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This document is archived and should only be used as a historical reference and should not be used for new deployments for one of the following reasons: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 design guide is the recommended alternative. **(Coming Soon)**

- This document is outdated.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 Design Summary

September 2017

Table of Contents

WAN Strategy

This guide provides a high-level overview of the Cisco Intelligent WAN (IWAN) architecture, followed by a discussion of the different design models and the IWAN best practices. The intended audience is a technical decision maker who wants to compare Cisco's wide-area network (WAN) offerings and learn more about the best practices for each technology. This guide should be used as a roadmap for how to use the companion IWAN deployment guides.

Reader Tip

For more information about deploying the Intelligent WAN, see the Design Zone for Branch WAN.

The days of conducting business with information stored locally on your computer are disappearing rapidly. The trend is for users to access mission-critical information by connecting to the network and downloading the information or by using a network-enabled application. Users depend upon shared access to common secured storage, web-based applications, and cloud-based services. Users may start their day at home, in the office, or from a coffee shop, expecting to log on to applications that they need in order to conduct business, update their calendar, or check email—all important tasks that support your business. Connecting to the network to do your work has become as fundamental as turning on a light switch to see your desk; it's expected to work. Taken a step further, the network becomes a means to continue to function whether you are at your desk, roaming over wireless local-area network (WLAN) within the facility, or working at a remote site, and you still have the same access to your applications and information.

Now that networks are critical to the operation and innovation of organizations, workforce productivity enhancements are built on the expectation of nonstop access to communications and resources. As networks become more complex in order to meet the needs of any device, any connection type, and any location, networks incur an enhanced risk of downtime caused by poor design, complex configurations, increased maintenance, or hardware and software faults. At the same time, organizations seek ways to simplify operations, reduce costs, and improve their return on investment by exploiting their investments as quickly and efficiently as possible.

With increasing mobile traffic from employee devices, an organization must plan for expanded WAN bandwidth at remote sites and larger router platforms to accommodate the higher capacity links.

There are many ways an organization can benefit by deploying a CVD enterprise IWAN architecture:

- Flexibility with multiple design models in order to address a variety of IWAN technologies and resiliency options
- Increased reliability with multiple remote-site designs that provide for resiliency through the addition of WAN links and WAN routers, depending on business requirements
- Scalability provided by using a consistent method for remote-site LAN connectivity based on the CVD enterprise campus architecture
- Reduced cost of deploying a standardized design based on Cisco-tested and supported best practices
- Summarized and simplified design choices so that IT workers with a CCNA certification or equivalent experience can deploy and operate the network

Using a modular approach to building your network with tested, interoperable designs allows you to reduce risks and operational issues and to increase deployment speed.

Hybrid WAN Designs: IWAN vs Traditional WAN

Hybrid WAN designs are becoming increasing popular because they allow an organization to choose the best transport options for their particular situation. Your organization can spend money on multiprotocol label switching (MPLS) services when the business needs require it. You can use Internet services when more bandwidth is needed for larger data transport requirements. There are some key differences between IWAN and traditional WAN hybrid designs, which are highlighted in the figure below.

Figure 1 Hybrid WAN designs

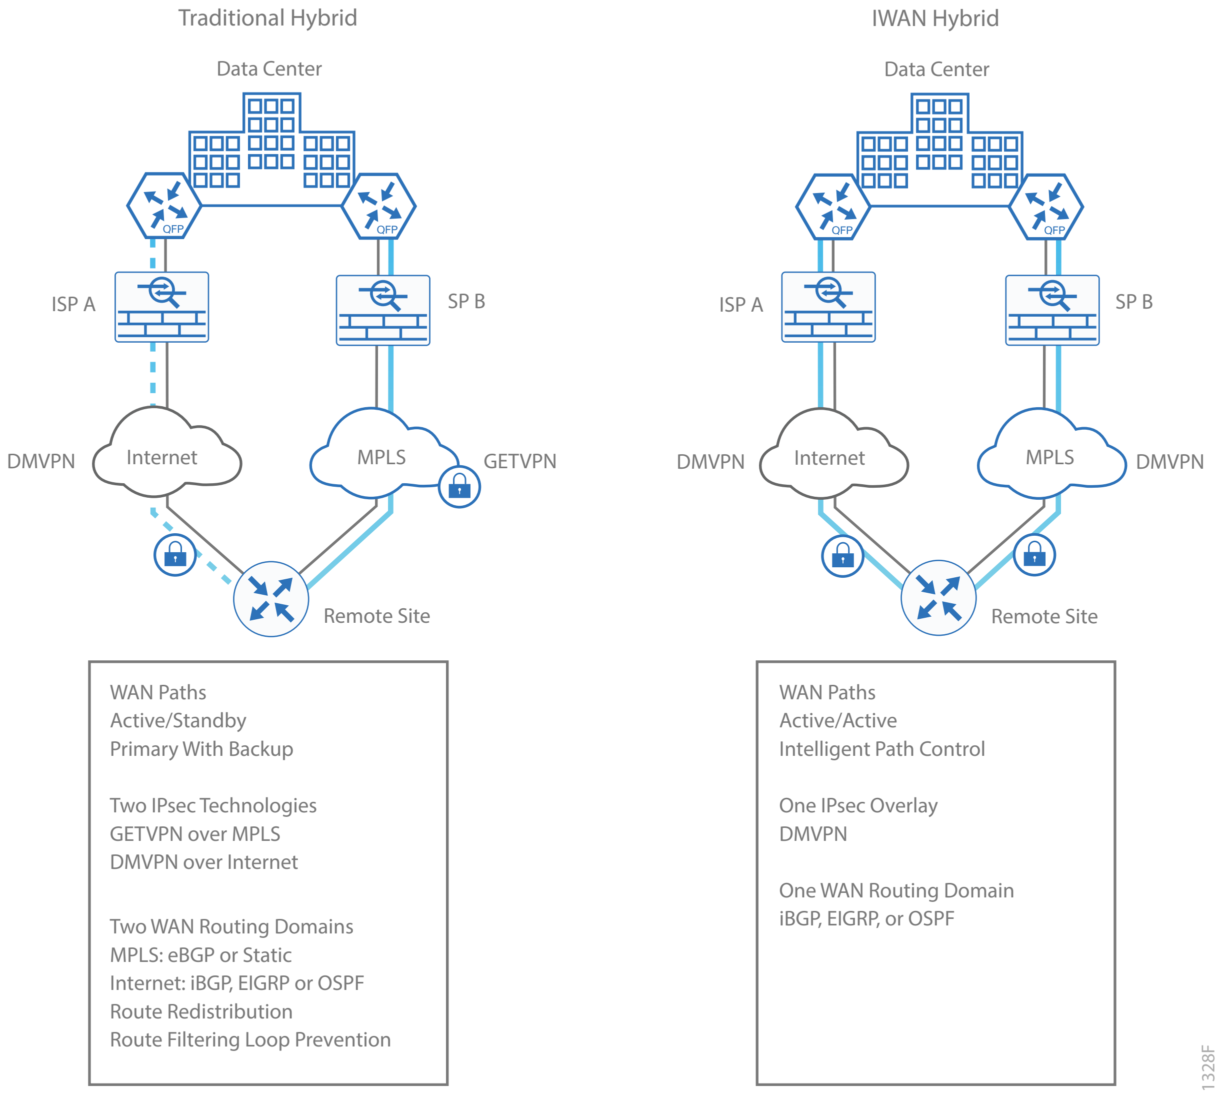

WAN Paths
Active/Standby
Primary With Backup

Two IPsec Technologies
GETVPN over MPLS
DMVPN over Internet

Two WAN Routing Domains
MPLS: eBGP or Static
Internet: iBGP, EIGRP or OSPF
Route Redistribution
Route Filtering Loop Prevention

WAN Paths
Active/Active
Intelligent Path Control

One IPsec Overlay
DMVPN

One WAN Routing Domain
iBGP, EIGRP, or OSPF

The IWAN design provides an active/active path for all WAN links and uses a single IPsec technology, which is not dependent on the underlying transport. The design also uses a single WAN routing domain without route redistribution or route filtering. The IWAN design is prescriptive in order to reduce the possible combinations, which lowers the cost and complexity for customers who want a simplified approach.

The traditional WAN hybrid design provides an active/standby path and two IPsec technologies based on the type of transport chosen. The design uses two WAN routing domains, which require route redistribution and route filtering for loop prevention. A traditional design has more transport options for customers who have varied needs, but because of the additional flexibility, the complexity is higher.

Reader Tip

For more information about traditional WAN design, see Traditional WAN Design Summary and the associated WAN deployment guides.

When planning your WAN strategy, Cisco recommends that you:

- Overprovision the WAN as much as possible
- Replace some or all of your MPLS bandwidth with Internet bandwidth
- Grow your existing WAN bandwidth with Internet bandwidth
- Keep quality of service (QoS) as simple as possible
- Use SDWAN management tools to automate and virtualize WAN connectivity

IWAN Introduction

The Cisco IWAN solution provides design and implementation guidance for organizations looking to deploy WAN transport with a transport-independent design (TID), 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.

BUSINESS USE CASES FOR IWAN

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 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 VPNs, Internet, Cellular (4G 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 design helps organizations connect remote sites over private (MPLS) and public (Internet) IP networks, efficiently and securely.

This design enables the following network capabilities:

- Secure, encrypted communications solutions for up to 2000 locations by 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 NAT is implemented
- QoS for WAN traffic such as voice, video, critical data applications, bulk data applications and management traffic

Use Case: Scale and High Availability

This design helps organizations scale their IWAN deployments beyond a single hub border router per DMVPN. It also provides high availability for hub site locations.

This design enables the following capabilities:

- Horizontal scaling across multiple border routers on a single DMVPN to utilize all WAN and router capacity
- Convergence across hub routers only when all channels in a hub location fail or reach their maximum bandwidth limits
- If the current channel to a remote site fails, convergence to an alternate channel on the same network
- Redundant hub master controller using Anycast IP

Use Case: Multiple Data Centers

This design helps organizations scale their IWAN deployments beyond a single point of presence (POP) location.

This design enables the following capabilities:

- Two or more data centers advertise different or the same set of prefixes.
- Data centers are reachable across the WAN core for each transit site.
- Remote sites can access any data center across either POP location, and data centers can reach any remote site across multiple transit sites.

Use Case: Multiple WAN Transports

This design helps organizations scale their IWAN deployments beyond a single pair of WAN transports at a POP location.

This design enables the following capabilities:

- Up to nine WAN transports at each POP with one designated as a path of last resort
- Convergence across WAN transports when all channels in a given transport fail or reach their maximum bandwidth limits
- Up to three WAN transports at a single-router remote site
- Up to five WAN transports at a dual-router remote site

Use Case: Multiple Virtual Routing and Forwarding Instances

This design helps organizations segment their traffic across their WAN using multiple virtual routing and forwarding instances (VRFs) in the overlay of their WAN transports.

This design enables the following capabilities:

- End-to-end traffic isolation
- Inter-VRF route leaking at the hub or transit site
- Up to twenty VRFs at a hub or transit site
- Up to seven VRFs at a remote site

BUSINESS USE CASES FOR THE IWAN APPLICATION

The IWAN application for the Application Policy Infrastructure Controller - Enterprise Module (APIC-EM) simplifies WAN deployments by providing a highly intuitive, policy-based interface that abstracts network complexity and design for business intent. The business policy is translated into network policies that are propagated across the network. The IWAN application is prescriptive of the Cisco Validated Design and provisioning of its core pillars for a large number of sites from a centralized location.

The IWAN application supports the following features and benefits:

- **Plug and Play—**The network is used to deploy Cisco routers in new sites. When the APIC-EM controller scanner discovers a new router, it creates a Network Information Database entry for it and then automatically configures it. This capability (zero-touch deployment) eliminates manual intervention, saving time and helping prevent errors. All you need to do is connect the cable and power up the device.

- **Centralized policy automation—**The IWAN application has a centralized policy automation engine that guarantees all sites run the business policies intended by the administrator. The IWAN application is also designed to allow the administrator to specify the business needs in terms of application delivery in a drag-and-drop intuitive fashion.

- **Public-key-infrastructure (PKI) certificate—**The IWAN application uses the APIC-EM Trust Manager service. This service automates the lifecycle management of issuing, renewing, and revoking the PKI X.509 certificate for IWAN. With this feature, the IWAN application greatly simplifies the process of establishing and keeping trust in the network.

- **Centralized hybrid WAN management—**The IWAN application defines business-level preferences by application or groups of applications in terms of the preferred path for hybrid WAN links. This feature allows for cost savings by helping guarantee delivery of application experience over any connection and using otherwise inactive or backup links.

- **QoS deployment and change of management—**The IWAN application can enforce QoS priority policies facing the WAN as well as the LAN. You can categorize applications into business-critical, default, or best-effort. This feature helps application traffic behave consistently and in accordance with your QoS service-level agreements (SLAs).

- **Network wide visibility and segmentation with Application Visibility and Control (AVC) —**The IWAN application provides a prepopulated set of common applications and lets you create a profile for custom applications. This feature helps you apply QoS and path control by application or set of applications and by business needs and priority.

- **DMVPN deployment and change of management—**The IWAN application fully automates the provisioning of the Cisco DMVPN. This automation includes the management of DMVPN internal IP address allocations as well as setting the high encryption policies. With the application, the provisioning of DMPVN is highly automated and simplified.

The following use cases are supported by the IWAN application.

Use Case: Hub Deployment (I-Block methodology)

The IWAN application enables the following capabilities:

- Seamless insertion of IWAN hub routers into customer network (I-Block)
- Single or dual DC set up with up to four WAN transports
- Hub validation steps
- Bring up the hub site with network wide settings, IP address management for overlay, SP profiles and the addition/deletion of POP locations
- Loopback interface as DMVPN source for MPLS hub
- NAT support for INET hub

Use Case: Zero Touch Branch Deployment

The IWAN application enables the following capabilities:

- Bootstrap process for branch routers
- Automatic device discovery with Plug and Play
- Plug and Play profiles for different WAN types
- Geo location settings
- WAN and LAN deployment
    - Combination of branch models supported (dual and single router designs with up to three WAN transports)
    - Automation of WAN and LAN set up for the branch

Use Case: Brownfield Branch Migration

The IWAN application enables the following capabilities:

- Conversion of existing non-IWAN Branches into IWAN
- Validation checks for branch sites
- Scale the process using Discovery App
- Layer 2 and layer 3 LAN in branch with VLANs
- LAN routes are automatically learned
- LAN prefixes are automatically redistributed into WAN and PfR database.

Use Case: Advanced Hierarchy

The IWAN application enables the following capabilities:

- Site affinity in Multi DC set up
- Path of Last Resort

Use Case: Application Policy

The IWAN application enables the following capabilities:

- Move applications between application groups
- Create custom applications
- Change business relevance on applications

Use Case: SLA optimization

The IWAN application enables the following capabilities:

- Configure application performance and path preference
- Change thresholds, preferred paths and categories for custom applications

The IWAN architecture is based on open interfaces, a software-defined networking services plane, and device-layer abstraction. The IWAN application allows full policy-directed deployment and operation of the network. The network, including all routers, is abstracted and the design engineer can focus on the business priorities without being concerned about the underlying topology.

For more information about the IWAN Application, see the Cisco IWAN Application on APIC-EM User Guide.

IWAN Architecture

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 were 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, 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 2 *Cisco IWAN solution components*

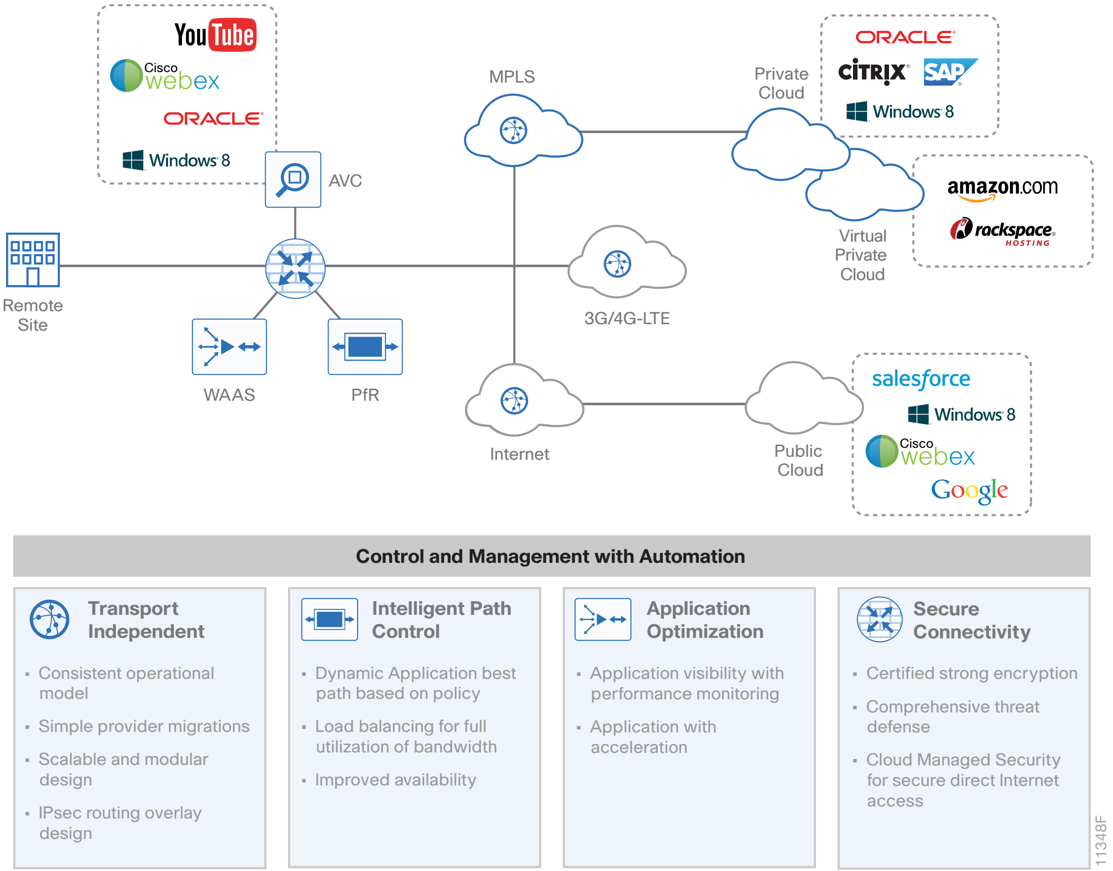

Control and Management with Automation

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 with acceleration

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 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.

TRANSPORT-INDEPENDENT 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 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 delivers the foundation for PfR to route around throughput fluctuations in the service providers’ network.

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.

Dynamic Multipoint VPN

DMVPN is the recommended solution for building scalable site-to-site VPNs that support a variety of applications. DMVPN is required for IWAN deployments because it provides a tight integration with PfRv3 and simplifies route control across any transport.

DMVPN was selected for the secure overlay IWAN solution because it supports on-demand full mesh connectivity over any carrier transport with a simple hub-and-spoke configuration. 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.

INTELLIGENT PATH CONTROL

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. At the hub site, the MC and BR are configured on separate routers for scalability. At a remote site, the MC and BR can be combined in the same router as shown in the figures below.

Figure 3 Cisco Performance Routing: Hub location

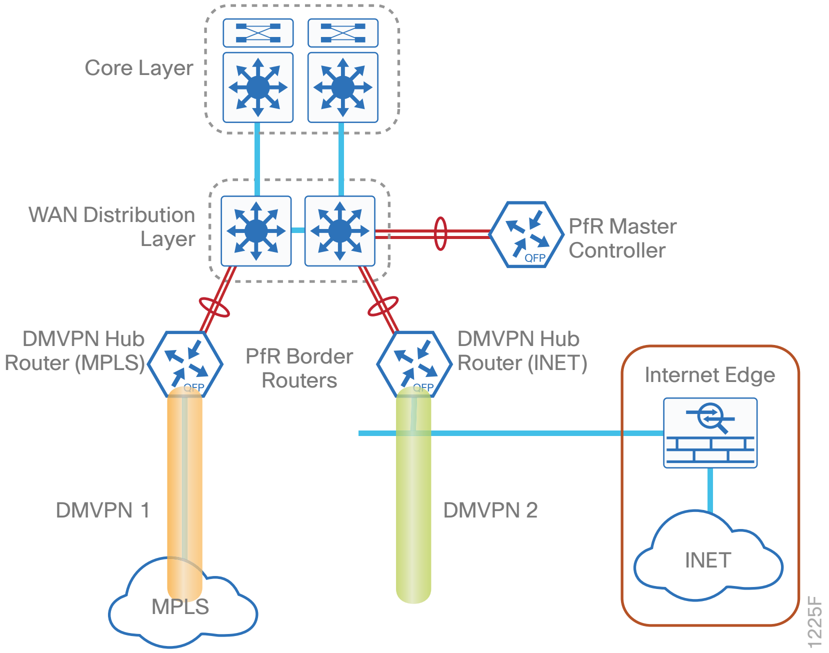

Figure 4 Cisco Performance Routing: Remote site options

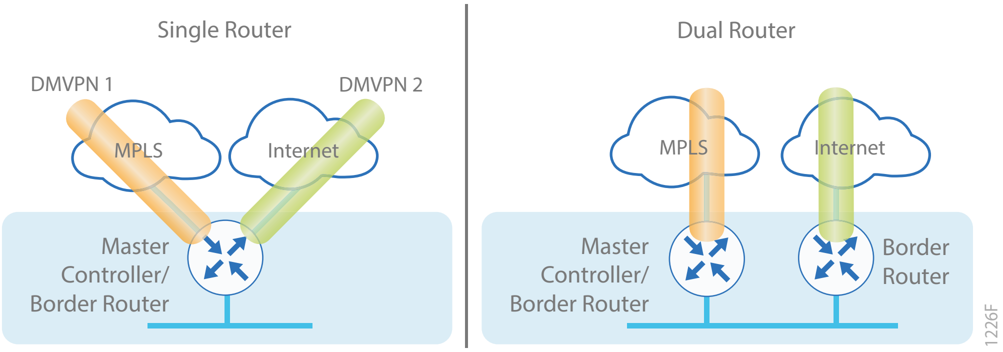

IWAN DESIGN

This guide describes two base IWAN design models and four advanced IWAN design models, but there are many options a customer can configure using the same underlying principles. The main goal of any IWAN design is to pick a preferred WAN path for your critical traffic. After establishing the preferred path, add more bandwidth for non-critical traffic using additional transports. The number of transports at any given remote site has to be at least two, but it can be as many as five with a dual-router design.

The first design model discussed is the IWAN Hybrid, which uses MPLS paired with Internet as the WAN transports. This is the primary use case for IWAN and the most common design being deployed. In this design model, the MPLS provides bandwidth for the critical classes of services needed for key applications and provides SLA guarantees for these applications. The non-critical classes use Internet bandwidth or whatever additional transports are available at each remote site location.

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. In this design model, you still have to determine a preferred path for your critical classes of traffic. In most cases, the preferred path will be the provider with the most bandwidth, but you might also want to choose the one that has the most favorable peering agreements or the one where the majority of your remote sites have direct connections. The non-critical classes can use the secondary Internet bandwidth available at each location.

The first advanced design model is the IWAN Dual Hybrid with Path of Last Resort (PLR), which has two MPLS transports, two Internet transports, and a fifth transport used as the final option when the other four are not available. In this design model, you still have to determine a preferred path for your critical classes of traffic. In some remote locations, the preferred path can be different from others locations. This multiple transport design model is not limited to two MPLS, two Internet, and one PLR transport, but the concepts described within this series of guides will give you an understanding of how to deploy IWAN with more than two transports. A customer can choose to deploy any combination of three to nine transports at their data center locations knowing they can support up to three transports at a single-router remote site and up to five at a dual-router remote site.

Figure 5   *Cisco IWAN design models*

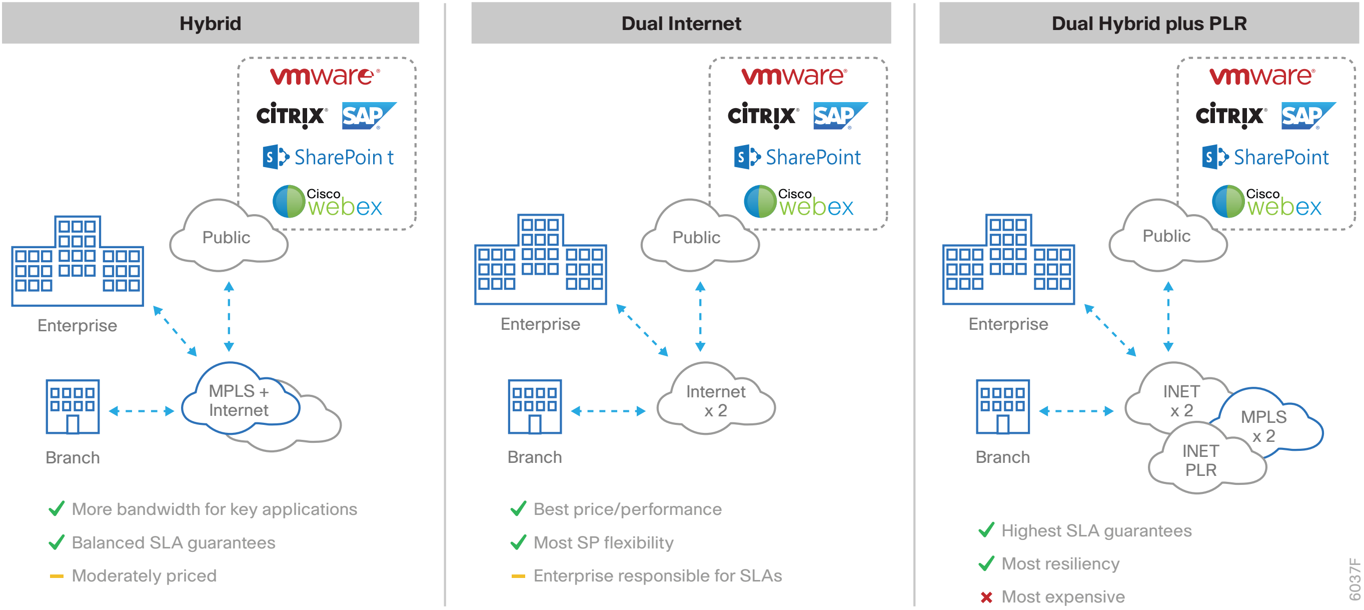

The IWAN WAN-aggregation includes at least two WAN edge routers, and each design model can scale up to 2000 remote sites.

Regardless of the design model, the WAN aggregation routers always connect into a pair of distribution layer switches using port-channel interfaces for additional bandwidth capacity and redundancy. Each of the design models has 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

This design allows you to move non-critical traffic off your MPLS, which gives your critical traffic better performance over the same amount of bandwidth. You use less expensive Internet bandwidth for your non-critical traffic classes. It provides balanced SLA guarantees and is moderately priced.

Figure 6 *WAN aggregation: IWAN hybrid design model*

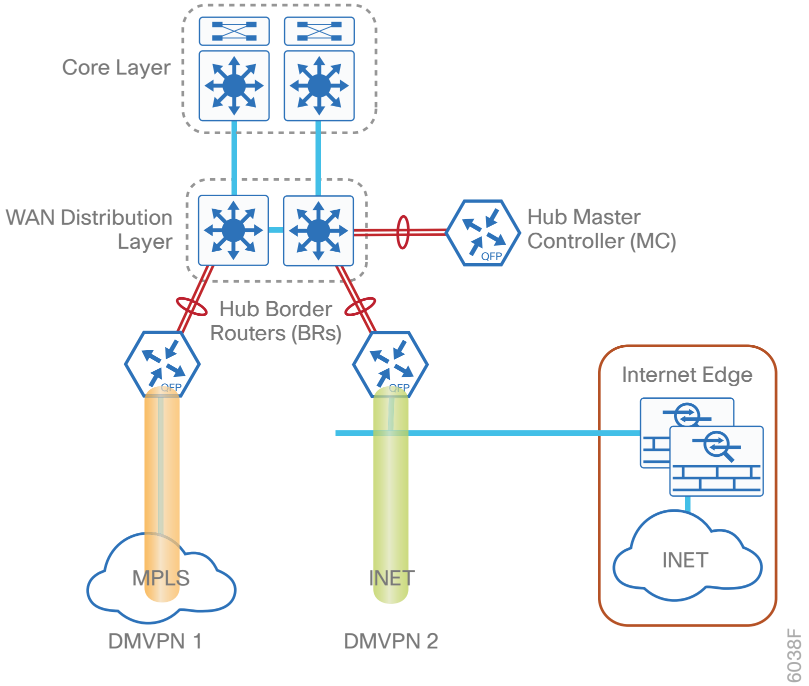

This model has the following capabilities:

- Uses at least one MPLS carrier
- Uses at least one Internet carrier
- Uses front-door virtual routing and forwarding (FVRF) on both MPLS and Internet links, with static default routing within the FVRF
- Scales to 2000 remote sites

IWAN Dual Internet Design Model

The dual Internet design has the least expensive monthly recurring cost, and it provides the most flexibility when choosing service providers. It is up to the enterprise to provide the SLA, because there are no bandwidth guarantees when using the Internet.

Figure 7 WAN aggregation: IWAN dual Internet design model

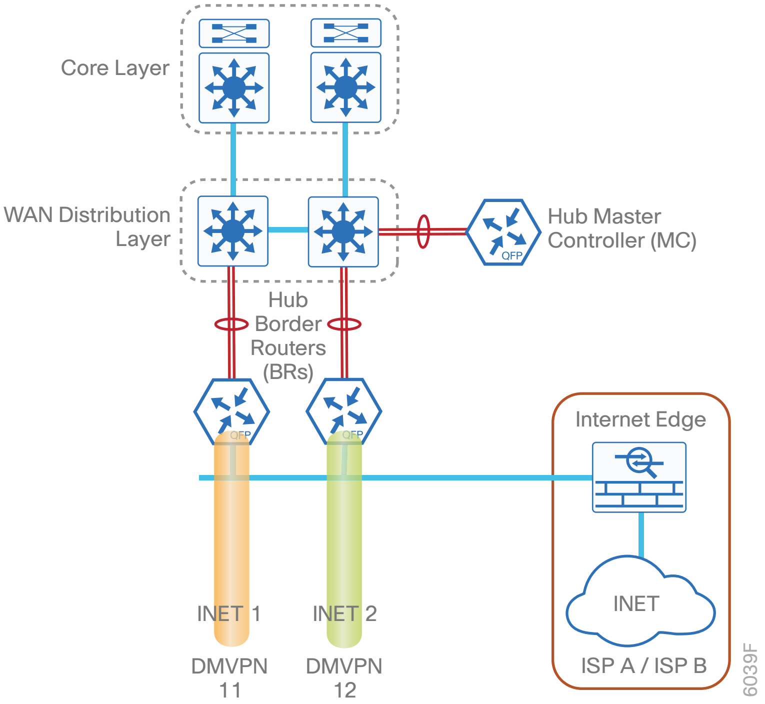

This model has the following capabilities:

- Uses at least two Internet carriers
- Uses FVRF on both Internet links, with static default routing within the FVRF
- Scales to 2000 remote sites

IWAN Scale and High Availability Design Model

This advanced design builds on previous design models by adding hub borders routers for horizontal scaling at a single data center. This design also has an option to add a second hub MC at a single data center for high availability.

The scale and high availability (HA) design models are as follows:

- A single data center with multiple borders for horizontal scalability
- Redundant Hub MC using Anycast IP

Single Data Center with Multiple Borders

In the following figure, two DMVPN hub border routers are used in a single data center for each service provider. There are two paths and two next-hops to the hub site from each remote site. To differentiate traffic from different ISP paths, a path-id is added on each DMVPN path. A path-id is a unique 32-bit number for a path between two sites.

Figure 8 Single data center with multiple borders and redundant hub MC

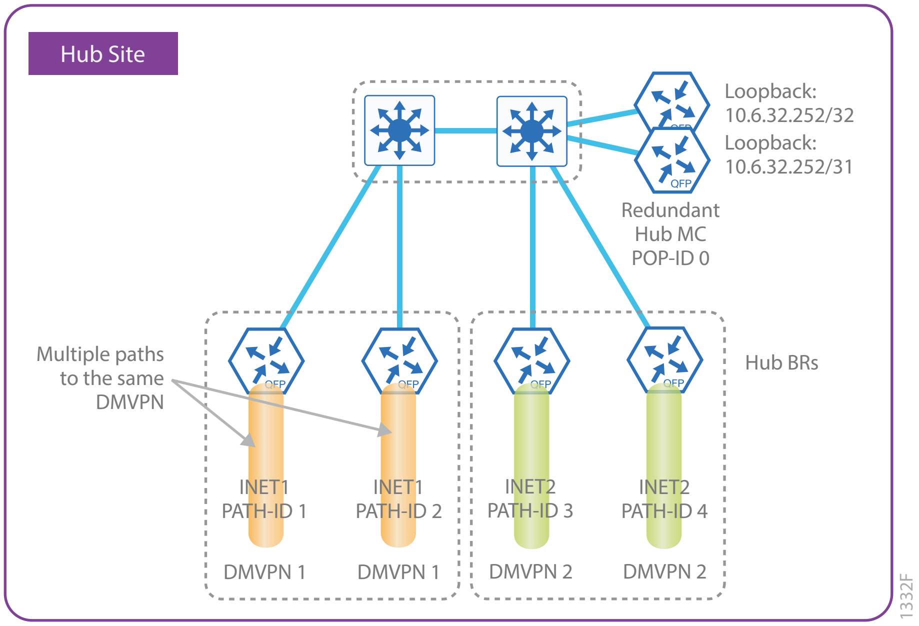

This model has the following capabilities:

- Distribute traffic across multiple border routers on a single (DMVPN) network in order to utilize all WAN and router capacity. Convergence across networks should only occur when all channels fail or when they reach their maximum bandwidth limits.
- If the current channel to a remote site fails, converge to an alternate channel on the same (DMVPN1) network. If both channels fail, converge over to the alternate (DMVPN2) network.
- Add a redundant Hub MC using Anycast IP and a secondary loopback interface with the same IP address and a /31 network mask. This HA concept works with all of the IWAN design models and it can be used with any standalone master controller, such as a transit master controller at a second data center or a standalone branch MC at a large remote site. Configure a second MC with the same base configuration as the first one and make a few minor changes to allow it to take over when the first MC goes offline. The two MCs must be kept in sync manually, but the failover occurs automatically within a few minutes depending, on the size of your IWAN implementation.

Tech Tip

The Hub MC HA feature is used to protect against the failure of the MC device at a single location. The redundant hub MC cannot be at a different location.

IWAN Multiple Data Center Design Model

This advanced design builds on previous design models with data center redundancy. The multi-data center or the transit site support feature enables organizations to scale their network infrastructure and load-balance the traffic when required. The multi-data center support enables multiple POP sites to be connected with the remote sites in an enterprise network. For example, in a use case scenario, an organization with two data centers and a single remote site, the remote site can communicate with the data centers through the next-hops (border routers) located at the hub and transit sites. If one border router is down, then the remote site can still communicate through the other border routers. To differentiate the traffic from different POP sites, a pop-id is configured by location and a path-id is configured on the interface of the hub and transit site border routers. The remote site router determines the inbound traffic based on the pop-id and path-id of the border routers.

The transit site design models are as follows:

- Multiple data centers with different prefixes
- Multiple data centers with shared prefixes

This guide documents multiple IWAN advanced deployment designs, and they are based on various combinations of data centers and advertised prefixes for specific requirements of service levels and redundancy.

Multiple Data Centers with Different Prefixes

In the following illustration, the two data centers are connected to all of the remote sites. You can use the data centers in active/active mode and use separate prefixes for each data center. To differentiate the traffic originating from the data centers, a pop-id is assigned to the MC at a data center. The valid range for a pop-id is from 1 to 62. By default, a pop-id of 0 is assigned to the hub MC.

Figure 9 *Multiple data centers with different prefixes*

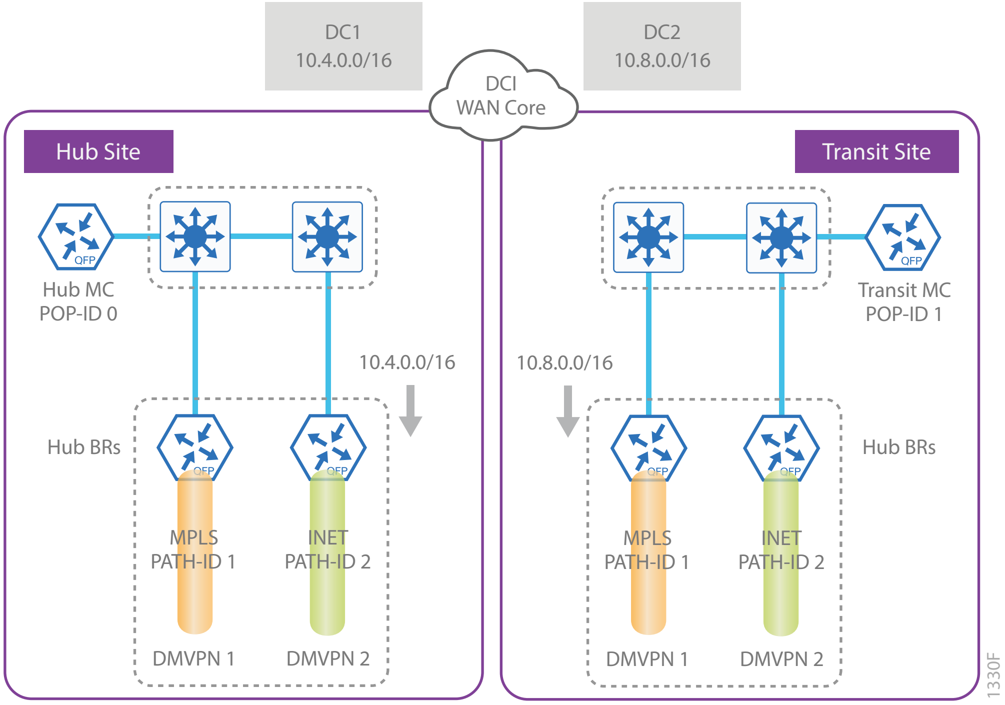

This model has the following characteristics:

- Two or more transit sites which advertise different prefixes.
- The data centers may be collocated with the transit sites.
- The data centers are reachable across the WAN core for each transit site.
- Multiple border routers per DMVPN network may be required for crypto and bandwidth horizontal scaling.

Multiple Data Centers with Shared Prefixes

In the following illustration, two data centers are connected to the remote sites. However, in this scenario both the data centers are active and load-balance the traffic. If one data center is down, then traffic is routed through the other data center. The IWAN hub site and transit site advertise the same set of prefixes.

Figure 10 *Multiple data centers with shared prefixes*

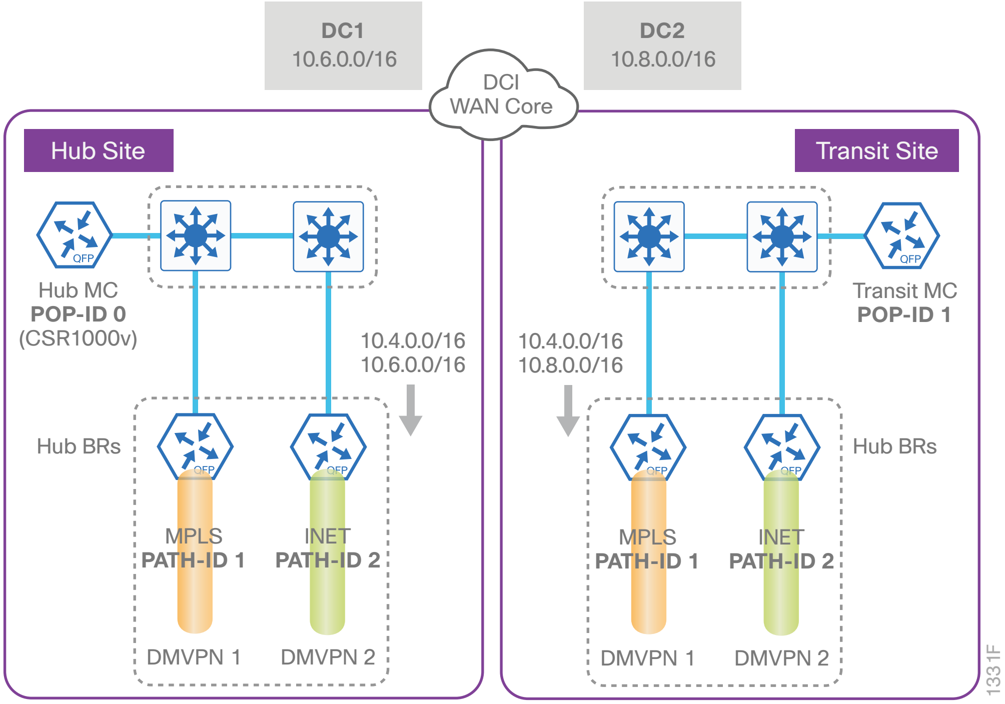

This model has the following characteristics:

- Two or more transit sites advertise a shared set of prefixes.
- The data centers may not be collocated with the transit sites.
- The data centers are reachable across the WAN core for each transit site.
- Remote site can access any data center across either hub border router. Data centers can reach any remote site across multiple transit site border routers.
- Multiple border routers per DMVPN network may be required for crypto and bandwidth horizontal scaling.

IWAN Dual Hybrid with PLR Design Model

This advanced design adds multiple WAN transports to any of the previous design models. The multiple transport design model is not limited to two MPLS, two Internet and one PLR transport, but this specific design will be used to discuss the underlying principles. The same concepts can be applied to other multiple transport designs.

Figure 11 *WAN aggregation: IWAN dual hybrid with PLR design model*

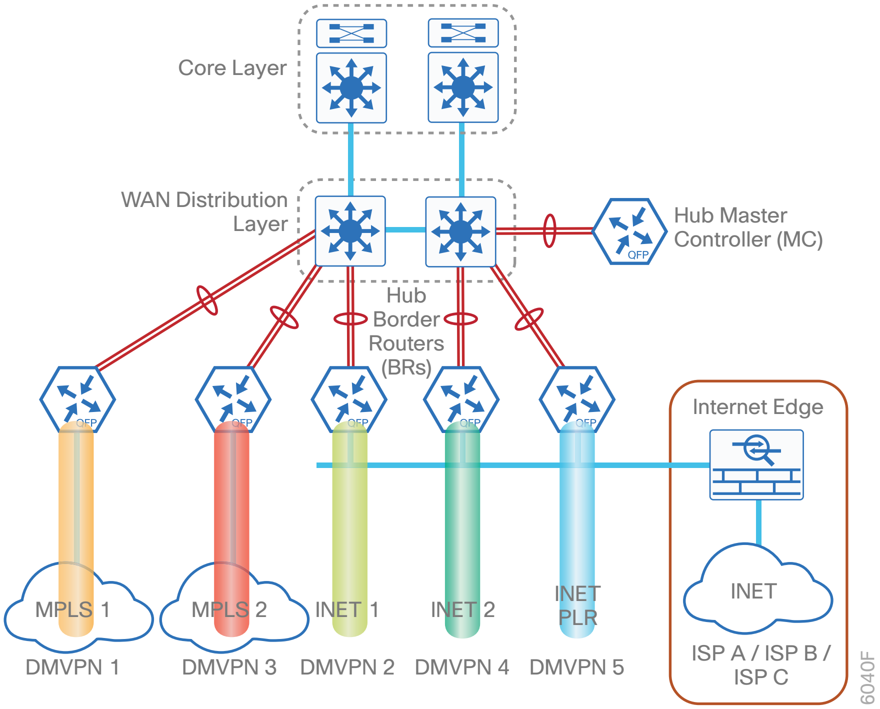

This model has the following characteristics:

- Uses two MPLS carriers
- Uses two Internet carriers
- Uses an additional Internet carrier as a PLR
- Scales to 2000 remote sites

The PLR feature provides the ability to designate a transport such that when the primary and fallback transports become unavailable or are out of bandwidth, traffic is routed over the path of last resort. This feature is used for metered links where data is charged on a usage basis and the path is only used when no other transports are available.

IWAN Multiple Tunnel Termination Design Model

This advanced design model allows you to add network access resiliency at your hub or transit site without adding network devices. The DMVPN Multiple Tunnel Termination (MTT) feature provides support for multiple tunnel terminations (interfaces) in the same VRF on the same hub device. The feature also provides transport resilience to DMVPN. Using one tunnel per-transport provides better visibility to PfR about the conditions in the underlying transport.

Figure 12 Multiple tunnel termination with IWAN dual hybrid design model

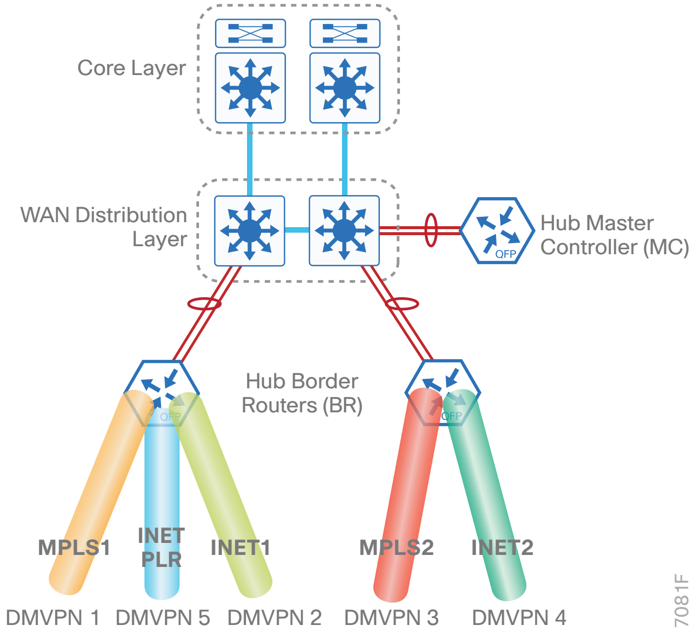

This model has the following characteristics:

- Uses two or more transports per hub or transit site border router
- Scales to 2000 remote sites

The DMVPN Multiple Tunnel Termination feature also brings in support for secondary paths for the supported routing protocols in the RIB. The routing protocols are configured in such a way that there is only one primary/regular path and one or more secondary paths for a network. When PfR is used in conjunction with this feature, the primary and secondary paths are active-active.

IWAN Multiple VRF Design Model

This advanced design model adds virtual routing and forwarding instances in the IWAN overlay for isolating end-to-end traffic among multiple independent networks inside the same data center. Similarly, at the remote location, devices in each VRF co-exist inside the same physical hardware as the common employee network.

Figure 13 *Multiple VRF with IWAN dual hybrid design model*

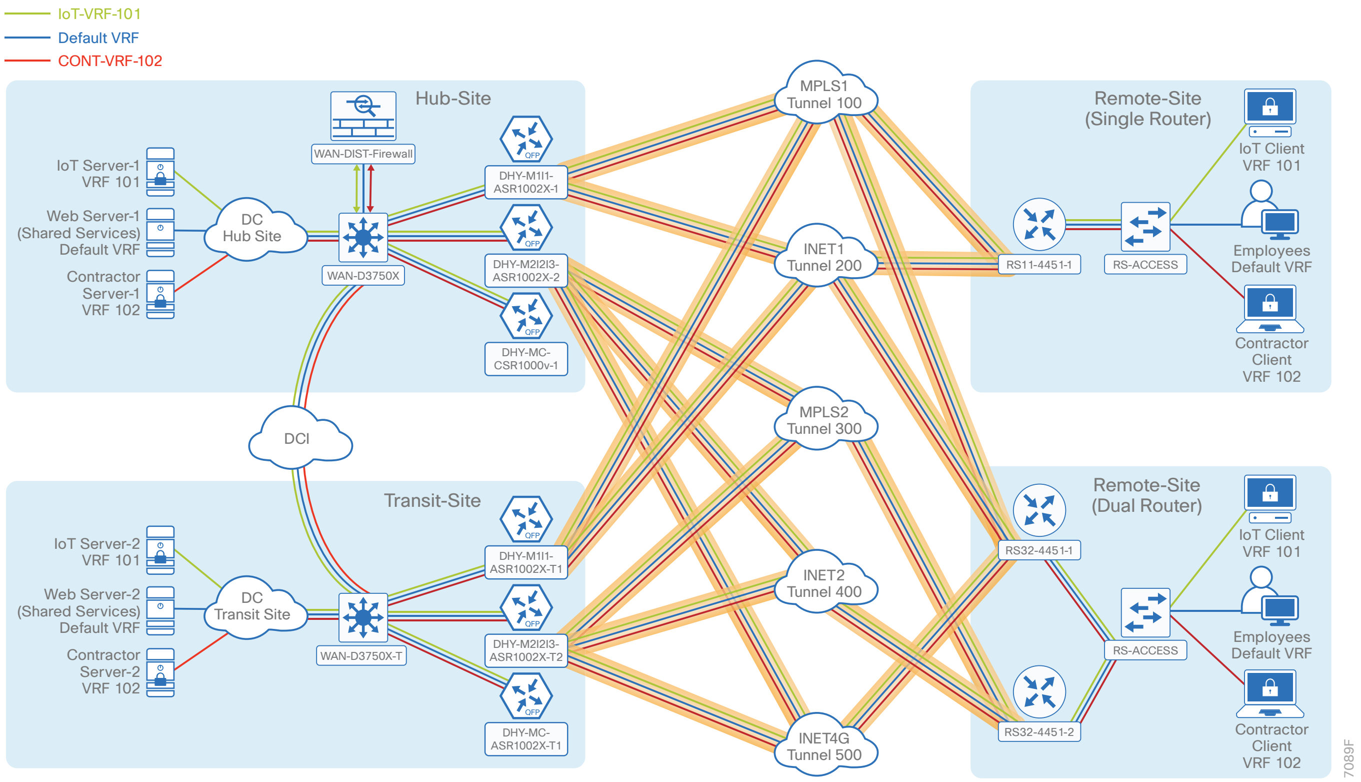

The model is not limited to three VRFs, but this specific design is used to show the underlying principles for multiple VRFs. The multiple VRF design can be used with any of the previous design models.

This model has the following characteristics:

- User traffic is segmented across the WAN
- Hub and transit sites support up to twenty VRFs
- Remote sites support up to seven VRFs

Inter-VRF route leaking is used to share global services such as DNS, DHCP, the company web portal, etc. between the different isolated segments. This design uses a firewall, attached to the WAN distribution switch at the hub-site location, to import and export individual VRF route information to and from the global routing tables. This gives the network administrator a secure device to manage the separation between the logical networks.

In all 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 Internet 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 hub router connects to a MPLS carrier.

IWAN REMOTE-SITE DESIGN

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 14 IWAN remote-site design options

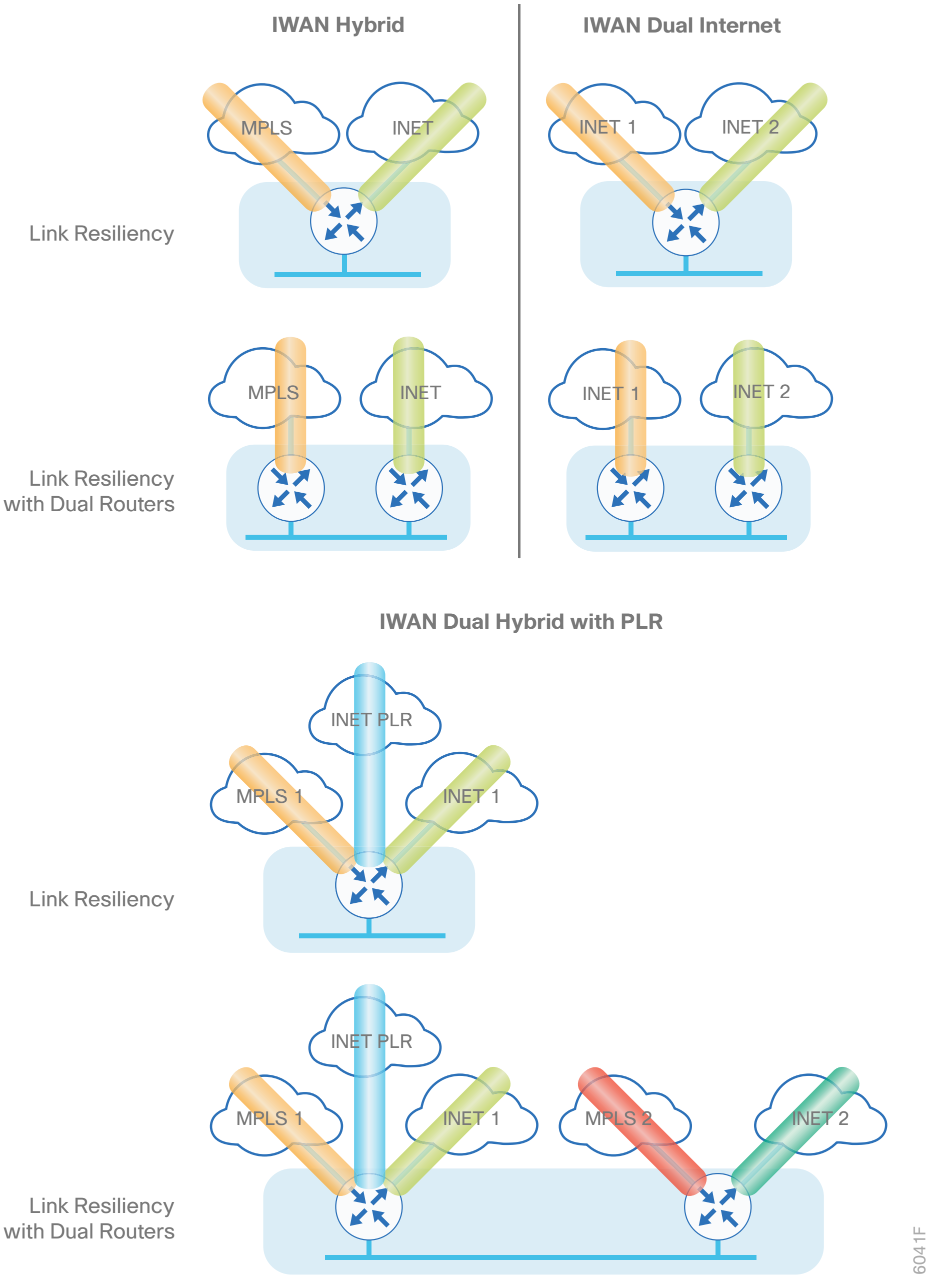

The remote-site designs include single or dual WAN edge routers. The remote-site routers are DMVPN spokes to the primary site hubs.

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.

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

Design model	WAN remote-site routers	WAN transports	Primary transport	Secondary transport(s)
Hybrid	Single	Dual	MPLS	Internet
Hybrid	Dual	Dual	MPLS	Internet
Dual Internet	Single	Dual	Preferred Internet	Internet
Dual Internet	Dual	Dual	Preferred Internet	Internet
Dual Hybrid	Single	Up to three	MPLS	Any available
Dual Hybrid	Dual	Up to five	MPLS	Any available

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

Design model	WAN remote- site routers	WAN transports	Primary transport	Secondary transport	PLR transport
Hybrid	Single	Dual w/ fallback	MPLS	Internet	4G LTE
Dual Internet	Single	Dual w/ fallback	Preferred 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.

Remote-site LAN

The primary role of the WAN is to interconnect primary site and remote-site LANs. The LAN discussion within this design 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 Wired LAN Technology Design 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 are shown in the following table.

Table 3 Remote-site LAN topology

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

For consistency and modularity, all WAN remote sites use the same VLAN assignment scheme, which is shown in the following table. This design 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 Remote-site 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)

Remote-site 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 Wired LAN Technology Design 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 15 *Remote-site with flat layer 2 LAN (single router)*

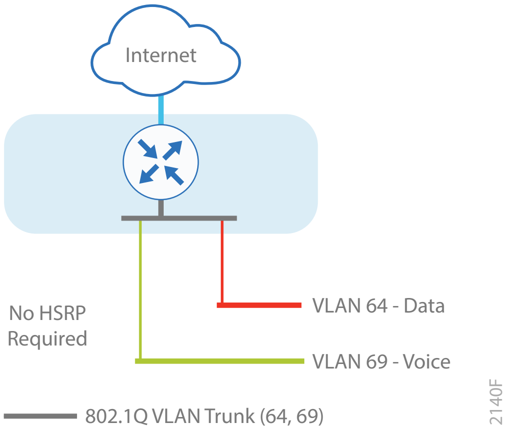

A similar LAN design can be extended to a dual-router edge as shown in the following figure. This design change introduces some additional complexity. The first requirement is to run a routing protocol. You need to configure enhanced interior gateway routing protocol (EIGRP) or open shortest path first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

Because there are now two routers per subnet, you must implement a first-hop redundancy protocol (FHRP). 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 16 *Remote-site with flat Layer 2 LAN (dual router)*

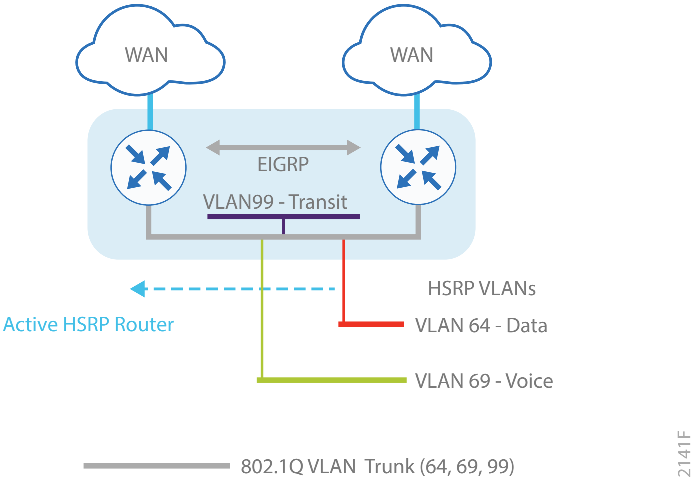

Enhanced object tracking (EOT) provides a consistent methodology for various router and switching features to conditionally modify their operation based on information objects available within other processes. The objects that can be tracked include interface line protocol, IP route reachability, and 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 preferred 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. If multiple WAN transports are available on the primary router, like in the case of a Dual Hybrid design, EOT tracks all of them using the **track list boolean or** command. The track list feature allows HSRP to remain on the primary router as long as at least one WAN tunnel interface is still operational.

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.

Remote-site 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 17 *IWAN single router remote-site: Connection to distribution layer*

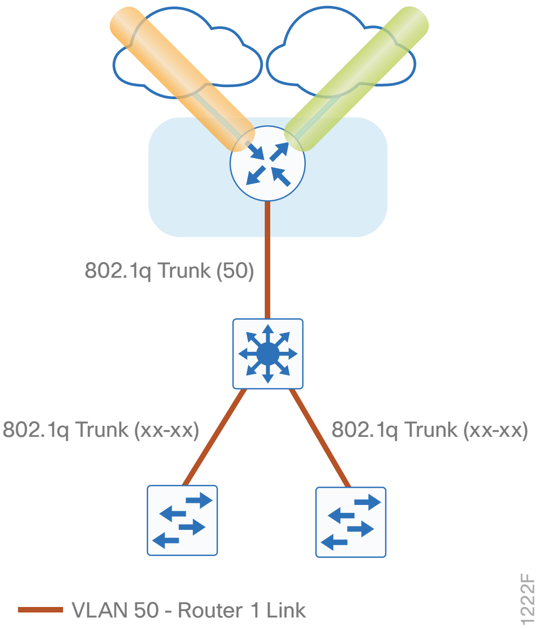

Figure 18 IWAN dual router remote-site: Connection to distribution layer

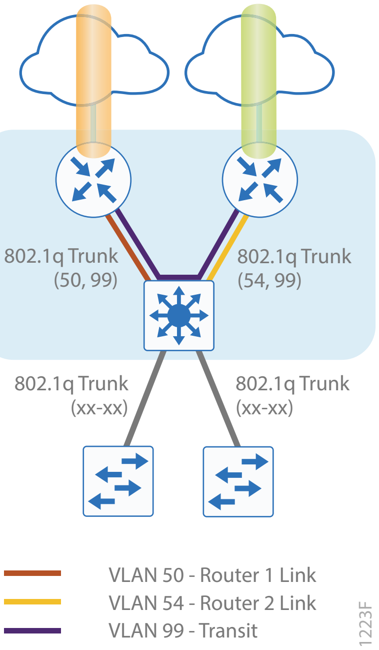

The distribution switch handles all access layer routing, with VLANs trunked to access switches. No HSRP is required when the design includes a distribution layer. A full distribution and access layer design is shown in the following figure.

Figure 19 IWAN dual router remote-site: Distribution and access layer

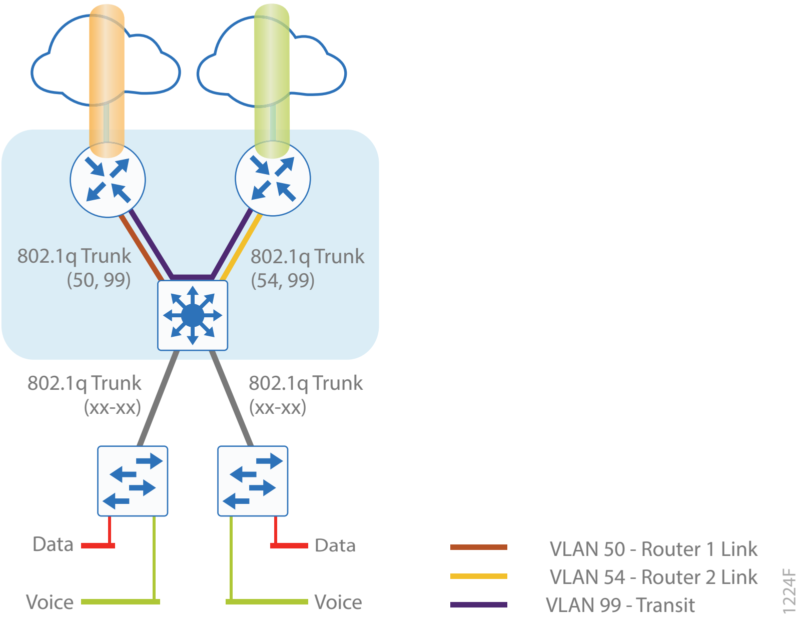

IP MULTICAST

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 (IGMP) 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.

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. 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.

There are twelve common traffic classes that are grouped together based on 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-class mappings in the following table are applied throughout this design by using an eight-class queuing model in the enterprise and a six-class, five-class, or four-class model in the service provider network as indicated in the figure below.

Figure 20 *QoS service 12-class mappings*

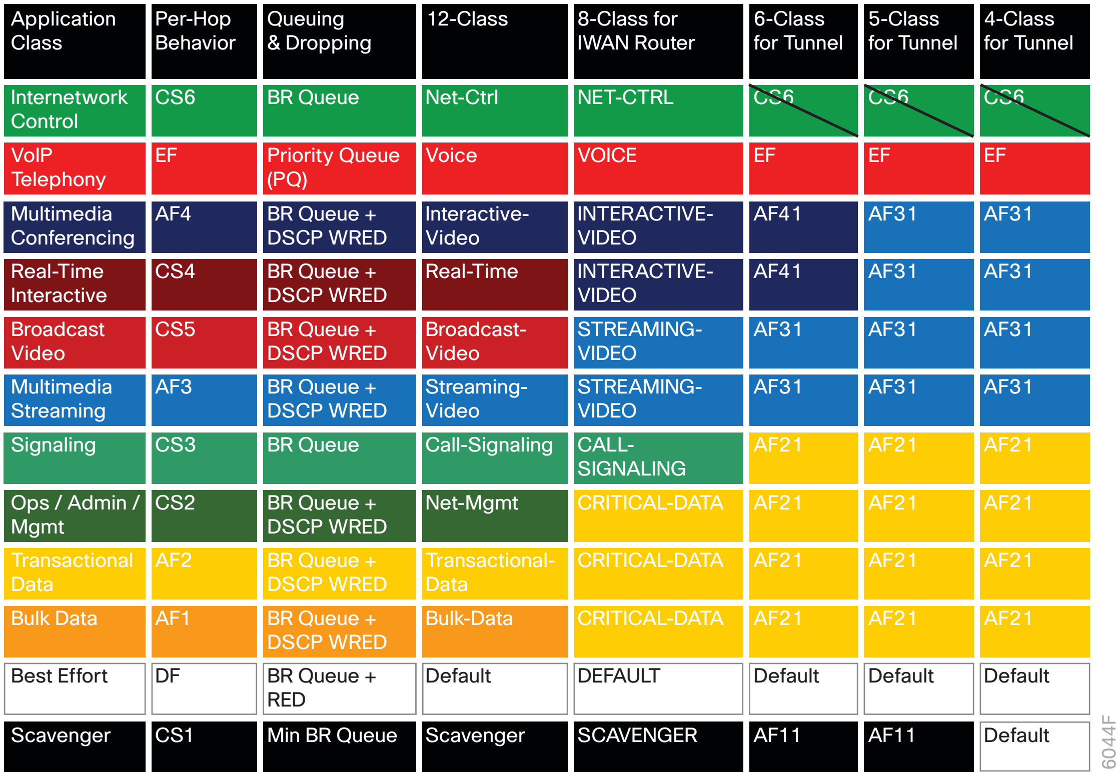

Application Class	Per-Hop Behavior	Queuing & Dropping	12-Class	8-Class for IWAN Router	6-Class for Tunnel	5-Class for Tunnel	4-Class for Tunnel
Internetwork Control	CS6	BR Queue	Net-Ctrl	NET-CTRL	~~CS6~~	~~CS6~~	~~CS6~~
VoIP Telephony	EF	Priority Queue (PQ)	Voice	VOICE	EF	EF	EF
Multimedia Conferencing	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

Hub Router QoS

The hub router uses a three-tier QoS scheduling hierarchy that consists of a child queuing policy and a parent shaping policy on the tunnel, along with a grandparent shaping policy on the physical interface. The child queuing policy does bandwidth sharing within the tunnel, while the parent shaping policy does bandwidth sharing between the tunnels and shaping for the remote sites inbound service rate. The grandparent shaping policy prevents the router from overrunning the outbound service rate purchased from the provider.

A three-tier hierarchy can lead to aggregate priority load issues, which happens if the sum of all the priority traffic from each child queuing policy exceeds the outbound service rate of the hub router. To avoid aggregate priority load problems, an explicit policer is recommended for priority traffic, rather than an implicit policer. The explicit policer is an always-on policer that does not require congestion for it to police the traffic to the defined value.

The DMVPN per-tunnel QoS feature on the hub router allows the configuration of the child queuing policy and the parent shaping policy to be signaled from the remote site router. With this type of simplified configuration, the hub site is prevented from sending more traffic than any single remote-site can handle. The figure below shows the three levels of QoS scheduling hierarchy for a hub router.

Figure 21 Hub Router QoS scheduling hierarchy

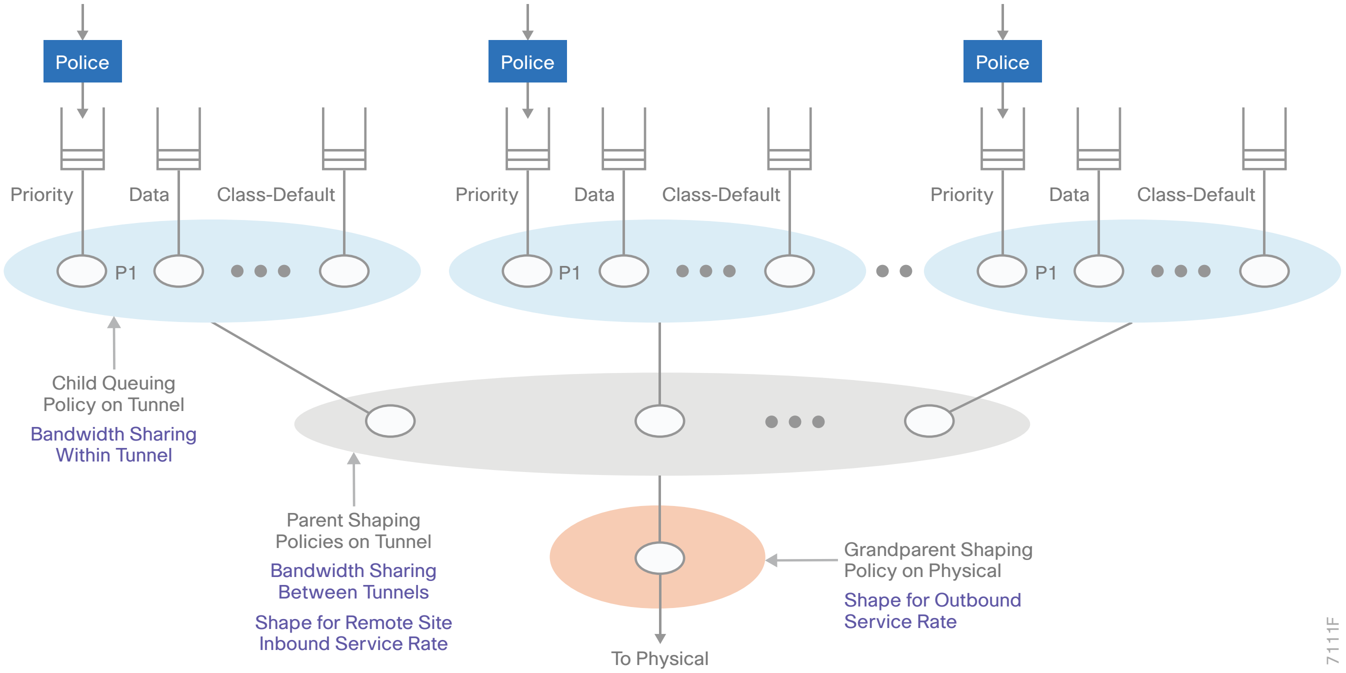

Remote Site Router QoS

The remote site router uses a two tier QoS scheduling hierarchy that consists of a child queuing policy and a parent shaping policy on the physical interface. The child queuing policy does bandwidth sharing, while the parent shaping policy shapes traffic to the remote sites outbound service rate. The figure below shows the two levels of QoS scheduling hierarchy for a remote site router.

Figure 22 Remote Site Router QoS scheduling hierarchy

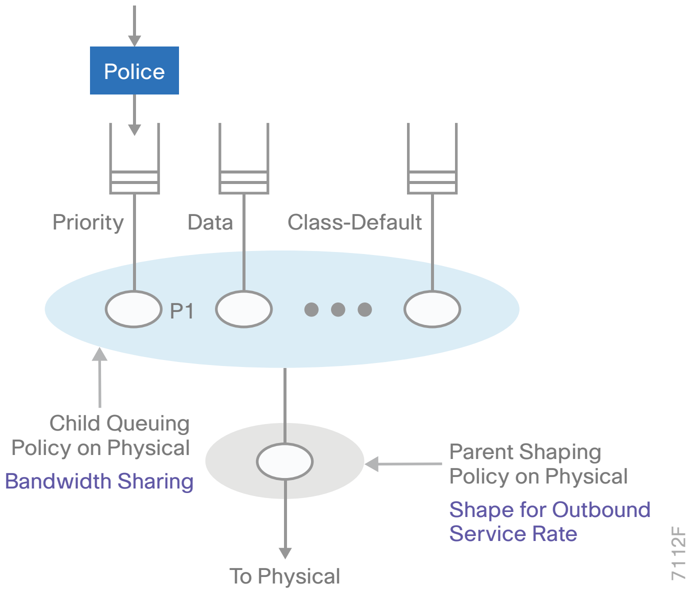

IWAN Best Practices

The next several sections cover best practices from the IWAN perspective.

ROUTING PRINCIPLES

It is a good practice to manipulate the routing protocols so that traffic flows across the preferred transport. Influencing the routing table ensures that when PfR is disabled, traffic will follow the Cisco Express Forwarding table derived from the Routing Information Base (RIB) and forward traffic to the DMVPN over the preferred tunnel.

PfRv3 always checks for a parent route of any destination prefix before creating a channel or controlling a traffic class. PfR selects next hops based on the following order of lookup:

- NHRP shortcut route (branch only)
- If not, check in the order of BGP, EIGRP, Static and RIB
- If at any point, an NHRP short cut route appears, PfRv3 would pick that up and relinquish using the parent route from one of the routing protocols.

It is essential to make sure all destination prefixes are reachable over all available paths so PfR can create the corresponding channels and control the traffic classes. Remember, PfR will check within the BGP or EIGRP topology table.

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
- Provide a solid underlying IP routed topology in order to support the Intelligent Path Control provided by Cisco PfR.
- Provide site-site remote routing via the primary WAN-aggregation site (hub-and-spoke model)
- Permit optimal direct site-site remote routing (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 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.

EIGRP Routing Overview

Cisco uses EIGRP as the primary routing protocol for IWAN because EIGRP 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 also supports wide metrics for 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.

This design uses a single EIGRP autonomous system (AS 400) for the LAN, WAN, and all of the remote sites. Every remote site is dual connected for resiliency. However, due to the multiple paths that exist within this topology, you must avoid routing loops and prevent remote sites from becoming transit sites if WAN failures occur. Interface delay is configured to make sure the WAN interfaces are always preferred and MPLS is the preferred WAN path in the hybrid design models.

There are several key recommendations for each type of site.

Hub and transit sites:

- **Disable split-horizon**—If split-horizon is enabled, the hub router will not advertise the networks learned from one spoke router to another spoke router.
- **Advertise site summary, enterprise summary and a default route to remote sites**—Scalability of the EIGRP routing domain is dependent upon summarization. Summarizing multiple routes to an aggregate reduces the size of the routing table and creates an EIGRP query boundary. EIGRP summarizes network prefixes on an interface basis.
- **Summary metrics**—Summary-metrics are used to reduce computational load on the DMVPN hub and transit border routers.
- **Ingress filter on tunnels**—An inbound prefix-list filter is applied on the DMVPN hub routers to prevent a DMVPN router from learning a default or summary route from a peer hub router which prevents sub-optimal or routing loops.
- **EIGRP hello and hold timers**—Increase the EIGRP hello interval to 20 seconds and the hold timer to 60 seconds. Increasing the timers allows for the DMVPN hub routers to handles a large number of remote sites. The hello and hold timers should match on the DMVPN hub and remote site routers.

Remote sites:

- **EIGRP stub-site**—The stub-site functionality builds on the EIGRP stub feature, which allows a router to advertise itself as a stub to peers on specified WAN interfaces. It also allows the router to exchange routes learned on LAN interface.
- **EIGRP hello and hold timers**—Increase the EIGRP hello interval to 20 seconds and the hold timer to 60 seconds. Increasing the timers allows for the DMVPN hub routers to handles a large number of remote sites. The hello and hold timers should match on the DMVPN hub and remote site routers.

The EIGRP stub-site feature provides the following key benefits:

- EIGRP neighbors on WAN links do not send EIGRP queries to the remote site when a route goes Active.
- Additional routers can be placed further in the site and still receive routes from the WAN through the stub-site router.
- Prevents the stub-site from becoming a transit router.
- Removes the need for a complex routing leaking with route tags and filtering.

BGP and OSPF Routing Overview

BGP can be deployed in the WAN overlay as an alternative routing protocol to EIGRP. BGP is a popular choice for network operators that require a rich set of features to customize path selection in complex topologies and large-scale deployments. Although BGP is traditionally positioned at the service provider WAN edge, recent enhancements such as BGP dynamic neighbors make it a viable choice for IWAN deployment. BGP dynamic neighbor support simplifies the configuration and allows peering to a group of remote neighbors that are defined by a range of IP addresses. Each range can be configured as a subnet IP address, which allows spokes to initiate the BGP peering without having to preconfigure remote peers on the route-reflectors.

This design uses a single iBGP routing process for the WAN overlay. OSPF is used on the LAN interfaces at the hub, transit, and remote sites.

There are several key recommendations for each type of site.

Hub and transit sites:

- **BGP route-reflectors**—DMVPN hub and transit routers function as BGP route-reflectors for the spokes.
- **BGP peering**—No BGP peering is configured between the route-reflectors.
- **BGP dynamic neighbors**—BGP dynamic neighbor support is configured on the route-reflectors.
- **Route advertisement**—Site specific prefixes, enterprise summary prefix, and the default route are advertised to the remote sites.
- **Local preference**—Set local preference for all prefixes based on the WAN transport hierarchy.
- **Route redistribution**—Redistribute BGP into OSPF with a defined metric cost to attract traffic from the central sites to the remote sites across MPLS.
- **BGP hello and hold timers**—Increase the BGP hello interval to 20 seconds and the hold timer to 60 seconds. Increasing the timers allows for the DMVPN hub routers to handles a large number of remote sites. The hello and hold timers should match on the DMVPN hub and remote site routers.
- **OSPF Area 0**—Configure an OSPF Area 0 for the LAN interfaces.

Remote sites:

- **BGP Peering**—Peering to hub and transit border routers for each DMVPN cloud.
- **Local preference**—Preferred path is chosen from the highest local preference.
- **Route redistribution**—Perform mutual redistribution of OSPF and BGP routes.
- **Route tagging**—Set a local route tag to identify routes in OSPF that were redistributed from BGP.
- **OSPF Area 0**—Configure an OSPF Area 0 for the LAN interfaces of dual-router remote sites or remotes sites with distribution layer switches.

When BGP is used, PfRv3 is able to check in the BGP database and will use the best path as computed by BGP. This path needs to be via an external WAN interface. If that is not the case, then PfRv3 will choose in sequence the path with biggest weight, then biggest local preference, and finally the path with the smallest IP address.

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 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. Remote sites with three to five transports must be able to tolerate the loss of multiple edge routers and WAN transports.

ENCRYPTION

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.

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 GRE.

Table 5 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.

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. 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 router 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. Cisco IOS provides anti-replay protection 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. In this 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 1024 packets.

IPsec uses a key exchange between the routers in order to encrypt/decrypt the traffic. You can exchange these keys by using a simple pre-sharing algorithm or 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

All use cases in the Cisco IWAN design have at least two DMVPN tunnel interfaces. The DMVPN MTT feature provides support for multiple tunnel terminations (interfaces) on the same hub border router. MTT is not required, but it is a feature which provides additional flexibility when deploying IWAN.

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 when site-to-site IP traffic flows require it.

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 indirect 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 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 but requires the use of IPsec transport mode to support a DMVPN hub behind static NAT.

Tech Tip

If the firewall is owned by the service provider, they will have to perform the same procedures and steps to allow DMVPN traffic into their DMZ, as described in the IWAN Deployment Guide.

The IWAN DMVPN design requires the use of Internet Key Management Protocol version 2 (IKEv2) keep alive intervals for dead peer detection (DPD), which is essential to facilitate fast re-convergence 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 and DPD timers reduces this convergence time to 40 seconds.

Figure 23 *DMVPN dual-cloud*

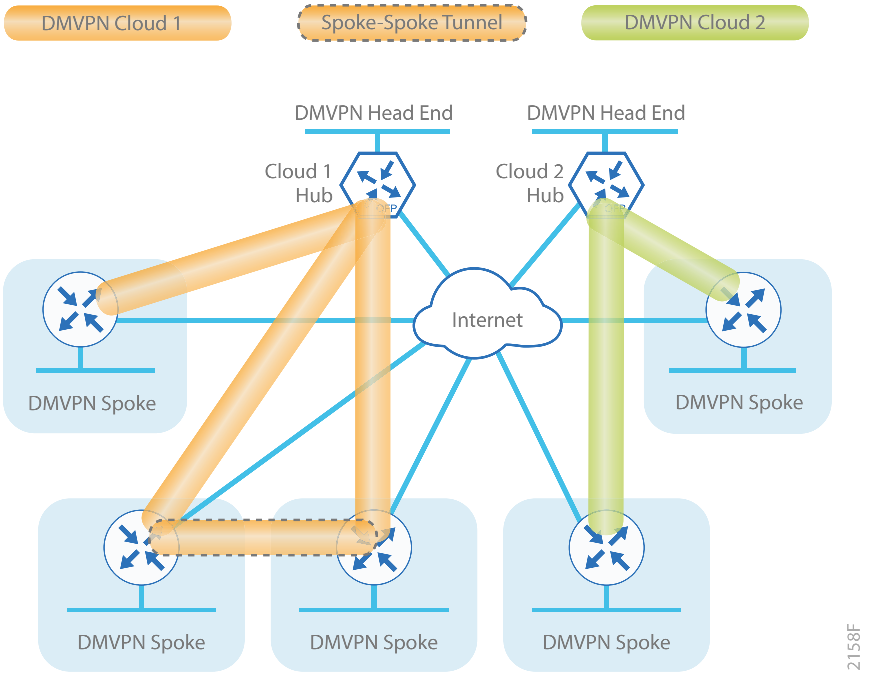

IWAN HUB BORDER ROUTERS

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 designs require at least two IWAN hub border routers to support a pair of DMVPN clouds. This pair of routers is required in order to provide resilient connections to all of the remote sites.

One of the key benefits of the DMVPN solution is that the spoke routers can use dynamically assigned addresses in their underlay network, 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 IWAN hub border 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.

IWAN REMOTE SITE ROUTERS

The WAN remote-site routing platform 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 single-router dual-link design provides a good level of redundancy for the remote site. This design can tolerate the failure of a link, which is automatically detected by the router and traffic is rerouted to the remaining path.

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 a router and traffic is rerouted to the remaining router.

The dual-router, five-link design provides the highest amount of redundancy and resiliency. This design can tolerate the loss of a router and all but one of the remaining paths.

Under normal conditions for all designs, the routing protocols are tuned to ensure the proper path selection and PfR takes care of channel selection for critical traffic.

Figure 24 *IWAN remote-site designs*

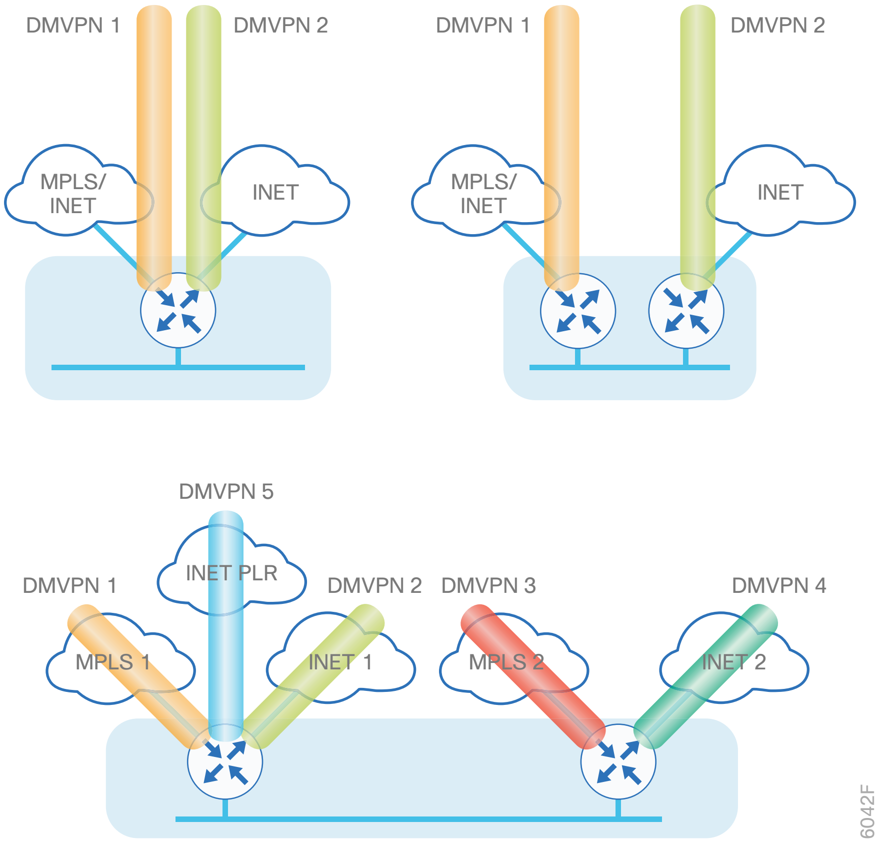

For the full list of IWAN supported routers for this version of the CVD, see the IWAN 2.2.1 release notes.

VRFS AND FRONT DOOR VRF

VRF is a technology used in computer networks that allows multiple instances of a routing table to co-exist within the same router. Because the routing instances are independent, you can use the same or overlapping IP Addresses without conflicting with each other.

VRFs are also important for isolating end-to-end user traffic among multiple independent networks inside the same data center—for example, private contractor networks and Internet of Things (IoT) networks operating in the same DC alongside the employee network. Similarly, at the remote location, private contractor computers and IoT devices can co-exist inside the same physical hardware as the common employee network. With the Multi-VRF feature in IWAN, the logical separation of traffic can be carried across the WAN.

IWAN uses VRFs to provide the following benefits:

- Default route separation between user traffic in the overlay network and the DMVPN tunnel establishment in the underlay network
- Control and data plane separation between inside and outside networks for security purposes
- Network segmentation of user traffic across the WAN in the IWAN Multi-VRF design model

The simplest form of a VRF implementation is VRF Lite. In this implementation, each router within the network participates in the virtual routing and forwarding 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 IWAN 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.

Reader Tip

For information about deploying Direct Internet Access from IWAN remote sites, see IWAN Direct Internet Access Design Guide.

The multiple default route conundrum is solved through the use of VRF Lite 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 in the IWAN overlay network, 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 the global VRF for user traffic routing in the overlay and a VRF for each WAN physical interface for DMVPN tunnel establishment in their respective underlay networks. 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.

Figure 25 *Front door VRF*

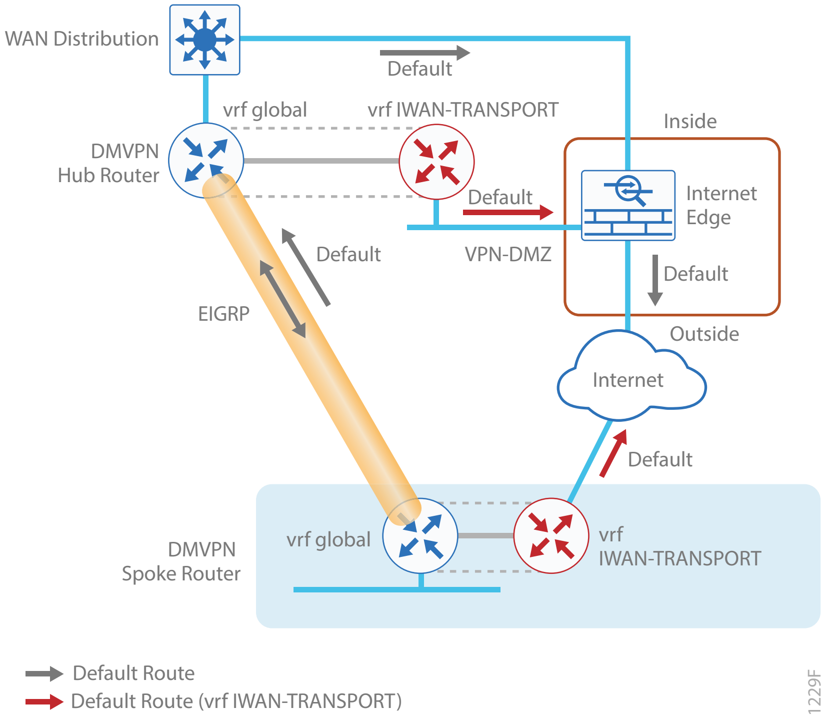

In the IWAN design models, the DMVPN hub routers must have sufficient IP-routing information in order to provide end-to-end reachability. Maintaining this routing information typically requires a routing protocol, and EIGRP or BGP are recommended for this purpose.

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 hub uses default routing to the MPLS provider edge router, and the Internet 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 26 *IWAN hybrid design model: FVRF default routing*

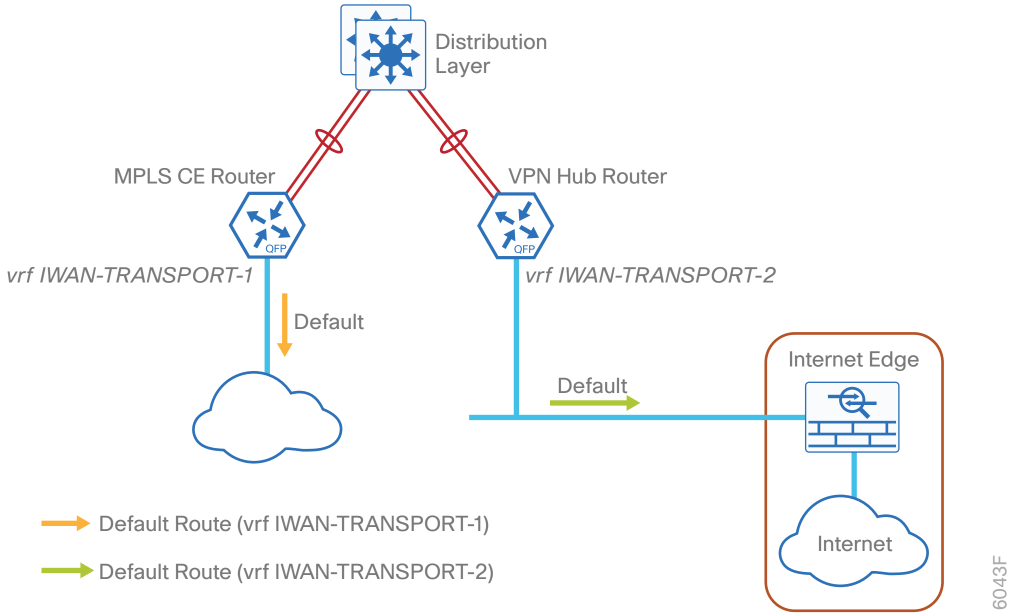

Figure 27 *IWAN dual Internet design model: FVRF default routing*

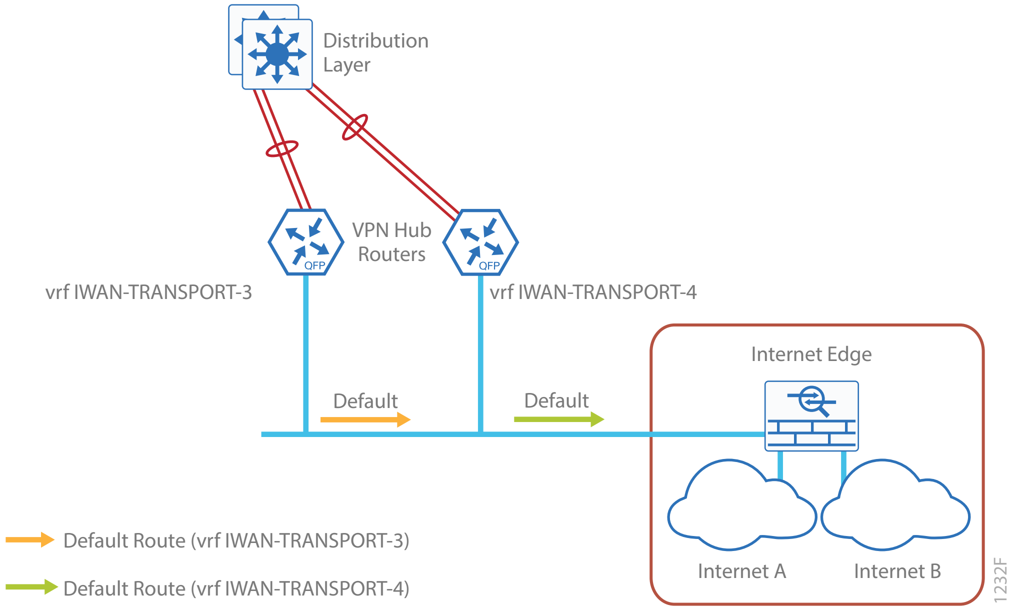

Summary for IWAN

Cisco enterprise IWAN architectures are proven solutions that scale to all remote-site sizes over MPLS, Internet and 4G/LTE transports. With rich application and security services on a single platform, IT can scale to hundreds of sites. Also, customers can maintain granular control, from the remote site, to the data center, and out to the public cloud. The traffic is dynamically routed based on application, endpoint, and network conditions in order to help ensure the best user experience. IT can confidently roll out critical business services such as consolidated data centers, SaaS, IP telephony, and video without overwhelming the WAN.

Reader Tip

For more information about deploying the Intelligent WAN, see the Design Zone for Branch WAN.

Appendix A: Changes

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

- Added the IWAN Multiple VRF design model
- Updated the QoS design information to align with recent changes

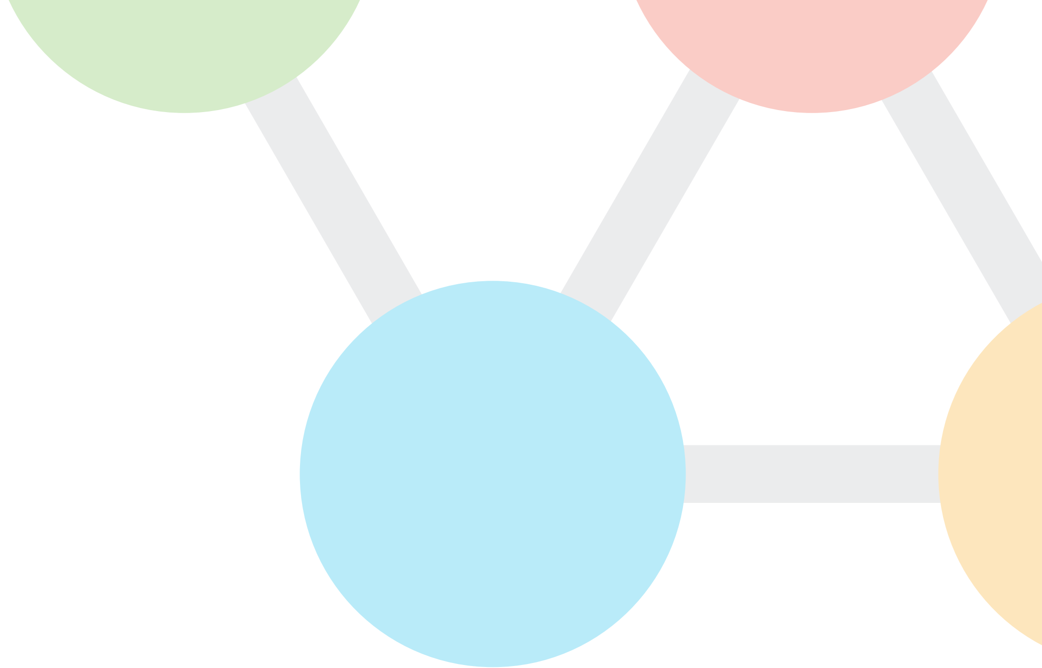

You can use the feedback form to send comments and suggestions about this guide.

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.