution Layer Virtual Switch Pair	Cisco Catalyst 6800 Series 6880-X Extensible Fixed Aggregation Switch (Standard Tables)	C6880-X-LE	15.2(1)SY1	IP Services
	Cisco Catalyst 6800 Series 6880-X Multi Rate Port Card (Standard Tables)	C6880-X-LE-16P10G		
Modular Distribution Layer Virtual Switch Pair	Cisco Catalyst 4500E Series 4507R+E 7-slot Chassis with 48Gbps per slot	WS-C4507R+E	3.7.1E(15.2.3E1)	Enterprise Services
	Cisco Catalyst 4500E Supervisor Engine 7-E, 848Gbps	WS-X45-SUP7-E	3.7.1E(15.2.3E1)	Enterprise Services
	Cisco Catalyst 4500E 12-port 10GbE SFP+ Fiber Module	WS-X4712-SFP+E		
	Cisco Catalyst 4500E 48-Port 802.3at PoE+ 10/100/1000 (RJ-45)	WS-X4748-RJ45V+E		
Fixed Distribution Layer Virtual Switch Pair	Cisco Catalyst 4500-X Series 32 Port 10GbE IP Base Front-to-Back Cooling WS-C4500X-32SFP+ 3.5.3E(15.2.1E3) Enterprise Services			
Stackable Distribution Layer Switch	Cisco Catalyst 3850 Series Stackable Switch with 12 SFP Ethernet	WS-C3850-12S	3.7.1E(15.2.3E1)	IP Services
	Cisco Catalyst 3850 Series 4 x 1GE Network Module	C3850-NM-4-1G		
	Cisco Catalyst 3850 Series 2 x 10GE Network Module	C3850-NM-2-10G		

Appendix B: Technical Feature Supplement

FRONT DOOR VRF FOR DMVPN

Building an IPsec tunnel requires reachability between the crypto routers. When you use the Internet, routers use a default route to contact their peers.

Figure 38 *IPsec tunnel*

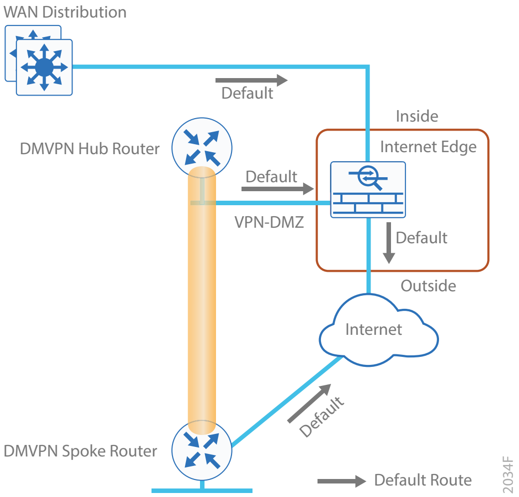

If you need to extend the internal network and the same default routing options that are available to internal users, you must advertise a default route to the VPN hub router. For details, see section A in the following figure.

Figure 39 *IPsec tunnel before/after default route injection*

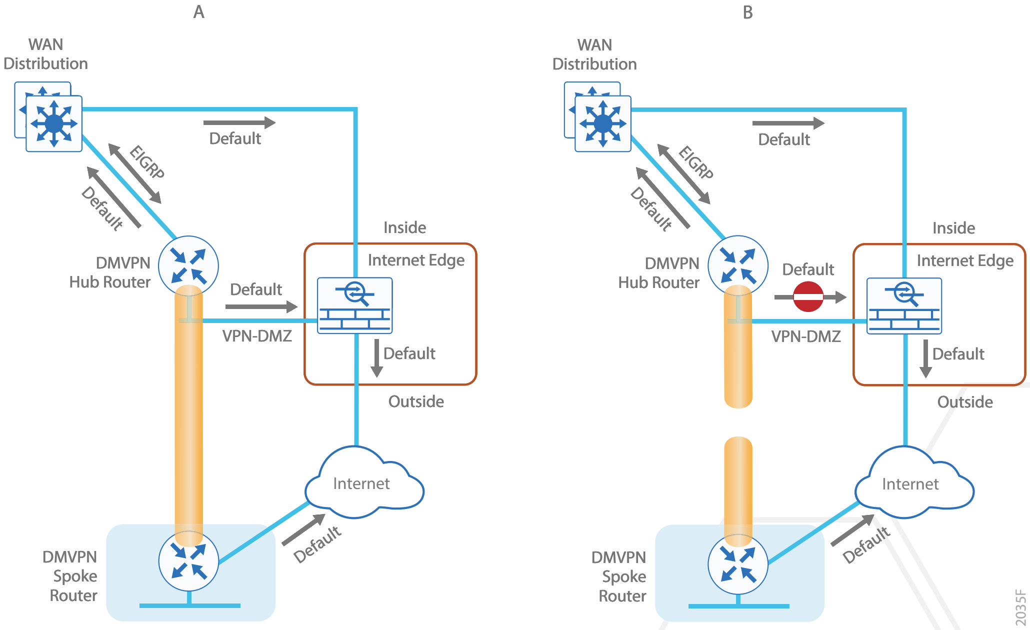

The advertisement of a default route to the hub router (with an existing default route) is problematic. This route requires a better administrative distance to become the active default, which then overrides the default route that is supporting the peer-peer IPsec tunnel connection. This routing advertisement breaks the tunnel as shown in section B in the previous figure.

Through the introduction of an external VRF INET-PUBLIC (shown in red), the hub router can support multiple default routes. The internal network remains in the global VRF. This is shown in section A of the following figure.

Tech Tip

Most additional features on the hub router do not require VRF-awareness.

Figure 40 *IPsec tunnel with FVRF aggregation*

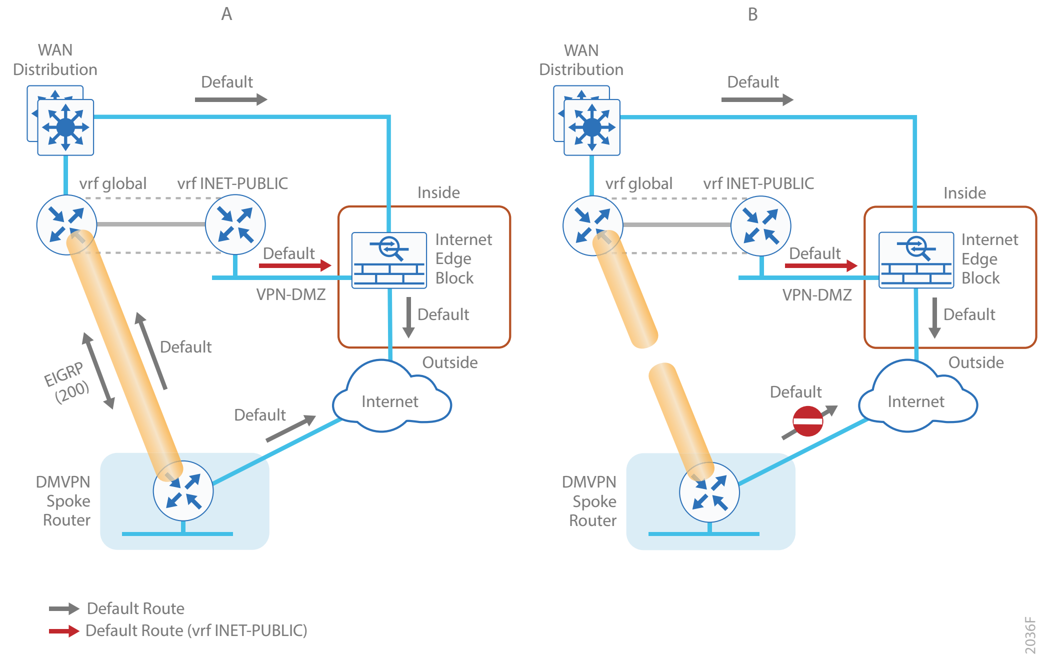

This configuration is referred to as *FVRF*, because the Internet is contained in a VRF. FVRF is sometimes referred to as *Front Side VRF*. The alternative to this design is Inside VRF (IVRF), where the internal network is in a VRF on the VPN hub and the Internet remains in the global VRF. This method is not documented in this guide.

It is now possible to reestablish the IPsec tunnel to the remote peer router. Because the remote-site policy requires central Internet access for end users, a default route is advertised through the tunnel. This advertisement causes a similar default routing issue on the remote router; the tunnel default overrides the Internet-pointing default and the tunnel connection breaks as shown in section B of the previous figure.

This configuration requires using FVRF on the remote-site router as well. The primary benefits of using this solution are:

- Simplified default routing and static default routes in the INET-PUBLIC VRFs.
- Ability to support default routing for end-users traffic through VPN tunnels.
- Ability to use dynamic default routing for sites with multiple WAN transports.
- Ability to build spoke-to-spoke tunnels with DMVPN with end-user traffic routed by default through VPN tunnels.

The final design that uses FVRF at both the WAN-aggregation site and a WAN remote-site is shown in the following figure.

Figure 41 FVRF: Final configuration

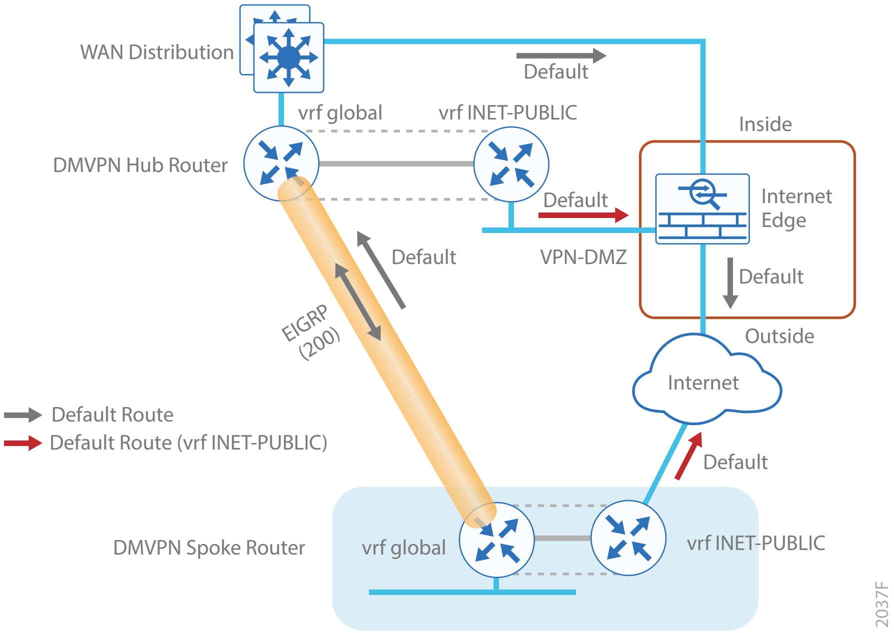

Appendix C: Common Procedures

The procedure in this Appendix is common to all routers.

Procedure 1 Configure the platform base features

Step 1: Configure the device host name. Make it easy to identify the device.

```
hostname [Hostname]
```

Step 2: Configure local login and password.

The local login account and password provide basic access authentication to a router, which provides only limited operational privileges. The enable password secures access to the device configuration mode. By enabling password encryption, you prevent the disclosure of plain text passwords when viewing configuration files.

```
username admin secret c1sco123
enable secret c1sco123
service password-encryption
aaa new-model
```

By default, https access to the router will use the enable password for authentication.

Step 3: (Optional) Configure centralized user authentication.

As networks scale in the number of devices to maintain, it poses an operational burden to maintain local user accounts on every device. A centralized AAA service reduces operational tasks per device and provides an audit log of user access for security compliance and root cause analysis. When AAA is enabled for access control, AAA controls all management access to the network infrastructure devices (SSH and HTTPS).

TACACS+ is the primary protocol used to authenticate management logins on the infrastructure devices to the AAA server. A local AAA user database is also defined (in Step 2) on each network infrastructure device in order to provide a fallback authentication source in case the centralized TACACS+ server is unavailable.

```
tacacs server TACACS-SERVER-1
 address ipv4 10.4.48.15
 key SecretKey

aaa group server tacacs+ TACACS-SERVERS
 server name TACACS-SERVER-1

aaa authentication login default group TACACS-SERVERS local
aaa authorization exec default group TACACS-SERVERS local
aaa authorization console
ip http authentication aaa
```

Step 4: Configure device management protocols.

Secure HTTPS and SSH are secure replacements for the HTTP and Telnet protocols. They use SSL and TLS in order to provide device authentication and data encryption.

Secure management of the network device is enabled through the use of the SSH and HTTPS protocols. Both protocols are encrypted for privacy and the nonsecure protocols, Telnet and HTTP, are turned off. SCP is enabled, which allows the use of code upgrades using Prime Infrastructure via SSH-based SCP protocol.

Specify the transport **preferred none** on vty lines in order to prevent errant connection attempts from the CLI prompt. Without this command, if the ip name-server is unreachable, long timeout delays may occur for mistyped commands.

```
ip domain-name cisco.local
ip ssh version 2
no ip http server
ip http secure-server
ip scp server enable
line vty 0 15
 transport input ssh
 transport preferred none
```

When synchronous logging of unsolicited messages and debug output is turned on, console log messages are displayed on the console after interactive CLI output is displayed or printed. With this command, you can continue typing at the device console when debugging is enabled.

```
line con 0
 transport preferred none
 logging synchronous
```

Enable SNMP in order to allow the network infrastructure devices to be managed by an NMS. SNMPv2c is configured both for a read-only and a read-write community string.

```
snmp-server community cisco RO
snmp-server community cisco123 RW
snmp-server ifindex persist ! IOS Classic Only
snmp ifmib ifindex persist  ! IOS XE Only
```

Step 5: (Optional) In networks where network operational support is centralized, you can increase network security by using an access list to limit the networks that can access your device. In this example, only devices on the 10.4.48.0/24 network will be able to access the device via SSH or SNMP.

```
access-list 55 permit 10.4.48.0 0.0.0.255
line vty 0 15
 access-class 55 in

snmp-server community cisco RO 55
snmp-server community cisco123 RW 55
```

Tech Tip

If you configure an access-list on the vty interface you may lose the ability to use ssh to login from one router to the next for hop-by-hop troubleshooting.

Step 6: Configure a synchronized clock.

NTP is designed to synchronize a network of devices. An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP then distributes this time across the organizations network.

You should program network devices to synchronize to a local NTP server in the network. The local NTP server typically references a more accurate clock feed from an outside source. By configuring console messages, logs, and debug output to provide time stamps on output, you can cross-reference events in a network.

```
ntp server 10.4.48.17


clock timezone PST -8
clock summer-time PDT recurring


service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
```

Step 7: Configure an in-band management interface.

The *loopback interface* is a logical interface that is always reachable as long as the device is powered on and any IP interface is reachable to the network. Because of this capability, the loopback address is the best way to manage the router in-band. Layer 3 process and features are also bound to the loopback interface to ensure process resiliency.

The loopback address is commonly a host address with a 32-bit address mask. Allocate the loopback address from the IP address block that the router summarizes to the rest of the network.

```
interface Loopback 0
  ip address [ip address] 255.255.255.255
  ip pim sparse-mode
```

Bind the device processes for SNMP, SSH, PIM, TACACS+ and NTP to the loopback interface address for optimal resiliency:

```
snmp-server trap-source Loopback0
ip ssh source-interface Loopback0
ip pim register-source Loopback0
ip tacacs source-interface Loopback0
ntp source Loopback0
```

Appendix D: Changes

This appendix summarizes the changes Cisco made to this guide since its last edition.

- We upgraded IOS software.
- We updated the QoS settings.
- We updated the EIGRP settings.
- We updated the NHRP settings.
- We updated the PfR policy settings.
- We simplified the IOS CA configuration.
- We added a virtual router hub master controller.
- We added hub master controller high availability.
- We added hub border router scalability.
- We added a second data center as a transit site.
- We added Prime Infrastructure for IWAN monitoring.

Appendix E: Device Configuration Files

To view the configuration files from the CVD lab devices that were used to test this guide, go to the following URL: http://cvddocs.com/fw/201i-16a

Please use the feedback form to send comments and suggestions about this guide.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2015 Cisco Systems, Inc. All rights reserved.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)