



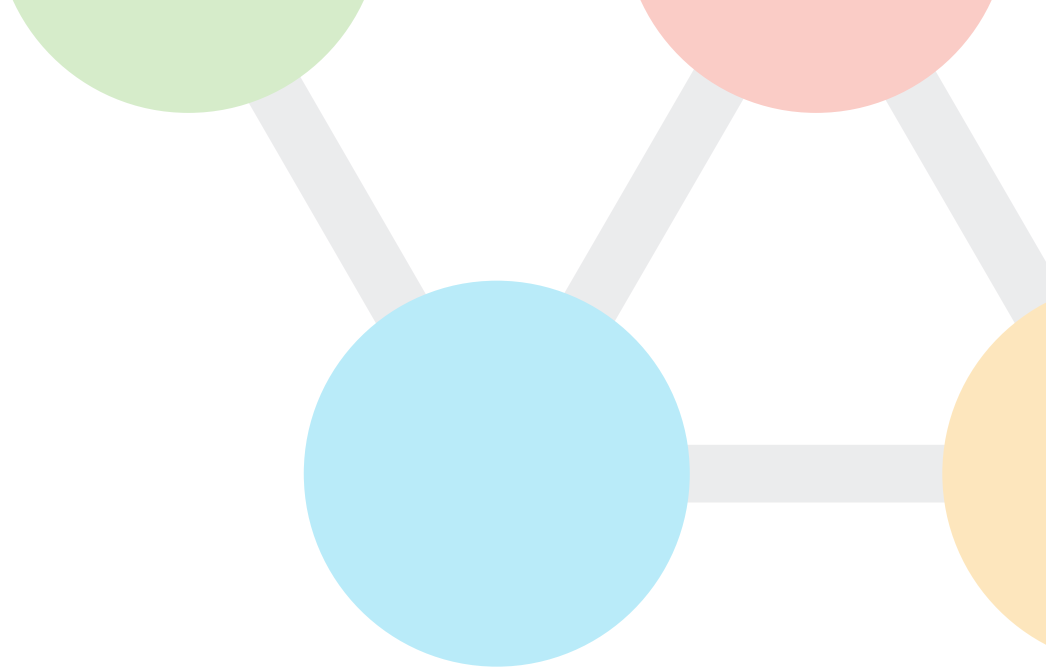
ARCHIVED DOCUMENT

This document is archived and should only be used as a historical reference and should not be used for new deployments for one of the following reasons:

- SD-WAN guides are the recommended alternative.
- This document is outdated. There are no plans to update the content.

For the latest guides, please refer to:

<https://cisco.com/go/cvd>



CISCO VALIDATED DESIGN

Intelligent WAN Multiple Data Center Deployment Guide

September 2017



Table of Contents

Deploying the Cisco Intelligent WAN.....	1
Deployment Details	1
Deploying Multiple Data Centers.....	2
Configuring Transit Border Routers	5
Configuring Transit Master Controller	23
Configuring PfR for Transit Location	27
Configuring Remote-Site Routers for Transit Site BRs	31
Deploying a Second DC without a DCI Link.....	38
Configuring Hub and Transit Border Routers	38
Deploying IWAN Quality of Service	42
Applying DMVPN QoS Policy to DMVPN Hub Routers.....	42
Applying QoS Configurations to Remote Site Routers.....	47
Appendix A: Product List.....	49
Appendix B: Changes.....	50

Deploying the Cisco Intelligent WAN

This guide is one in a series of IWAN advanced deployment guides that focus on how to deploy the advanced features of the Cisco Intelligent WAN (IWAN). These guides build on the configurations deployed in the [Intelligent WAN Deployment Guide](#) and are optional components of its base IWAN configurations.

The advanced guides are as follows:

- [IWAN High Availability and Scalability Deployment Guide](#)
- [IWAN Multiple Data Center Deployment Guide](#) (this guide)
- [IWAN Multiple Transports Deployment Guide](#)
- [IWAN Multiple VRF Deployment Guide](#)
- [IWAN Public Key Infrastructure Deployment Guide](#)
- [IWAN NetFlow Monitoring Deployment Guide](#)
- [IWAN Remote Site 4G LTE Deployment Guide](#)

For design details, see [Intelligent WAN Design Summary](#).

For configuration details, see [Intelligent WAN Configuration Files Guide](#).

For an automated way to deploy IWAN, use the APIC-EM IWAN Application. For more information, see the [Cisco IWAN Application on APIC-EM User Guide](#).

If want to use TrustSec with your IWAN deployment, see “Configuring SGT Propagation” in the [User-to-Data-Center Access Control Using TrustSec Deployment Guide](#).

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
```

Deploying Multiple Data Centers

Use this guide to deploy a second data center location as a transit site for geographic redundancy and scalability. This concept works with any of the IWAN design models.

This type of configuration offers the following benefits:

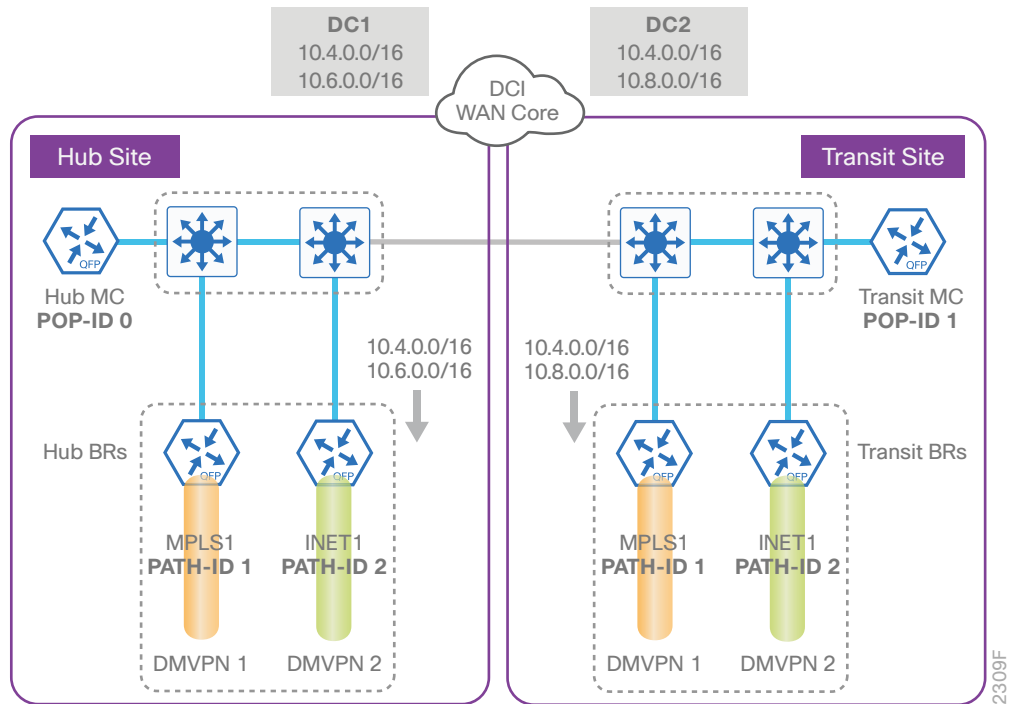
- Data centers are reachable across the WAN core for each transit site using a Data Center Interconnect (DCI).
- Remote sites can access any data center across either hub.
- Data centers can reach any remote site across any of the transit sites.
- Multiple hub BRs per DMVPN per site may be required for horizontal scaling, as noted in the [Intelligent WAN High Availability and Scalability Deployment Guide](#).
- Non-DCI designs are also supported, with additional routing protocol commands on hub and transit site BRs.

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.

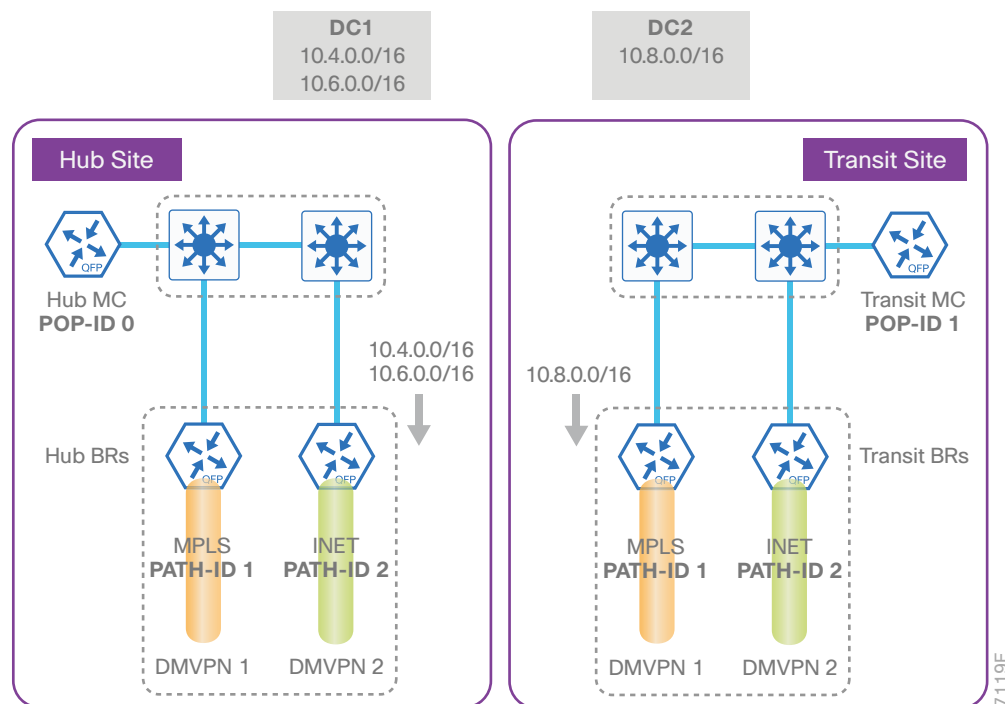
The following diagram shows the transit MC with two additional transit BRs and where they fit into the IWAN hybrid design model using a DCI link.

Figure 1 IWAN hybrid design model—Second data center as a transit site with DCI



The following diagram shows the transit MC with two additional transit BRs when not using a DCI link.

Figure 2 IWAN hybrid design model—Second data center without DCI



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 Cisco Performance Routing (PFR) uses the primary data center as its preference for all traffic.

If the MPLS1 path is primary and INET1 path is secondary in your design, the path preference will be as follows:

- Path #1 to 10.4.0.0/16 is MPLS1 path to DC#1
- Path #2 to 10.4.0.0/16 is INET1 path to DC#1
- Path #3 to 10.4.0.0/16 is MPLS1 path to DC#2
- Path #4 to 10.4.0.0/16 is INET1 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 MPLS1 path to DC#1
- Path #2 to 10.4.0.0/16 is MPLS1 path to DC#2
- Path #3 to 10.4.0.0/16 is INET1 path to DC#1
- Path #4 to 10.4.0.0/16 is INET1 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. Configure the routing protocol for the LAN
5. Connect to the MPLS WAN or Internet
6. Configure the mGRE tunnel
7. Configure the routing protocol for the WAN
8. 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 1 Path and IP addresses for hub BRs

Host name	Path	Path ID	Loopback IP address	Port-channel IP address	MPLS/Internet DMZ IP address
HY-MPLS1-ASR1002X-1	MPLS1	1	10.6.32.241/32	10.6.32.2/30	192.168.6.1/24
HY-INET1-ASR1002X-2	INET1	2	10.6.32.242/32	10.6.32.6/30	192.168.146.10/24

Table 2 Path and IP addresses for transit BRs

Host name	Path	Path ID	Loopback IP address	Port-channel IP address	MPLS/Internet DMZ IP address
HY-MPLS1-CSR1000v-T1	MPLS1	1	10.8.32.241/32	10.8.32.2/30	192.168.6.41/24
HY-INET1-CSR1000v-T2	INET1	2	10.8.32.242/32	10.8.32.6/30	192.168.146.11/24

Reader Tip

Whenever IWAN is designed with WAAS leveraging AppNav, please ensure that the Loopback IP address that is being used for PfR is not also used as the AppNav Service Controller address. This is applicable for any Hub IWAN router that is part of an AppNav Cluster.

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.

```
copy running-config ftp://cisco:cisco@10.4.48.27
Address or name of remote host [10.4.48.27]?
Destination filename [hy-mpls1-asr1002x-1-config]?
Writing hy-mpls1-asr1002x-1-config !
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.

You can also make the changes below in a text editor before pasting the configuration into the router.

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 HY-MPLS1-CSR1000v-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
```

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.

```
interface Port-channel1
  description IWAN-D3750X-T
  ip address 10.8.32.2 255.255.255.252
  ip pim sparse-mode
  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 GigabitEthernet1
  description IWAN-D3750X-T Gig1/0/1

interface GigabitEthernet2
  description IWAN-D3750X-T Gig2/0/1

interface range GigabitEthernet1, GigabitEthernet2
  no ip address
  cdp enable
  channel-group 1
  no shutdown
```

Procedure 4 Configure the routing protocol for the LAN

If you are planning to use EIGRP, choose option 1. If you are planning to use BGP on the WAN and OSPF on the LAN, choose option 2.

Option 1: EIGRP on the LAN

The following table shows the EIGRP LAN delay in use.

Table 3 EIGRP LAN delay for IWAN transit routers

LAN Interface	EIGRP LAN Delay (10 usec)
All LAN	50000

Step 1: 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
  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
```

Step 2: Configure the EIGRP interface.

Allow EIGRP to form neighbor relationships across the interface in order 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
```

Step 3: Configure the throughput delay on the LAN interface.

At the hub location where there are multiple border routers, the interface throughput delay setting should be set to influence the EIGRP routing protocol path preference.

Tech Tip

If you are using Port-channel interfaces with two Gigabit Ethernet members as recommended in this guide, you will have to double the LAN path delay to 500000 microseconds (usec), instead of the standard IWAN setting of 250000.

Set the internal LAN path to 500000 microseconds (usec). The delay command is entered in 10 usec units.

```
interface Port-channel1
  delay 50000
```

Option 2: OSPF on the LAN

Step 1: Configure OSPF Area 0 by using the loopback interface IP address as the router-id.

```
router ospf 100
  router-id 10.8.32.241
```

Step 2: Remove passive interface for the LAN interface.

```
router ospf 100
  no passive-interface Port-channel1
```

Procedure 5 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 GigabitEthernet3
  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

Step 1: The DMVPN design is using FVRF, so you must place the WAN interface into the VRF configured in Procedure 3, "Configure the WAN-facing VRF."

Step 2: Enable the interface, select the VRF, and assign the IP address.

```
interface GigabitEthernet3
  vrf forwarding IWAN-TRANSPORT-2
  ip address 192.168.146.11 255.255.255.0
  no shutdown
```

Step 3: 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 6 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 4 DMVPN tunnel parameters for transit BRs

Hostname	Tunnel type	Tunnel number	Tunnel IP address
HY-MPLS1-CSR1000v-T1	MPLS1	100	10.6.34.2/23
HY-INET1-CSR1000v-T2	INET1	200	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 Tunnel100
  ip address 10.6.34.2 255.255.254.0
```

Procedure 7 Configure the routing protocol for the WAN

If you are planning to use EIGRP, choose option 1. If you are planning to use BGP on the WAN and OSPF on the LAN, choose option 2.

Option 1: EIGRP on the WAN

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 Tunnel100
      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
    exit-af-interface
```

Step 2: Configure EIGRP summary metrics.

Step 3: If there are many component routes to be summarized and the component routes are frequently updated, the metrics are also updated frequently, which may cause a spike in the CPU usage. The **summary-metric** command explicitly sets the metric for the summary regardless of the component route metric, which reduces the computational load on a router.

The first value is the bandwidth metric in Kbits per second. The second value is the delay metric in 10 usecs. The third value is the reliability metric where 255 is 100% reliable. The fourth value is the effective bandwidth metric (loading) where 255 is 100% loaded. The fifth value is the MTU of the path.

Tech Tip

EIGRP uses the path's minimum bandwidth as part of the metric calculation. The path's minimum bandwidth is defined in a route advertisement in the minimum bandwidth path attribute. Setting the summary metric bandwidth (first value) to 10 Mbps essentially removes the ability to differentiate between a 10 Mbps tunnel (MPLS1) and a 100 Mbps circuit (INET1) because both paths have a minimum bandwidth of 10 Mbps. Setting the summary metric bandwidth to 10 Gbps as recommended in this guide allows the calculations on the branch router to differentiate tunnel bandwidth regardless of the size of each path.

Use the identical values for each summary address defined in the previous step.

```
router eigrp IWAN-EIGRP
  address-family ipv4 unicast autonomous-system 400
  topology base
    summary-metric 10.6.0.0/16 10000000 10000 255 1 1500
    summary-metric 10.7.0.0/16 10000000 10000 255 1 1500
    summary-metric 10.8.0.0/16 10000000 10000 255 1 1500
    summary-metric 10.255.240.0/21 10000000 10000 255 1 1500
  exit-af-topology
```

Step 4: Configure the throughput delay on the tunnel interface.

The tunnel interface throughput delay setting should be set to influence the EIGRP routing protocol path preference. Set the primary WAN path to 10000 usec and the secondary WAN path to 20000 usec to prefer one over the other. The delay command is entered in 10 usec units.

```
interface Tunnel100
  delay 1000
```


Step 5: Tag the routes for data center (POP) affinity.

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 POP site, tags are required in order to identify the different DMVPN hub router locations, which allows a remote site to prefer one POP over the other.

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

The following tables show specific route tags in use.

Table 5 Route tag information for hub BRs at POP1

DMVPN hub	DMVPN tunnel key	Tag tunnel
HY-MPLS1-ASR1002X-1	1100 (MPLS1)	1100 (All routes)
HY-INET1-ASR1002X-2	1200 (INET1)	1200 (All routes)

Table 6 Route tag information for transit BRs at POP2

DMVPN hub	DMVPN tunnel key	Tag tunnel
HY-MPLS1-CSR1000v-T1	1100 (MPLS1)	1600 (All routes)
HY-INET1-CSR1000v-T2	1200 (INET1)	1700 (All routes)

The following examples show the hub and transit border routers in the IWAN hybrid design model.

Example: POP1 MPLS hub border router–HY-MPLS1-ASR1002X-1

```

route-map SET-TAG-ALL permit 10
  description Tag all routes advertised through the tunnel
  set tag 1100

router eigrp IWAN-EIGRP
  address-family ipv4 unicast autonomous-system 400
  topology base
  distribute-list route-map SET-TAG-ALL out Tunnel100

```

Example: POP1 INET hub border router–HY-INET1-ASR1002X-2

```
route-map SET-TAG-ALL permit 10
description Tag all routes advertised through the tunnel
set tag 1200

router eigrp IWAN-EIGRP
address-family ipv4 unicast autonomous-system 400
topology base
distribute-list route-map SET-TAG-ALL out Tunnel200
```

Example: POP2 MPLS transit border router–HY-MPLS1-CSR1000v-T1

```
route-map SET-TAG-ALL permit 10
description Tag all routes advertised through the tunnel
set tag 1600

router eigrp IWAN-EIGRP
address-family ipv4 unicast autonomous-system 400
topology base
distribute-list route-map SET-TAG-ALL out Tunnel100
```

Example: POP2 INET transit border router–HY-INET1-CSR1000v-T2

```
route-map SET-TAG-ALL permit 10
description Tag all routes advertised through the tunnel
set tag 1700

router eigrp IWAN-EIGRP
address-family ipv4 unicast autonomous-system 400
topology base
distribute-list route-map SET-TAG-ALL out Tunnel200
```

Option 2: BGP on the WAN

The following table shows the tunnel DMVPN IP subnets, local preferences, community strings, and metrics in use.

Table 7 Tunnel IPs, local preferences, community strings, and metrics for hub BRs

DMVPN hub router	DMVPN Tunnels	BGP local preference	BGP community string	OSPF metric preferred POP	OSPF metric secondary POP
HY-MPLS1-ASR1002X-1	10.6.34.0/23	800 (MPLS1)	65100:100	1000	2000
HY-INET1-ASR1002X-2	10.6.36.0/23	780 (INET1)	65100:200	1200	2200

Table 8 Tunnel IPs, local preferences, community strings, and metrics for transit BRs

DMVPN hub router	DMVPN Tunnels	BGP local preference	BGP community string	OSPF metric preferred POP	OSPF metric secondary POP
HY-MPLS1-CSR1000v-T1	10.6.34.0/23	600 (MPLS1)	65100:101	1000	2000
HY-INET1-CSR1000v-T2	10.6.36.0/23	580 (INET1)	65100:202	1200	2200

Step 1: Configure BGP values for the tunnel interface.

Use a private AS number for the BGP process. Assign this router's loopback address as the BGP router-id. Log the neighbor changes. Create a listen range that includes the subnet range of the tunnel interface. For internal BPG, use the same AS number for the remote sites. Create the route reflector and use the tunnel as the update source interface. Adjust the BGP hello and hold timers to 20 seconds and 60 seconds, respectively.

```
router bgp 65100
  bgp router-id 10.8.32.241
  bgp log-neighbor-changes
  bgp listen range 10.6.34.0/23 peer-group MPLS1-SPOKES
  neighbor MPLS1-SPOKES peer-group
  neighbor MPLS1-SPOKES remote-as 65100
  neighbor MPLS1-SPOKES description MPLS1 Spoke Route Reflector
  neighbor MPLS1-SPOKES update-source Tunnel100
  neighbor MPLS1-SPOKES timers 20 60
```

Step 2: Create the static null routes for the enterprise summary prefix and the site-specific prefixes.

```
ip route 10.4.0.0 255.252.0.0 Null0 254
ip route 10.8.0.0 255.255.0.0 Null0 254
ip route 10.4.0.0 255.255.0.0 Null0 254
```

Step 3: Configure the BGP address family.

Define the network statements for the default network, the enterprise summary prefix, the site-specific prefixes, and the local MC loopback IP address the router will advertise to the remote sites. Configure BGP dynamic neighbors for the remote sites. Set the BGP distance and redistribute the internal networks.

```
router bgp 65100
  address-family ipv4
    bgp redistribute-internal
    network 0.0.0.0
    network 10.4.0.0 mask 255.252.0.0
    network 10.4.0.0 mask 255.255.0.0
    network 10.8.0.0 mask 255.255.0.0
    network 10.8.32.251 mask 255.255.255.255
    neighbor MPLS1-SPOKES activate
    neighbor MPLS1-SPOKES send-community
    neighbor MPLS1-SPOKES route-reflector-client
    neighbor MPLS1-SPOKES next-hop-self all
    neighbor MPLS1-SPOKES weight 50000
    neighbor MPLS1-SPOKES soft-reconfiguration inbound
    distance bgp 201 19 200
  exit-address-family
```

Step 4: Create the prefix lists for BGP.

Define the prefix-lists for the default network, the enterprise summary prefix, the site-specific prefixes, the local MC loopback IP address, and the subnet ranges for the DMVPN tunnels.

```
ip prefix-list DEFAULT-ROUTE seq 10 permit 0.0.0.0/0
ip prefix-list ENTERPRISE-PREFIX seq 10 permit 10.4.0.0/14
ip prefix-list LOCALDC-PREFIX seq 10 permit 10.4.0.0/16
ip prefix-list LOCALDC-PREFIX seq 20 permit 10.8.0.0/16
ip prefix-list LOCALMCLOOPBACK seq 10 permit 10.8.32.251/32
ip prefix-list TUNNEL-DMVPN seq 10 permit 10.6.34.0/23
```

Step 5: Create and apply the prefix route maps for BGP.

Define the route map to block prefixes inbound on the tunnel interface. Define the route map to allow prefixes to go out on the tunnel interface. Set the local preference and the community string for this DMVPN hub router. Apply the route maps to the BGP address family. Configure BGP to display communities in the format AA:NN.

Example: MPLS transit border router–HY-MPLS1-CSR1000v-T1

```
ip bgp-community new-format

route-map MPLS1-IN deny 10
  description All Blocked Prefixes to come IN on BGP
  match ip address prefix-list DEFAULT-ROUTE ENTERPRISE-PREFIX LOCALDC-PREFIX LO-  
CALMCLOOPBACK TUNNEL-DMVPN

route-map MPLS1-IN permit 1000
  description Allow Everything Else

route-map MPLS1-OUT permit 10
  description All Allowed Prefixes to Go OUT on BGP to Spokes
  match ip address prefix-list DEFAULT-ROUTE ENTERPRISE-PREFIX LOCALDC-PREFIX LO-  
CALMCLOOPBACK
  set local-preference 600
  set community 65100:101

router bgp 65100
  address-family ipv4
    neighbor MPLS1-SPOKES route-map MPLS1-IN in
    neighbor MPLS1-SPOKES route-map MPLS1-OUT out
  exit-address-family
```

Example: INET transit border router–HY-INET1-CSR1000v-T2

```
ip bgp-community new-format

route-map INET1-IN deny 10
  description All Blocked Prefixes to come IN on BGP
  match ip address prefix-list DEFAULT-ROUTE ENTERPRISE-PREFIX LOCALDC-PREFIX LO-  
CALMCLOOPBACK TUNNEL-DMVPN

route-map INET1-IN permit 1000
```

```

description Allow Everything Else

route-map INET1-OUT permit 10
description All Allowed Prefixes to Go OUT on BGP to Spokes
match ip address prefix-list DEFAULT-ROUTE ENTERPRISE-PREFIX LOCALDC-PREFIX LO-  
CALMCLOOPBACK
set local-preference 580
set community 65100:201

router bgp 65100
address-family ipv4
neighbor INET1-SPOKES route-map INET1-IN in
neighbor INET1-SPOKES route-map INET1-OUT out
exit-address-family

```

Step 6: Create and apply the BGP to OSPF redistribution route map for hub BRs.

When there are two or more POP sites, there might be certain remote sites that want to prefer one POP over the other. This preference choice is done using a community string value, which is sent by the remote site router to indicate which POP they prefer.

This example uses a community string in the form of AS:NN with AS being the BGP autonomous system number and NN being the value that selects the preferred POP.

Example:

```
65100:10 to prefer POP 1 (hub site)
```

```
65100:20 to prefer POP 2 (transit site)
```

The hub and transit BRs use the community string value they receive from the remote site to determine the OSPF metric for each location.

Define the community list to classify the remote sites as preferring POP1 or POP 2. Define the route map to block null routes from being distributed into OSPF. Set the metric to the appropriate value for the POP chosen by the remote site community string value. Apply the route map to the OSPF process when redistributing BGP.

The hub location matches the POP2 community string to set the higher metric values.

Example: POP1 MPLS1 border router–HY-MPLS1-ASR1002X-1

```

ip community-list standard POP2-SPOKES permit 65100:20

route-map REDIST-BGP-TO-OSPF permit 10
description Secondary POP2 with higher Metric
match community POP2-SPOKES
set metric 2000

```

```
set metric-type type-1

route-map REDIST-BGP-TO-OSPF deny 20
description Block Null routes to be distributed from BGP to OSPF
match ip address prefix-list DEFAULT-ROUTE ENTERPRISE-PREFIX LOCALDC-PREFIX

route-map REDIST-BGP-TO-OSPF permit 1000
description Prefer POP1 with lower Metric
set metric 1000
set metric-type type-1

router ospf 100
redistribute bgp 65100 subnets route-map REDIST-BGP-TO-OSPF
```

Example: POP1 INET1 border router–HY-INET1-ASR1002X-2

```
ip community-list standard POP2-SPOKES permit 65100:20

route-map REDIST-BGP-TO-OSPF permit 10
description Secondary POP2 with higher Metric
match community POP2-SPOKES
set metric 2200
set metric-type type-1

route-map REDIST-BGP-TO-OSPF deny 20
description Block Null routes to be distributed from BGP to OSPF
match ip address prefix-list DEFAULT-ROUTE ENTERPRISE-PREFIX LOCALDC-PREFIX

route-map REDIST-BGP-TO-OSPF permit 1000
description Prefer POP1 with lower Metric
set metric 1200
set metric-type type-1

router ospf 100
redistribute bgp 65100 subnets route-map REDIST-BGP-TO-OSPF
```

Step 7: Create and apply the updated BGP to OSPF redistribution route map for transit BRs.

The POP preference route map changes from the previous step have to be applied to the corresponding transit BRs at your POP2 location.

The transit location matches the POP1 community string to set the higher metric values.

Example: POP2 MPLS1 border router–HY-MPLS1-CSR1000v-T1

```
ip community-list standard POP1-SPOKES permit 65100:10

route-map REDIST-BGP-TO-OSPF permit 10
  description Secondary POP1 with higher Metric
  match community POP1-SPOKES
  set metric 2000
  set metric-type type-1

route-map REDIST-BGP-TO-OSPF deny 20
  description Block Null routes to be distributed from BGP to OSPF
  match ip address prefix-list DEFAULT-ROUTE ENTERPRISE-PREFIX LOCALDC-PREFIX

route-map REDIST-BGP-TO-OSPF permit 1000
  description Prefer POP2 with lower Metric
  set metric 1000
  set metric-type type-1

router ospf 100
  redistribute bgp 65100 subnets route-map REDIST-BGP-TO-OSPF
```

Example: POP2 INET1 border router–HY-INET1-CSR1000v-T2

```
ip community-list standard POP1-SPOKES permit 65100:10

route-map REDIST-BGP-TO-OSPF permit 10
  description Secondary POP1 with higher Metric
  match community POP1-SPOKES
  set metric 2200
  set metric-type type-1

route-map REDIST-BGP-TO-OSPF deny 20
  description Block Null routes to be distributed from BGP to OSPF
```



```

match ip address prefix-list DEFAULT-ROUTE ENTERPRISE-PREFIX LOCALDC-PREFIX

route-map REDIST-BGP-TO-OSPF permit 1000
description Prefer POP2 with lower Metric
set metric 1200
set metric-type type-1

router ospf 100
redistribute bgp 65100 subnets route-map REDIST-BGP-TO-OSPF

```

Procedure 8 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 9 DMVPN NAT address mapping

Hostname	DMVPN hub router DMZ address	DMVPN hub router public address (externally routable after NAT)
HY-INET1-CSR1000v-T2	192.168.146.13	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 10 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 11 Private DMZ firewall network objects

Network object name	Object type	IP address	Description
dmz-dmvpn-T2	Host	192.168.146.13	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: 192.168.146.13)

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-dm-vpn-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
4. Configure the routing protocol on 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 12 Hub and transit site MC IP addresses

Host name	Pop ID	Loopback IP address	Port-channel IP address
HY-MC-ASR1002X-1	0	10.6.32.251/32	10.6.32.151/25
HY-MC-CSR1000v-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.

```
copy running-config ftp://cisco:cisco@10.4.48.27
Address or name of remote host [10.4.48.27]?
Destination filename [hy-mc-asr1002x-1-config]?
Writing hy-mc-asr1002x-1-config !
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.

You can also make the changes below in a text editor before pasting the configuration into the router.

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 HY-MC-CSR1000v-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 GigabitEthernet1
  description IW-WAN-D3750X-T Gig1/0/3

interface GigabitEthernet2
  description IW-WAN-D3750X-T Gig2/0/3

interface range GigabitEthernet1, GigabitEthernet2
  no ip address
  cdp enable
  channel-group 21
  no shutdown
```

Procedure 4 Configure the routing protocol on the LAN

If you are planning to use EIGRP, choose option 1. If you are planning to use BGP on the WAN and OSPF on the LAN, choose option 2.

Option 1: EIGRP on the LAN

Step 1: Configure IP unicast routing using EIGRP named mode.

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
  network 10.8.0.0 0.1.255.255
  eigrp router-id 10.8.32.151
  exit-address-family
```

Step 2: 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
```

Option 2: OSPF on the LAN

Step 1: Configure OSPF Area 0 by using the network summary addresses and the loopback interface IP address as the router-id.

```
router ospf 100
  router-id 10.8.32.251
  network 10.8.32.128 0.0.0.63 area 0
  network 10.8.32.251 0.0.0.0 area 0
```

Step 2: Turn on passive-interface as the default and remove it for the LAN interface.

```
router ospf 100
  passive-interface default
  no passive-interface Port-channel21
```

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

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 13 Remote-site loopback IP address ranges

IWAN design model	Tunnel type	Loopback 0 address range
Hybrid-Primary Router	MPLS1	10.255.241.0/24
Hybrid-Secondary Router	INET1	10.255.242.0/24

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.

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      10.255.241.31/32 [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
```

show ip route | include 10.255.242

```
D      10.255.242.12/32 [90/25613440] via 10.8.32.129, 1w1d, Port-channel21
D      10.255.242.32/32 [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 cisco123
  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 14 Transit BR path and IP addresses

Host name	Path	Path ID	Loopback IP address	Zero SLA
HY-MPLS1-CSR1000v-T1	MPLS1	1	10.8.32.241/32	No
HY-INET1-CSR1000v-T2	INET1	2	10.8.32.242/32	Yes (optional)

Reader Tip

Whenever IWAN is designed with WAAS leveraging AppNav, please ensure that the Loopback IP address that is being used for PfR is not also used as the AppNav Service Controller address. This is applicable for any Hub IWAN router that is part of an AppNav Cluster.

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 cisco123
```

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 100 with MPLS as the provider.

```
interface Tunnel100
domain iwan path MPLS1 path-id 1
```

Step 3: (Optional) This example is the secondary hub BR using Tunnel 200 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 Tunnel200
 shutdown
 domain iwan path INET1 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.

PROCESS

Configuring Remote-Site Routers for Transit Site BRs

1. Configure NHRP at remote site
2. Configure POP selection and secondary paths at remote site

There are additional commands you need to configure at a remote site to begin using the transit site BRs.

Procedure 1 Configure NHRP at remote site

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 15 NHRP parameters

Hostname	Tunnel type	Tunnel number	Tunnel IP address	MPLS/public IP address
HY-MPLS1-CSR1000v-T1	MPLS1	100	10.6.34.2	192.168.6.41
HY-INET1-CSR1000v-T2	INET1	200	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. 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 remote site for hybrid-RS11-2921

```
interface Tunnel100
 ip nhrp nhs 10.6.34.2 nbma 192.168.6.41 multicast

interface Tunnel200
 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`.**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.

Procedure 2 Configure POP selection and secondary paths at remote site

If you are planning to use EIGRP, choose option 1. If you are planning to use BGP on the WAN and OSPF on the LAN, choose option 2.

Option 1: EIGRP on the WAN

The following tables show specific EIGRP route tags in use from the previous procedure.

Table 16 Route tag information for hub location

Tunnel interface	DMVPN tunnel key	Tag tunnel	Metric
Tunnel 100 (DMVPN 1)	1100 (MPLS1)	1100 (All routes)	+10000
Tunnel 200 (DMVPN 2)	1200 (INET1)	1200 (All routes)	+20000

Table 17 Route tag information for transit location

Tunnel interface	DMVPN tunnel key	Tag tunnel	Metric
Tunnel 100 (DMVPN 1)	1100 (MPLS1)	1600 (All routes)	+10000
Tunnel 200 (DMVPN 2)	1200 (INET1)	1700 (All routes)	+20000

Set the EIGRP metric value higher for the routes tagged from the non-preferred site.

Step 1: Define the route maps to identify the tags from border routers in POP1 and POP 2.

Example: Single-router remote site that prefers POP1

```

route-map POP-SELECT permit 10
  description Prefer POP1 for MPLS1
  match tag 1600
  set metric +10000

route-map POP-SELECT permit 50
  description Prefer POP1 for INET1
  match tag 1700
  set metric +20000

route-map POP-SELECT permit 100
  description Allow the rest

```

Example: Single-router remote site that prefers POP2

```

route-map POP-SELECT permit 10
  description Prefer POP2 for MPLS1
  match tag 1100
  set metric +10000

route-map POP-SELECT permit 50
  description Prefer POP2 for INET1
  match tag 1200
  set metric +20000

route-map POP-SELECT permit 100
  description Allow the rest

```

Step 2: Apply the POP select route map on the inbound tunnel interfaces.

```

router eigrp IWAN-EIGRP

  address-family ipv4 unicast autonomous-system 400

  topology base
    distribute-list route-map POP-SELECT in Tunnel100
    distribute-list route-map POP-SELECT in Tunnel200
  exit-af-topology

```

Step 3: Configure the maximum secondary paths.

The multi-data center feature adds support for secondary paths in the routing information base (RIB) of the supported routing protocols. The routing protocols are configured with one primary path and one or more secondary paths for a network. PfR is used for the primary path, as well the secondary paths, so they are all active-active.

Use the **maximum-secondary-paths** command to limit the number of additional entries in the RIB to the number of tunnels terminated on the remote site router. The path value is set to one minus the total number of WAN links multiplied by the number of DCs. For example, if there are 3 WAN links and 2 DCs configured on a remote site router, the number of secondary paths is set to 5.

The example below is for a remote site router with two WAN links into two DCs.

```

router eigrp IWAN-EIGRP

  address-family ipv4 unicast autonomous-system 400

  topology base
    maximum-secondary-paths 3
  exit-af-topology

```

Step 4: Repeat this process for each remote site that will use the transit BRs.

Option 2: BGP on the WAN

Table 18 Local preferences, community strings, and metrics for hub BRs at POP1

Transport	BGP local preference	BGP community string	OSPF metric preferred POP	OSPF metric secondary POP
MPLS1	800	65100:100	1000	2000
INET1	780	65100:200	1200	2200

Table 19 Local preferences, community strings, and metrics for transit BRs at POP2

Transport	BGP local preference	BGP community string	OSPF metric preferred POP	OSPF metric secondary POP
MPLS1	600 (MPLS1)	65100:101	1000	2000
INET1	580 (INET1)	65100:202	1200	2200

Step 1: Configure BGP to display communities in the format AA:NN.

```
ip bgp-community new-format
```

Step 2: Define the community lists to identify the border routers from POP1 and POP 2.

```
ip community-list standard POP1-MPLS1 permit 65100:100
ip community-list standard POP1-INET1 permit 65100:200

ip community-list standard POP2-MPLS1 permit 65100:101
ip community-list standard POP2-INET1 permit 65100:201
```

Step 3: Create the inbound route maps and update the outbound route map.

Update the outbound route map with a community string to signal the POP preference to the border routers.

Example:

```
65100:10 to prefer POP 1 (hub site)
65100:20 to prefer POP 2 (transit site)
```

Use a community string in the form of AS:NN with AS being the BGP autonomous system number and NN being the value that selects the preferred POP.

On the inbound route maps, set the local preference higher for preferred POP border routers.

Example: Single-router remote site that prefers POP1

```
route-map SPOKE-OUT permit 10
  description Prefer POP1 with community 65100:10
  set community 65100:10

route-map POP-SELECT permit 100
  description Prefer POP1 with higher LP
  match community POP1-MPLS1
  set local-preference 800

route-map POP-SELECT permit 120
  description Prefer POP1 with higher LP
  match community POP1-INET1
  set local-preference 780

route-map POP-SELECT permit 200
  match community POP2-MPLS1
  set local-preference 600

route-map POP-SELECT permit 220
  match community POP2-INET1
  set local-preference 580

route-map POP-SELECT permit 1000
  description If no match do not set LP
```

Example: Single-router remote site that prefers POP2

```
route-map SPOKE-OUT permit 10
  description Prefer POP2 with community 65100:20
  set community 65100:20

route-map POP-SELECT permit 100
  match community POP1-MPLS1
  set local-preference 600

route-map POP-SELECT permit 120
```



```

match community POP1-INET1
set local-preference 580

route-map POP-SELECT permit 200
description Prefer POP2 with higher LP
match community POP2-MPLS1
set local-preference 800

route-map POP-SELECT permit 220
description Prefer POP2 with higher LP
match community POP2-INET1
set local-preference 780

route-map POP-SELECT permit 1000
description If no match do not set LP

```

Step 4: Apply the POP select route map on the inbound WAN transports.

```

router bgp 65100
address-family ipv4
neighbor MPLS1-HUB route-map POP-SELECT in
neighbor INET1-HUB route-map POP-SELECT in

```

Step 5: Configure the maximum secondary paths.

The multi-data center feature adds support for secondary paths in the RIB of the supported routing protocols. The routing protocols are configured with one primary path and one or more secondary paths for a network. PfR is used for the primary path, as well the secondary paths, so they are all active-active.

Use the **maximum-secondary-paths** command to limit the number of additional entries in the RIB to the number of tunnels terminated on the remote site router. The path value is set to one minus the total number of WAN links multiplied by the number of DCs, and the **ibgp** keyword indicates the router is using Internal BGP peering between its neighbors. For example, if there are 3 WAN links and 2 DCs configured on a remote site router, the number of secondary paths is set to 5.

The example below is for a remote site router running iBGP with two WAN links into two DCs.

```

router bgp 65100
address-family ipv4
maximum-secondary-paths ibgp 3
exit-address-family

```

Step 6: Repeat this process for each remote site that will use the transit BRs.

Deploying a Second DC without a DCI Link

Optional

If you are not planning to use a DCI link between your hub and transit site locations, the following changes are needed on your border routers. If you are planning to use a DCI link, skip to the next process.

PROCESS

Configuring Hub and Transit Border Routers

1. Configure the routing protocol for the WAN
2. Configure NHRP at Hub and Transit Sites

This process applies only to DMVPN WAN Aggregation routers.

Procedure 1 Configure the routing protocol for the WAN

If you are planning to use EIGRP, choose option 1. If you are planning to use BGP on the WAN and OSPF on the LAN, choose option 2.

Option 1: EIGRP on the WAN

Step 1: To avoid issues with PfR probe routing, create IP prefix lists for blocking summary routes used at the remote sites.

Summarize the remote site LAN and remote site loopback addresses advertised into the hub and transit sites.

```
ip prefix-list HUB-SUMMARY seq 10 permit 10.7.0.0/16
ip prefix-list HUB-SUMMARY seq 20 permit 10.255.240.0/21
```

Step 2: Create a route map to block the summaries.

Use the IP prefix lists from the previous step to block the remote site LAN and remote site loopback summary addresses.

```
route-map BLOCK-HUB-SUMMARY deny 5
  description Block the remote site summary routes
  match ip address prefix-list HUB-SUMMARY

route-map BLOCK-HUB-SUMMARY permit 10
  description Permit the rest
```

Step 3: Add the inbound distribute list to the tunnel interface.

Use the route map from the previous step to block the summaries.

```
router eigrp IWAN-EIGRP
  address-family ipv4 unicast autonomous-system 400
  topology base
    distribute-list route-map BLOCK-HUB-SUMMARY in Tunnel100
  exit-af-topology
exit-address-family
```

Step 4: Repeat this process for each hub and transit site BR using EIGRP.

Option 2: BGP on the WAN

Step 1: Create a route map to control the updates between the hub and transit site BRs.

Use the existing IP prefix lists for local DC prefixes and the local MC loopback address.

```
route-map HUB-TO-HUB-OUT permit 10
  description Permit routes forwarded to Peer Hub
  match ip address prefix-list LOCALDC-PREFIX LOCALMCLOOPBACK
```

Step 2: Configure the BGP values for the neighbor relationship with the corresponding BR over the same transport.

For internal BPG, use the same AS number for the remote sites. Use the tunnel as the update source interface. Adjust the BGP hello and hold timers to 20 seconds and 60 seconds, respectively. Add the adjacent BR as a member of the peer-group.

Example: MPLS1 hub border router-HY-M1-ASR1002X-1

```
router bgp 65100
  neighbor NON-DCI-MPLS1 peer-group
  neighbor NON-DCI-MPLS1 remote-as 65100
  neighbor NON-DCI-MPLS1 description NON-DCI MPLS1 Hub BR Peers
  neighbor NON-DCI-MPLS1 update-source Tunnel100
  neighbor NON-DCI-MPLS1 timers 20 60
  neighbor 10.6.34.2 peer-group NON-DCI-MPLS1
```

Step 3: Configure the address family values for the neighbor relationship with the corresponding BR over the same transport.

Use the route map from the previous step to limit the updates between the adjacent BRs and activate the connection.

Example: MPLS1 hub border router–HY-M1-ASR1002X-1

```
address-family ipv4
  neighbor NON-DCI-MPLS1 send-community
  neighbor NON-DCI-MPLS1 next-hop-self all
  neighbor NON-DCI-MPLS1 weight 50000
  neighbor NON-DCI-MPLS1 soft-reconfiguration inbound
  neighbor NON-DCI-MPLS1 route-map HUB-TO-HUB-OUT out
  neighbor 10.6.34.2 activate
```

Step 4: Repeat this process for each hub and transit site BR using BGP.

Procedure 2 Configure NHRP at Hub and Transit Sites

Hub and transit site BRs without a DCI connection require an additional configuration statement in order to create an adjacency between them. This statement includes the NBMA definition for the adjacent router tunnel endpoint. The routing protocol relies on a multicast transport and requires that NHRP automatically add routers to the multi-cast NHRP mappings.

Step 1: Configure NHRP.

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 externally routable IP address because the hub and transit routers are in different locations. This design uses the values shown in the following table.

Table 20 NHRP parameters

Hostname	Tunnel type	Tunnel number	Tunnel IP address	MPLS / External IP address
HY-M1-ASR1002X-1	MPLS1	100	10.6.34.1	192.168.6.1
HY-I1-ASR1002X-2	INET1	200	10.6.36.1	172.16.140.1
HY-M1-CSR1000v-T1	MPLS1	100	10.6.34.2	192.168.6.41
HY-I1-CSR1000v-T2	INET1	200	10.6.36.2	172.16.140.2

The two corresponding BRs must point at each other in order to allow an adjacency to be formed. For the nbma address on the adjacent BRs, use the externally routable IP address for MPLS or INET.

Example: MPLS1 hub border router–HY-M1-ASR1002X-1

```
interface Tunnel100
  ip nhrp nhs 10.6.34.2 nbma 192.168.6.41 multicast
```

Example: MPLS1 transit border router–HY-M1-CSR1000v-T1

```
interface Tunnel100
  ip nhrp nhs 10.6.34.1 nbma 192.168.6.1 multicast
```

Tech Tip

If you are using BGP on the WAN, the `ip nhrp nhs` commands have to be added after the BGP neighbor commands to avoid dynamically creating the neighbor relationship ahead of time.

```
router bgp 65100
  neighbor 10.6.34.2 peer-group NON-DCI-MPLS1
% Cannot configure for dynamically created neighbor
```

If this happens, you will need to shut down the tunnel interface and clear the BGP routes for the new peer-group on the second BR. Because you have to shut down the tunnel interface, make sure you have IP connectivity to the BR through some other path prior to taking this action.

```
interface Tunnel100
  shut
end
clear bgp all peer-group NON-DCI-MPLS1
```

Rerun the `neighbor` command, and then perform a `no shut` on the tunnel interface to bring up the connection.

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

Applying DMVPN QoS Policy to DMVPN Hub Routers

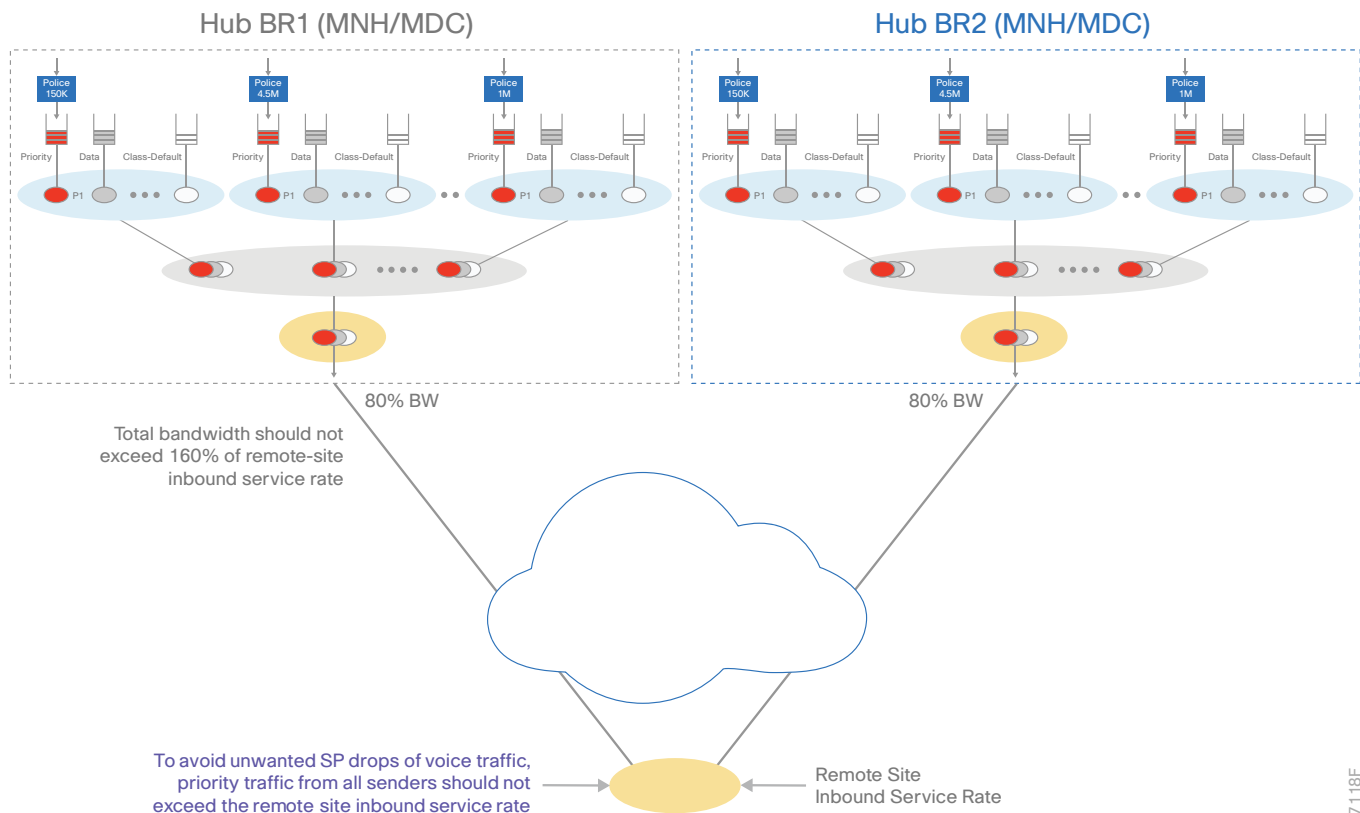
1. Configure per-tunnel QoS policies for DMVPN hub router
2. Configure per-tunnel QoS NHRP policies on DMVPN hub router

This process applies only to DMVPN WAN Aggregation routers.

For multiple DC deployments, the QoS configuration has to take into account each hub BR sending traffic to the same remote site router at the same time. In a multiple sender environment, the following rules are applied:

- Total bandwidth for all senders should not exceed 160% of remote-site inbound service rate, because over-subscription traffic will be dropped in the SP cloud.
- Bandwidth has to be divided equally between the hub BRs because there is only one NHRP group command on each remote site tunnel interface.
- QoS child policies do not have to be the same per sender, but DSCP markings must match for PfR Traffic Channel channels to establish.
- To avoid unwanted SP drops of voice traffic, priority traffic from all senders should not exceed the remote site inbound service rate.
- As the number of senders increase, the percentages used per site must come down accordingly based on the network administrator's knowledge of their traffic patterns.

Figure 3 Multiple sender QoS for hub routers



7-118F

Procedure 1 Configure per-tunnel QoS policies for DMVPN hub router

This procedure is based on the QoS configurations from the [IWAN Deployment Guide](#). Only the changes needed for this advanced guide will be discussed.

In the example below, configure the traffic at 80% of the remote site inbound service rate bandwidth.

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 get all of their assigned bandwidth, while the higher bandwidth sites 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 21 Per-tunnel QoS policies for 80% of the bandwidth

Policy name	Class	Bandwidth bps	Bandwidth remaining ratio
RS-GROUP-300MBPS-80-POLICY	class-default	240000000	240
RS-GROUP-200MBPS-80-POLICY	class-default	160000000	160
RS-GROUP-100MBPS-80-POLICY	class-default	80000000	80
RS-GROUP-50MBPS-80-POLICY	class-default	40000000	40
RS-GROUP-30MBPS-80-POLICY	class-default	24000000	24
RS-GROUP-20MBPS-80-POLICY	class-default	16000000	16
RS-GROUP-10MBPS-80-POLICY	class-default	8000000	8
RS-GROUP-4G-80-POLICY	class-default	6000000	6

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 previously for all of your sites.

The shape average value is entered in bits per second (bps). If all of your bandwidth values are greater than 5 Mbps, enter the bandwidth remaining ratio as shape average bandwidth/1 Mbps. If any of your bandwidth values are 5 Mbps or less, enter the bandwidth remaining ratio as shape average bandwidth/100 Kbps.

```
policy-map [policy-map-name]
  class class-default
    shape average [bandwidth (bps)]
    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 border router using 80% policies for all remote site types

```
policy-map RS-GROUP-300MBPS-80-POLICY
  class class-default
    shape average 240000000
    bandwidth remaining ratio 240
    service-policy WAN
policy-map RS-GROUP-200MBPS-80-POLICY
  class class-default
    shape average 160000000
    bandwidth remaining ratio 160
```



```
    service-policy WAN
policy-map RS-GROUP-80MBPS-POLICY
class class-default
  shape average 80000000
  bandwidth remaining ratio 80
    service-policy WAN
policy-map RS-GROUP-50MBPS-80-POLICY
class class-default
  shape average 40000000
  bandwidth remaining ratio 40
    service-policy WAN
policy-map RS-GROUP-30MBPS-80-POLICY
class class-default
  shape average 24000000
  bandwidth remaining ratio 24
    service-policy WAN
policy-map RS-GROUP-20MBPS-80-POLICY
class class-default
  shape average 16000000
  bandwidth remaining ratio 16
    service-policy WAN
policy-map RS-GROUP-10MBPS-80-POLICY
class class-default
  shape average 8000000
  bandwidth remaining ratio 8
    service-policy WAN
policy-map RS-GROUP-4G-80-POLICY
class class-default
  shape average 6000000
  bandwidth remaining ratio 6
    service-policy WAN
```

Procedure 2 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]
  nhrp map group [NHRP GROUP Policy Name] service-policy output [policy-map name]
```

Example: Hub border router

```
interface tunnel100
  nhrp map group RS-GROUP-300MBPS-80 service-policy output RS-GROUP-300MBPS-80-POLICY
  nhrp map group RS-GROUP-200MBPS-80 service-policy output RS-GROUP-200MBPS-80-POLICY
  nhrp map group RS-GROUP-100MBPS-80 service-policy output RS-GROUP-100MBPS-80-POLICY
  nhrp map group RS-GROUP-50MBPS-80 service-policy output RS-GROUP-50MBPS-80-POLICY
  nhrp map group RS-GROUP-30MBPS-80 service-policy output RS-GROUP-30MBPS-80-POLICY
  nhrp map group RS-GROUP-20MBPS-80 service-policy output RS-GROUP-20MBPS-80-POLICY
  nhrp map group RS-GROUP-10MBPS-80 service-policy output RS-GROUP-10MBPS-80-POLICY
  nhrp map group RS-GROUP-4G-80 service-policy output RS-GROUP-4G-80-POLICY
```

PROCESS

Applying QoS Configurations to Remote Site Routers

1. Configure per-tunnel QoS NHRP policy on remote-site routers
2. Verify QoS policy on physical interfaces of remote site router
3. Verify DMVPN per-tunnel QoS from hub routers

This process applies only to DMVPN remote site routers.

Procedure 1 Configure per-tunnel QoS NHRP policy on remote-site routers

This procedure configures the remote-site router to reference the updated 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 the previous process. Configure the **bandwidth receive** command on the interface to match the NHRP group policy chosen and the inbound service rate. The bandwidth value is entered in kilobits per second (Kbps)

```
interface Tunnel[value]
  bandwidth receive [value in Kbps]
  nhrp group [NHRP GROUP Policy Name]
```

Example: Remote site router with dual-link for hybrid

This example shows a remote-site using 20 Mbps and 50 Mbps inbound service rate policies. The bandwidth receive command should reflect the actual amount of your inbound service rate.

```
interface Tunnel100
  bandwidth receive 20000
  nhrp group RS-GROUP-20MBPS-80

interface Tunnel200
  bandwidth receive 50000
  nhrp group RS-GROUP-50MBPS-80
```

Procedure 2 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.

Step 2: Repeat the previous step for each interface configured with QoS.

Tech Tip

If you experience tail-drops in your class class-default, a potential work-around is to increase the size of the queue-limit.

On an interface with bandwidth of less than 15 Mbps, the default queue-limit is 64 packets. Increasing this value will add latency to the traffic in the default-class but will also reduce the number of tail-drops.

```
policy-map WAN
  class class-default
    queue-limit 512 packets
```

Procedure 3 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.

Step 2: Repeat the previous step for each DMVPN hub router.

Appendix A: Product List

To view the full list of IWAN-supported routers for this version of the CVD, see [Supported Cisco Platforms and Software Releases](#). All master controllers and border router devices at a common site must use the same version of software.

This guide was validated using the software in this appendix. When deploying, you should always use the Cisco IOS Software Checker tool to see if there are software vulnerabilities applicable for your environment. This tool is available at the following location:

<https://tools.cisco.com/security/center/selectIOSVersion.x>

Appendix B: Changes

This appendix summarizes the changes Cisco made to this guide since its last edition.

- Routing updates:
 - Added non-DCI configurations for hub and transit site BRs
 - Added maximum secondary paths to remote site routers
 - Updated the tunnel interface and tunnel ID numbering to match other guides
 - Updated the EIGRP tag numbering to match other guides
- QoS update:
 - Added support for multiple sender QoS
- Platform update:
 - Added support for CSR1000v as BRs at transit site



You can 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.

© 2017 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)