Overview of the Cisco Next-Generation WAN Architecture

Last Updated: April 17, 2013

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

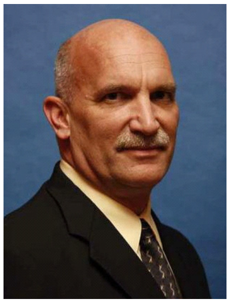

Tom Schepers

Scott Van de Houten

Tom Schepers, Director, Borderless Networks Architecture, Cisco Systems

Tom Schepers is the Director for the Borderless Networks Architecture team. Tom is responsible for next generation enterprise network architectures and sales enablement that relates to Borderless Networks architecture. Prior to his current assignment, Tom was a Distinguished Systems Engineer for the US Enterprise segment, focused on WAN architecture and technology. During his 17-year tenure at Cisco, Tom has also worked as a Systems Engineer and Consulting Systems Engineer, focused on large-scale network design for Enterprise.

Scott Van de Houten, Distinguished Architect, Borderless Networks Architecture, Cisco Systems

Scott Van de Houten is a Distinguished Communications Architect on the Borderless Networks technical strategy team focusing on enterprise network architectures. Scott has been with Cisco for 20 years and has over 25 years of industry experience. Scott has a broad range of expertise and has held many different positions at Cisco including Product Line Manager for IOS routing protocols as well as multiple Systems Engineering roles in Cisco Federal Sales.

About Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit http://www.cisco.com/go/designzone.

Overview of the Cisco Next-Generation WAN Architecture

Overview of the Cisco Next-Generation WAN Architecture

The Cisco Next-Generation WAN (NGWAN or Next Gen WAN) for enterprises is a comprehensive architecture developed for large enterprise and public-sector entities that targets large-scale routed wide area network (WAN) deployments. The architecture encompasses branch and metro connectivity, IP and Multiprotocol Label Switching (MPLS) core backbones, and emerging enterprise edge functions.

Related Documents

This document provides the initial details of the architecture, the component architectures, and the integration of Cisco Borderless Networks services on IPv4 and IPv6 infrastructures.

Several companion documents located under the Regional WAN and Enterprise Edge tabs provided at http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns816/landing_overall_wan.html provide a more detailed definition of the component architectures and their designs. Deployment and configuration guides are provided as well as specific supplemental guides for detailed deployment recommendations for specific items. The release of some of these documents will follow in subsequent phases.

- Regional WAN At-A-Glance
- Regional WAN Deployment Guide
- Best Practice Guide (Cisco PfR, AVC, and WAAS)
- IPv6 Migration Deployment Guide
- Regional WAN Management Supplemental Guide
- Video Deployment Guide
- Remote Access VPN Design Guide
- Remote Access VPN Deployment Guide
- IPv6 Deployment at the Internet Edge At-A-Glance
- DMVPN-to-GET VPN Migration Guide
- Cisco Performance Routing Deployment Guide (Coming Soon)
- CWS Deployment Guide (Coming Soon)

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

- RWAN TrustSec Supplemental Design Guide (Coming Soon)
- Unified Collaboration Supplemental Guide (Coming Soon)
- AppNav Design Guide (Coming Soon)
- LTE Deployment Guide (Coming Soon)
- BYOD For the Branch Whitepaper (Coming Soon)
- VDI-In-A-Box Guide (Coming Soon)
- CSR Whitepaper (Coming Soon)
- WAN Core Solution Overview (Coming Soon)
- Metro Solution Overview (Coming Soon)
- Enterprise Edge Solution Overview (Coming Soon)

Introduction

This document provides an architecture overview of the NGWAN for enterprises and takes an architectural approach to designing and deploying routed WAN solutions for Borderless Networks. This document provides a breakdown of the architectural components and an overview of the Borderless Networks services that are integrated within the architecture.

This document is intended for the reader who is interested in these topics:

- Large national or global WANs
- Regional branch networks with up to 15,000 sites
- MPLS VPN, Internet, and third and fourth generation (3G and 4G) transport for branch connectivity
- Metro networks using Carrier Ethernet or self-deployed optical networks of up to 100 sites
- Enterprise edge with collaboration (voice and video) or public cloud requirements
- WAN core networks within theater (country) and global, based on self-deployed MPLS or IP
- Incorporating security for privacy and regulatory requirements
- Addressing the need for mobility of users and machines
- Migration requirements for IPv6
- Upgrading or addressing growth to new regions or countries

Architecture Benefits

NGWAN is a comprehensive architectural approach to defining the enterprise routed WAN. You can realize several benefits when you implement a Next Gen WAN:

- Reduced operating costs
- Increased scalability
- Improved security and risk management
- Improved application performance
- Maximized network availability
- Flexibility and “services on demand”

- Reduced overall complexity
- Faster time to market for new IT services
- Accelerated merger and acquisition return on investment (ROI)
- Regulatory compliance

Architecture Overview

The NGWAN architecture is modular and hierarchical, and it provides a scalable solution across the enterprise customer segments as shown in Figure 1.

Figure 1 *High-Level Architecture Abstract*

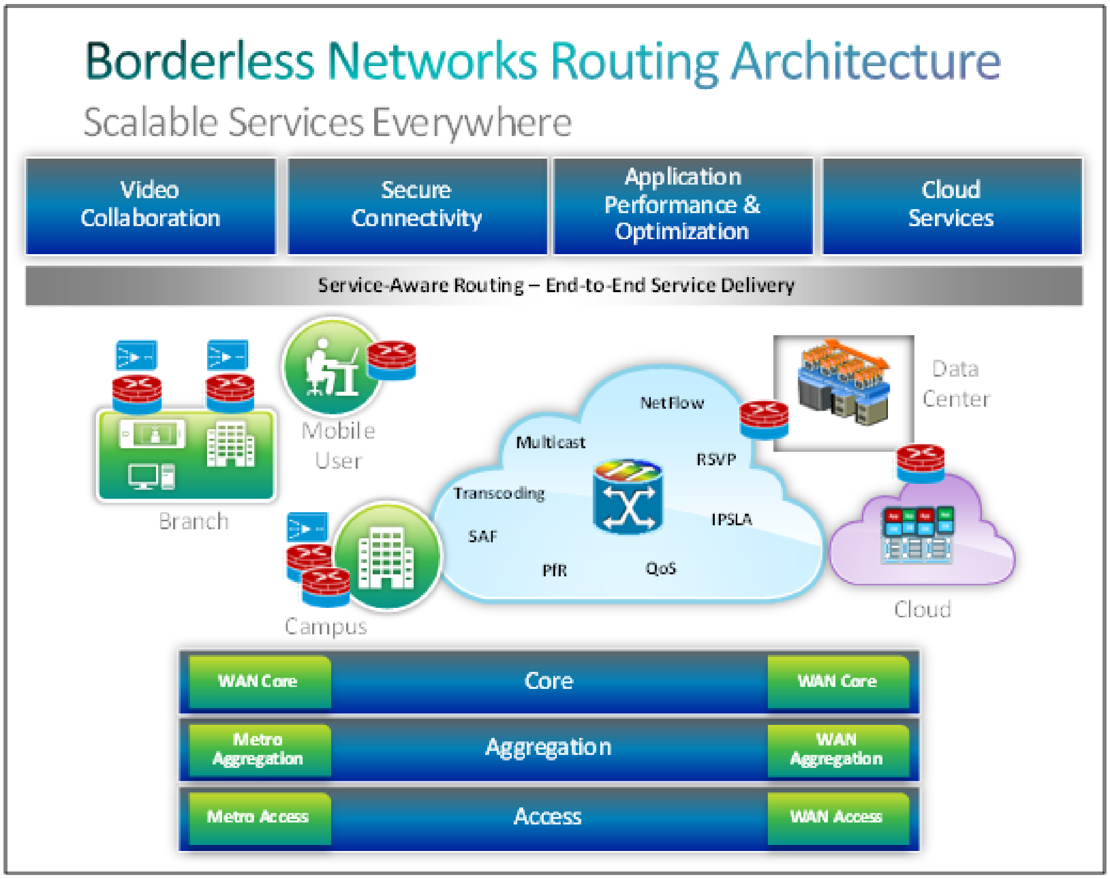

The NGWAN architecture is comprised of five core modules:

- Regional WAN: Used to connect branch offices and aggregate remote locations
- Metro: Used to connect remote offices and data centers across metro transports
- WAN core: Used to interconnect regional networks and data centers within a country, theater or globally
- Enterprise edge: Used to connect the enterprise network to other external networks and services
- Enterprise interconnect: Connects all WAN, campus, and data center network modules together

This modular and hierarchical approach allows a NGWAN to be designed that uses basic building blocks to increase availability and scalability. Availability is increased by providing fast convergence and fault isolation within a building block. The three-tier hierarchy allows the enterprise network to scale from a regional two-tier design up to a global three-tier network design. These NGWAN modules are tied together, along with campus and data center modules, using the enterprise interconnect, which acts as an enterprise distribution network. The WAN core peers with multiple regional interconnects to scale the network across the country or theater. In the largest cases, a top-tier global core backbone can interconnect such theater networks.

The NGWAN architecture also focuses on the enterprise edge and includes emerging strategies such as connectivity to the cloud (private, public, and hybrid) and collaboration services, as well as business-to-business rich-media capabilities such as telepresence.

The NGWAN architecture provides more than advanced routing functionality. Another major goal of the architecture is to deliver the applications and services that are relevant and fundamental to the enterprise business along with the routing elements that provide reliable delivery of those services. The architecture delivers services like these:

- medianet
- TrustSec
- Application Velocity and Control (AVC)
- cloud connectors
- IPv6

The incorporation of these services within the routing architecture is an important consideration when considering a WAN solution. The NGWAN architecture allows for services to be added as they are needed, and large-scale replacement of equipment or redesign is unnecessary.

Finally, the NGWAN architecture examines the linkages between the Cisco Virtualization, Collaboration, and Service Provider architectures. The NGWAN includes linkages to the enterprise Virtualization and Collaboration architectures, which provides additional functionality and true end-to-end value. The NGWAN leverages the Service Provider architectures to use carrier-class network technologies and products to build high-performance, highly available enterprise networks.

NGWAN Architecture Modules

The NGWAN architecture has multiple tiers:

- Access: regional WAN and metro
- Aggregation: interconnect
- WAN core: in-theater and global cores

An additional access module provides external access through the enterprise edge (next-generation Internet edge). Today the enterprise, or Internet, edge is centralized, but in the future we can expect a more distributed method for accessing public networks for cloud, mobility, and collaboration. This hierarchical structure allows the architecture to be separated into elements that are relevant to any customer environment. When a global footprint is required, all aspects of the architecture can be used. On the other hand, a footprint that is contained solely within a single theater does not require the global core, but it can be added when the network expands into other theaters.

The NGWAN architecture also introduces the concept of the enterprise interconnect. The enterprise interconnect connects the regional WAN, metro networks, enterprise edge, and WAN core to integrate them into a single architecture. The interconnect also can connect the local data center and campus that may be co-resident at the location. Figure 2 depicts the NGWAN architecture and the individual components.

Figure 2 Hierarchical Architecture Overview

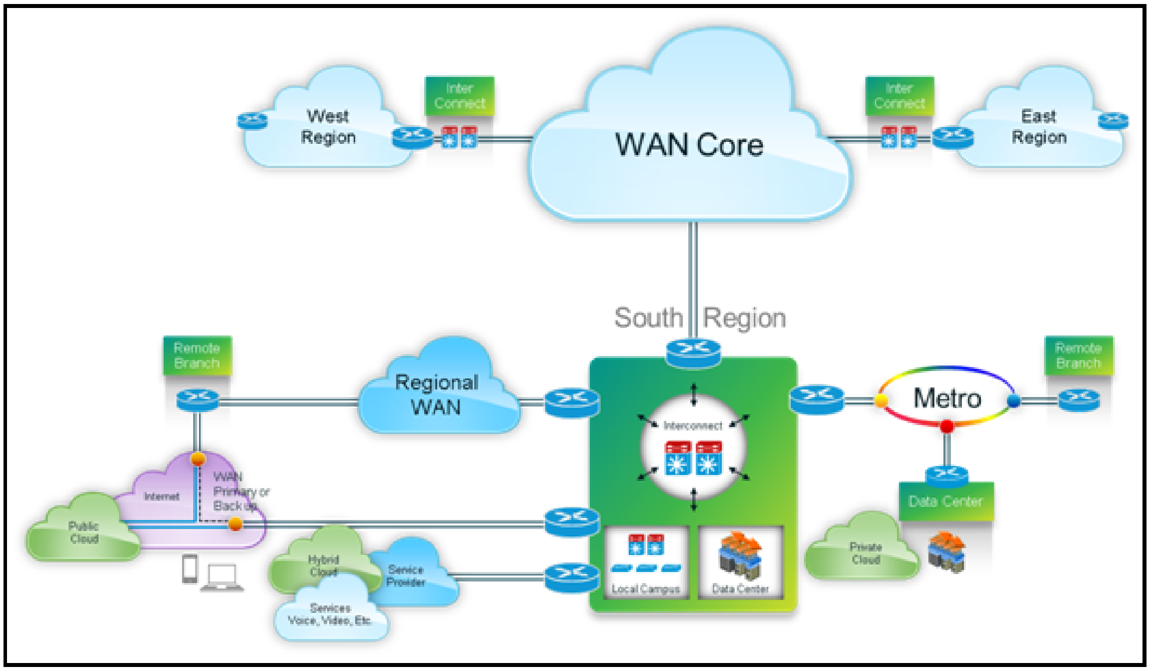

To fully understand the architecture and the ability to provide Borderless Network services, read the next sections, which describe what each module contains.

Regional WAN Module

The regional WAN combines branch access and WAN aggregation routing into a single network system. Branch locations in this context are remote locations that provide connectivity to a central site, typically campus or data center locations (see Figure 3 on page 10):

- Standard branch: A standard branch requires application availability and performance, but it does not require the higher-level of availability of a high-end branch. Although a standard branch may or may not have telepresence installed, it has requirements for collaboration and video services.
- High-end branch: A high-end branch is a branch where the customer demands the latest technology for collaboration and application availability, regardless of size. A high-end branch also encompasses more services that support critical applications, such as telepresence, IP video surveillance, and digital signage. This branch is critically important, so high availability and application performance capabilities must be provided.
- Ultra high-end branch: An ultra high-end branch requires the highest availability and performance, above those of the high-end branch. This location has a high user count, more like a remote campus. An ultra high-end branch has very high bandwidth and encryption requirements, usually at speeds of OC-12 or line-rate Gigabit Ethernet, and beyond. The ultra-high-end branch has such high

performance and availability requirements, so it employs appliances to deliver WAN optimization, public switched telephone network (PSTN) breakout, and other similar services. Support for multiple simultaneous high-definition (HD) video streams is also required.

- Mobile branch: A mobile branch takes advantage of a wireless WAN (WWAN) connection that allows for the entire branch to move. Temporary workspace and locations without wired connectivity are examples of the mobile branch. Mobility of the entire branch is the primary requirement, but application performance and security are also important over the high-latency links of cellular 3G, 4G, or satellite.

Figure 3 Regional WAN Architecture

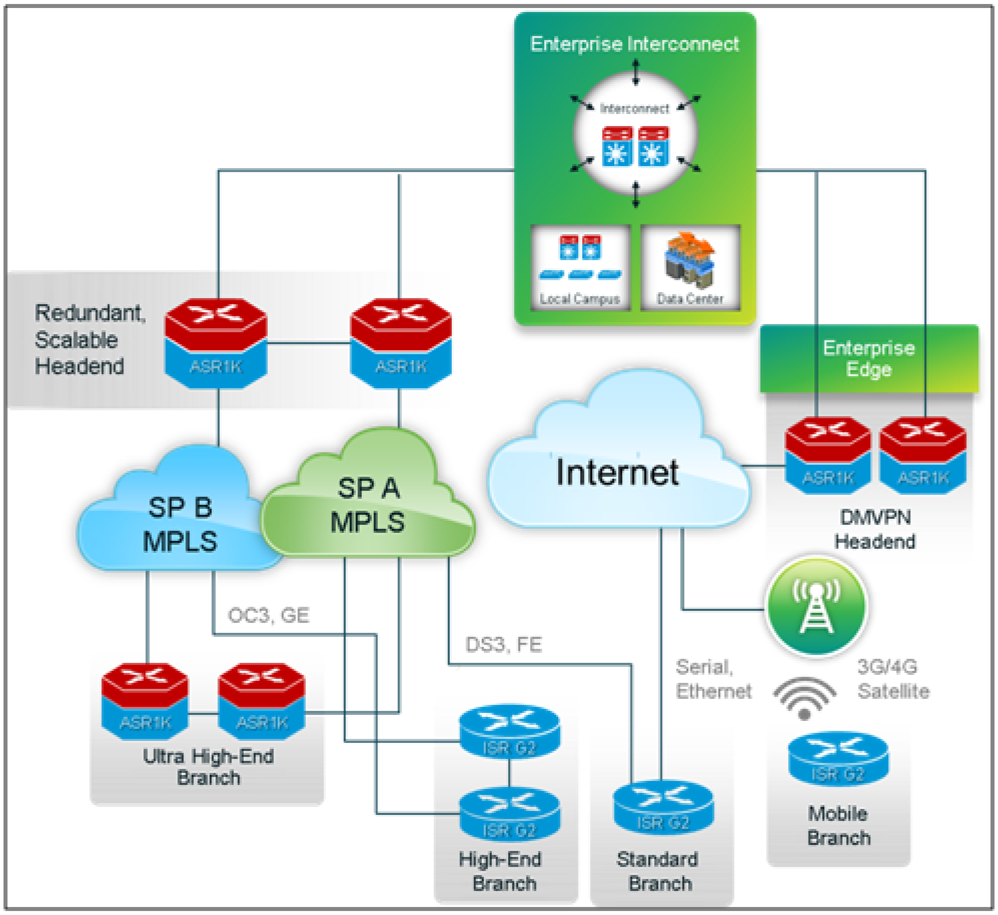

Metro Module

The metro networks are divided into the access and transport models. A typical metro access model is based on the regional WAN but uses a Layer 2 Ethernet hand-off instead of a Layer 3 IP hand-off from the service provider. Self-deployed metro transport models use privately owned or leased fiber assets to deliver a richer set of services as part of the NGWAN infrastructure. Enterprise metro networks get their primary requirements from the Cisco Data Center Interconnect (DCI) services used across the metro network.

- The metro access model provides connectivity to remote offices and data centers over Layer 2 Metro Ethernet service offerings (Figure 4). The Metro Ethernet service offering could be provided by a service provider or by the enterprise-owned metro network (Figure 5 on page 12). The metro access branch and aggregation design is very similar to the regional WAN design except that Metro Ethernet services are typically Layer 2 Ethernet and offer higher speeds.

Figure 4 Metro Network Access Models

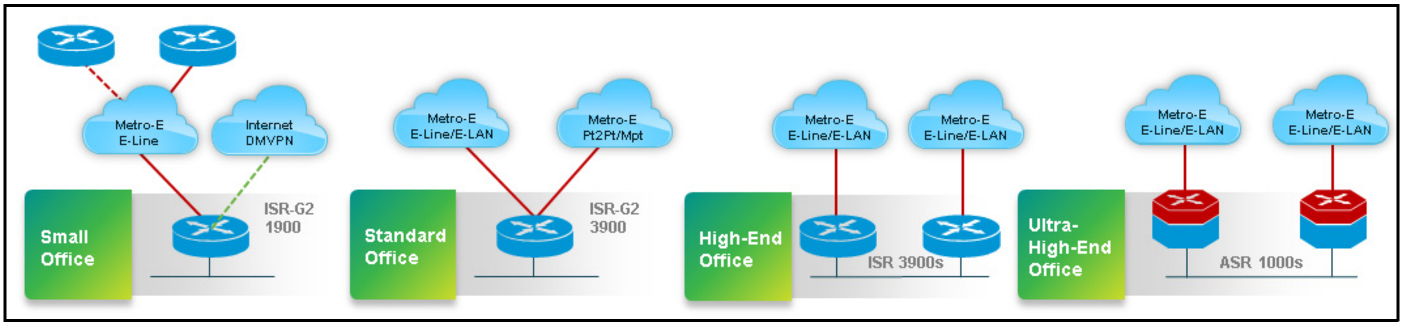

  - The small branch models use carrier-supported Metro-E over a service such as Ethernet line (E-Line) between sites. The typical deployment model employs E-Line point-to-point service with redundancy to a failover Ethernet circuit or to an IPsec DMVPN network in an active/standby model.
  - The standard branch model offers redundant active/active E-Line or Ethernet LAN (E-LAN) Ethernet connections. The E-LAN service is a multipoint Ethernet service that emulates a LAN between sites. E-LAN service-level agreements (SLAs) are strictly on "access" information rates into the Ethernet cloud with rate-limits on multicast and broadcast traffic.
  - The high-end branch model builds on the standard branch with redundant routers with active/active E-Line or E-LAN Ethernet. The ultra high-end, campus or critical branch model adds higher performance routers and higher bandwidth Ethernet circuits to meet the needs of the users at this site.

Figure 5 Metro Access Network Architecture

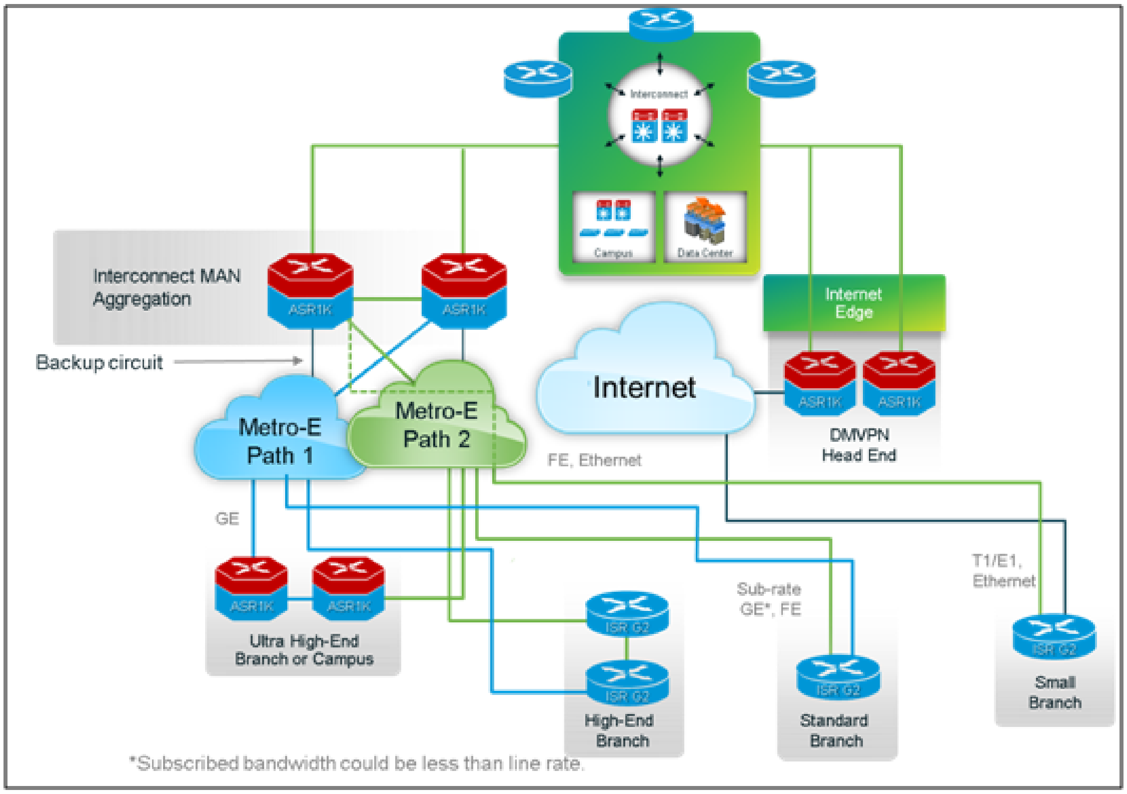

- The self-deployed metro transport models are used when the network service requirements justify a privately owned and managed metro network infrastructure. Critical drivers for metro networks include data center interconnect, legacy time-division multiplexing (TDM), storage networking, and HD video. Cost-effective access to fiber can easily justify deploying an enterprise-owned metro transport network for these drivers. The metro access models discussed would leverage an enterprise metro transport in the same manner as a service provider metro service. The self-deployed enterprise metro transport models in this architecture include the MPLS and Metro-E, optical, and enterprise service provider (ESP) designs (Figure 6 on page 13).
  - The enterprise MPLS and Metro-E deployment model includes an enterprise-owned and operated MPLS and Metro-E network to interconnect remote sites and data centers. These networks provide higher availability, faster change control, and flexible network virtualization options.
  - The enterprise optical deployment model provides benefits similar to those of the previous model plus optical transport of other non-IP traffic types, such as native Fibre Channel, native video, and TDM circuit services.
  - The ESP model combines a self-deployed optical and an MPLS and Metro-E metro network for the customers who need the maximum availability, capacity, virtualization, and transport flexibility. The scale and autonomous users of this type of network typically dictate a service provider operating model, thus the name ESP design.

Figure 6 Metro Transport Network Architecture

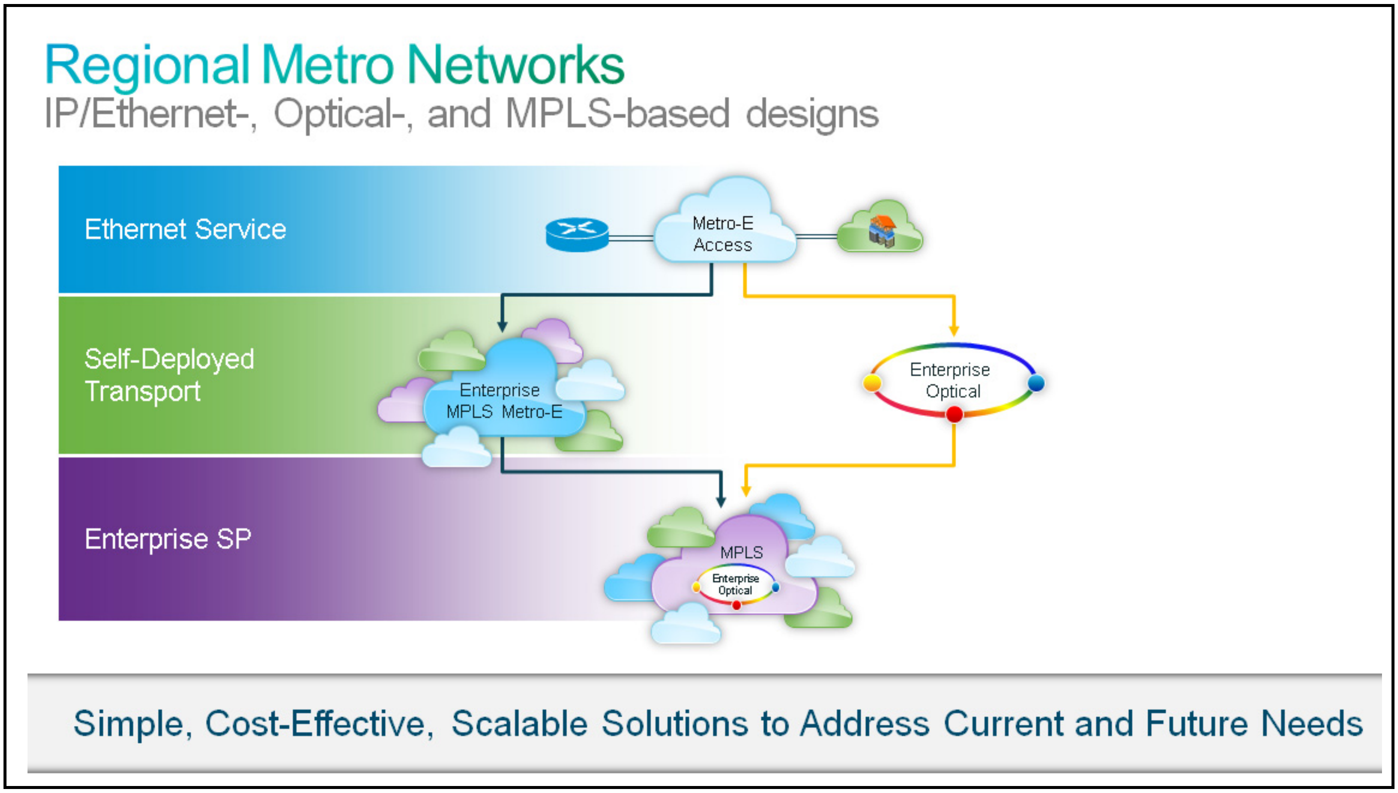

Enterprise WAN Core Module

Often, enterprises have multiple backbones for different services, such as, user, DCI, storage, and TDM. The convergence of applications and systems to IP and Ethernet with advances in network virtualization (Layer 2 and Layer 3 VPNs) has made it possible to collapse multiple isolated backbones into a single converged core network. Converged cores can reduce transport circuit costs, simplify operations, and increase the efficiency of network assets. Converged cores can be based upon IP or MPLS technologies. IP cores are very common and can address the majority of enterprise requirements. MPLS cores are common when the scale and operating requirements dictate service provider-like operating model. The additional capabilities of an MPLS core increase the complexity and operating costs as compared to an IP core.

The enterprise interconnect provides the connectivity between the regional WANs, MANs, data centers, enterprise edge, and campus networks. The WAN core network provides connectivity between regional enterprise interconnects within a theater as well as globally between theaters (Figure 7). The WAN core networks are sometimes referred to as the WAN backbone. Figure 2 on page 9 shows how the WAN core is positioned within the architecture.

Figure 7 Enterprise WAN Core Topology

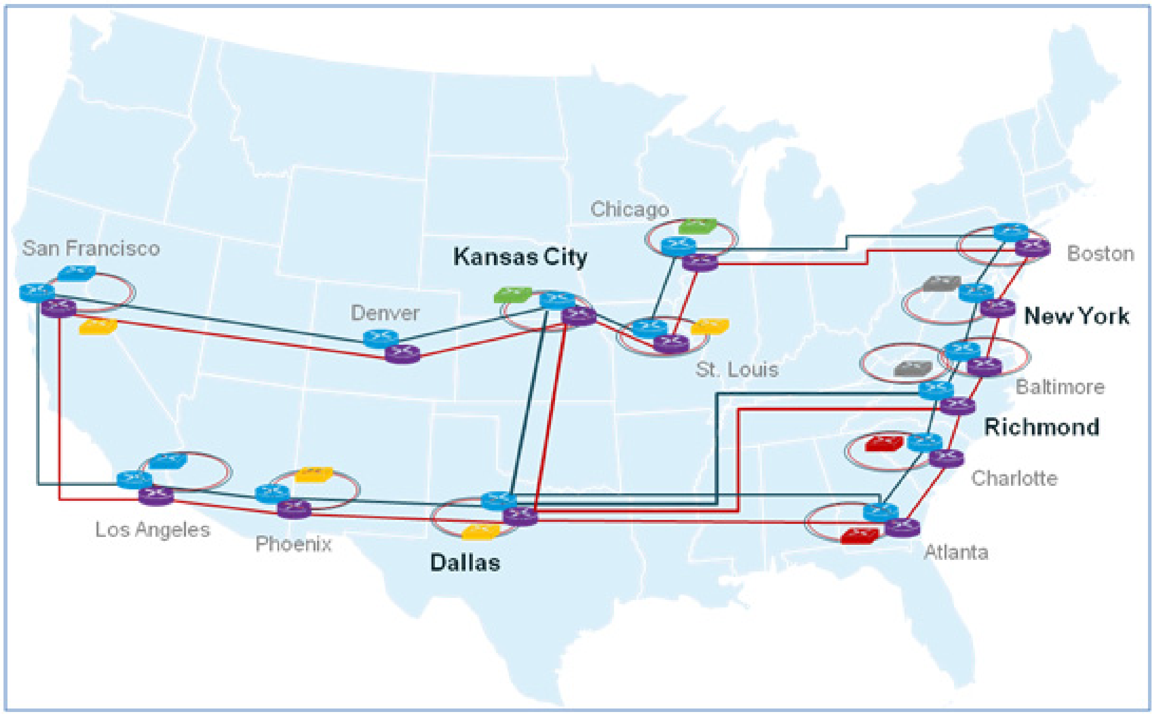

The in-theater core provides connectivity within a territory with common requirements, geography, service providers, and regulatory authorities. For example, the European core could provide connectivity between regional networks in Europe and the U.S. core could provide similar connectivity in the United States.

A global core provides connectivity between in-theater cores, which allows for scalability of the architecture on a global basis.

The vital nature of the enterprise core network requires that it be high-performance, highly available, and resilient. The enterprise core network must balance the costs of transport circuits, bandwidth, and network operations. The use of separate cores for in-theater and global networks provides additional scalability, better fault isolation, and policy control. In addition, multiple planes (multiplanar cores) can be used to improve availability of up to six nines for in-theater and global cores.

Enterprise Edge Module

Requirements for IP connectivity outside of the organization are changing. Enterprises must consider the outside edges of their networks to be something beyond simple, highly centralized Internet access. These new edge requirements mean new types of external peering relationships for voice, video, and cloud services that require new SLAs. These new requirements also introduce new security concerns and have more instances of external connectivity. This situation requires that the Internet edge evolve into a new multiservice “enterprise edge”.

The enterprise edge is defined as the interface between the controlled enterprise network, users, or resources and those that are outside of the enterprise’s control or visibility as shown in Figure 5 on page 12. Users are employees, partners, customers, and guests. Resources are Internet access, business-to-business connectivity, collaboration resources, remote user access, and application services. The enterprise edge is also a place where services such as security and collaboration acceleration reside (Figure 8).

Figure 8 Enterprise Edge

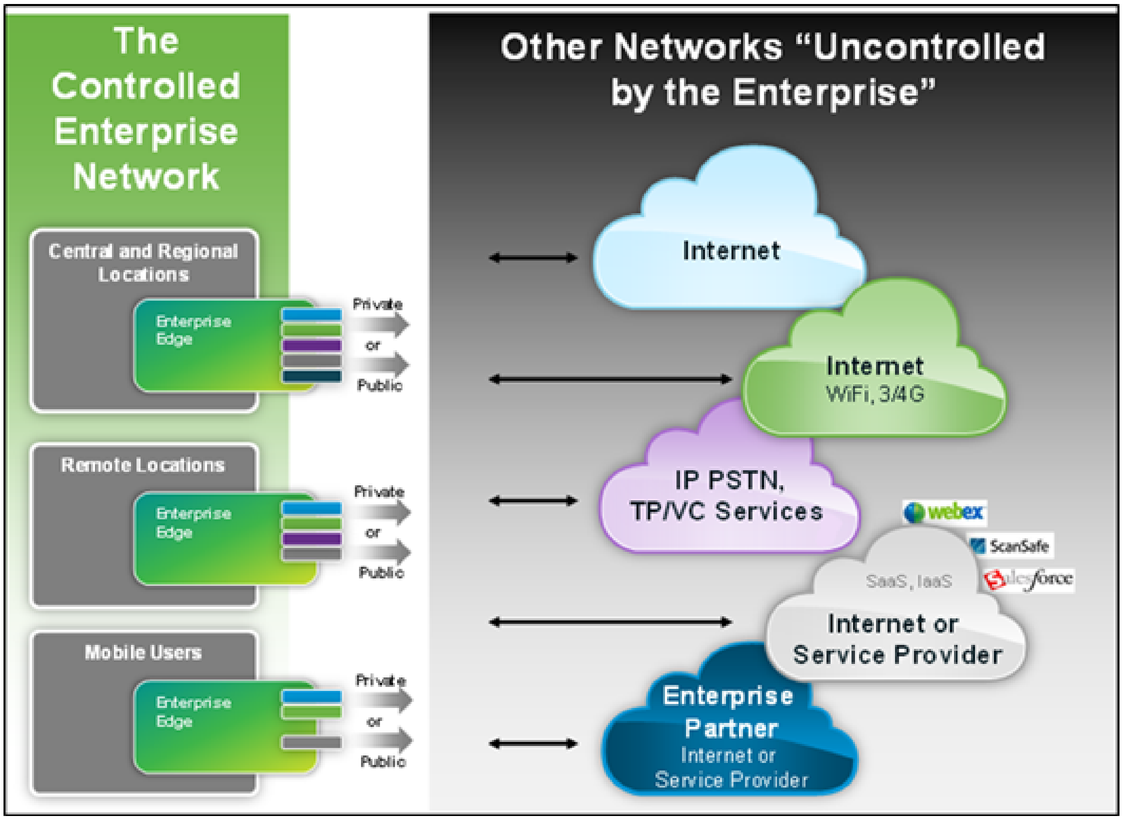

The enterprise edge is divided into five submodules:

- web services
- mobility services
- communications services
- cloud services
- business-to-business services

These submodules are installed optionally at each location where the enterprise edge connectivity is required, such as central sites, regional hubs, or branch offices (Figure 9). Mobile devices that have simultaneous connectivity to corporate and Internet resources can be considered an enterprise edge as well.

Figure 9 Enterprise Edge Component View

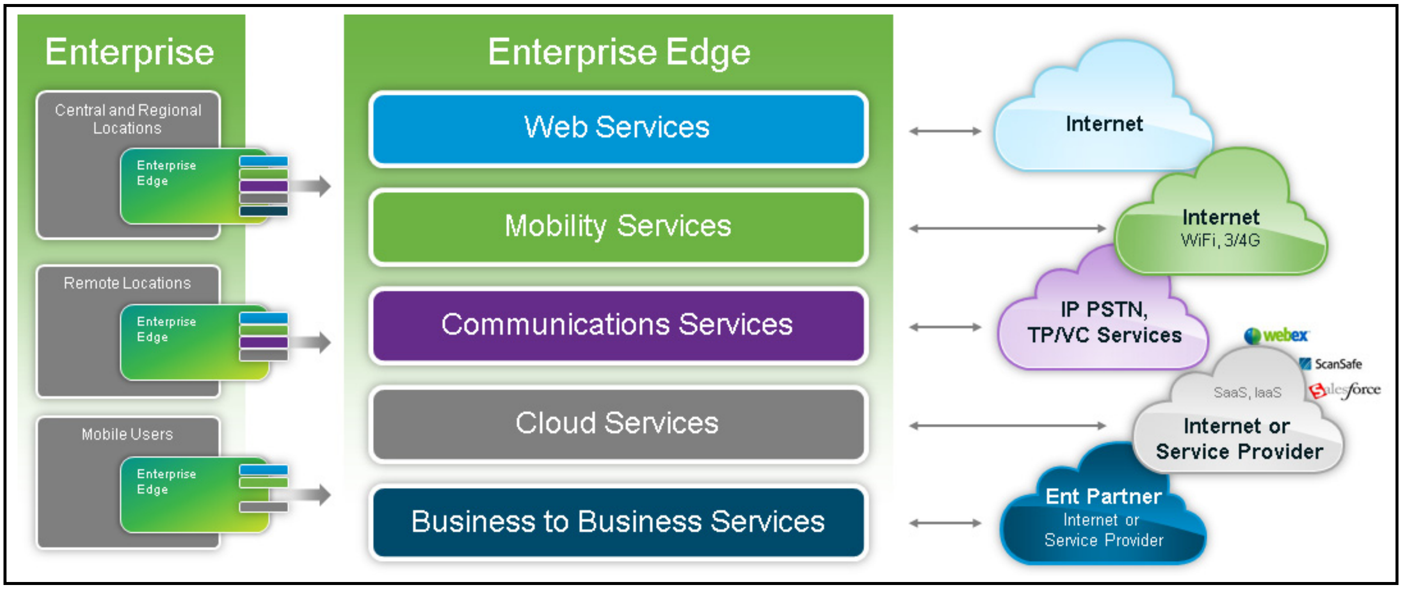

Enterprise Interconnect Module

The enterprise interconnect is the location in the NGWAN where the WAN core, regional WAN, metro, enterprise edge, and local data center or campus network modules are interconnected. The interconnect may consist of any or all of the architecture modules, depending upon the requirements and overall enterprise architecture (Figure 10). The interconnect is the network peering point between modules and not necessarily a physical part of the network infrastructure.

Figure 10 Enterprise Interconnect Framework

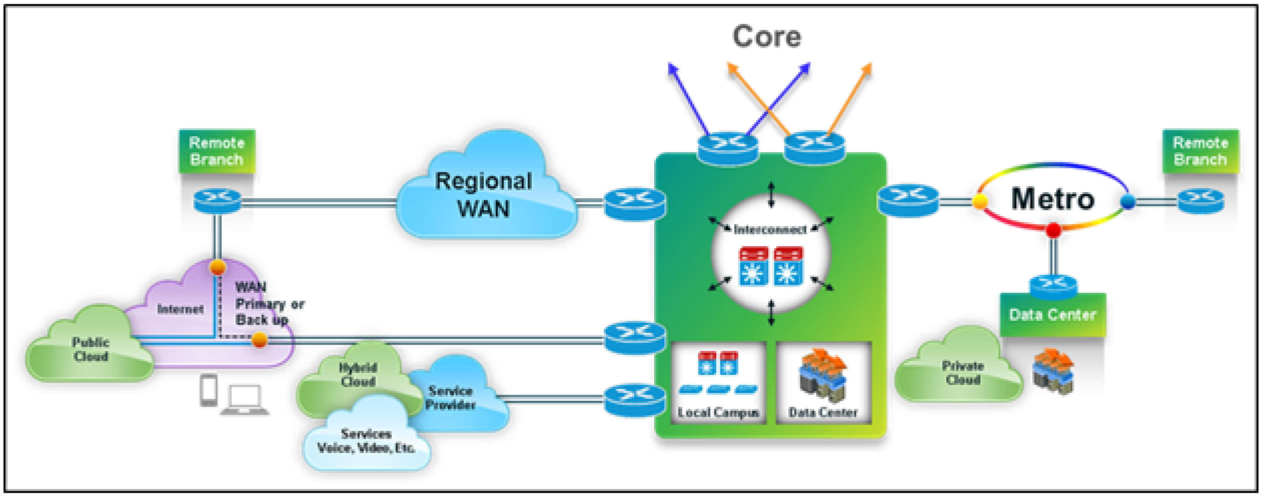

Within a regional location, the interconnect is the peering point between the regional WAN, metro, campus, and data center network modules. The WAN core module provides the connectivity between multiple regional network locations within a theater and globally across multiple theaters.

The enterprise interconnect infrastructure delivers services between the other component modules of the architecture, which provides the end-to-end integration of Borderless Network services, such as medianet. The interconnect supports any co-located data center or campus networks, but those component architectures (campus and data center) are well-defined elsewhere and are outside the scope of this document.

Borderless Network Services

Cisco has positioned a framework for the next-generation network called the Borderless Network Architecture. NGWAN is an underlying foundation for the Cisco Cloud Intelligent Network (CIN), which is one of the three pillars of the Borderless Networks framework in Figure 11. The NGWAN architecture enables the Borderless Network systems and services over integrated IPv4 and IPv6 infrastructures.

Figure 11 Borderless Network Architecture

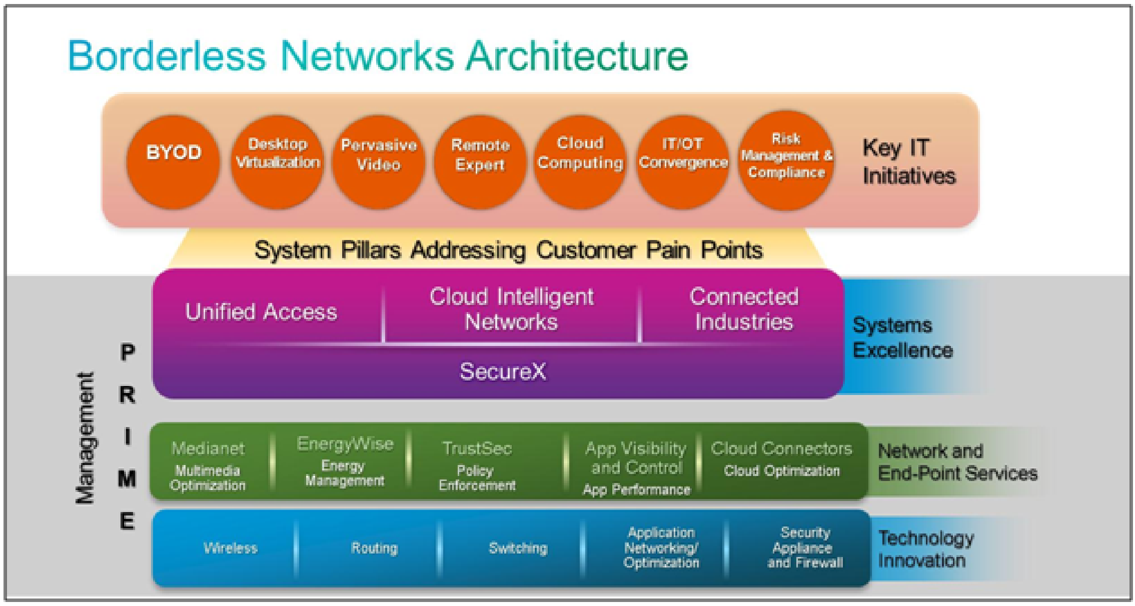

The key components of the Borderless Networks Architecture that are used in the NGWAN are described in the next sections.

Cisco Prime Management

Cisco Prime provides a portfolio of network management services that enables enterprises and service providers to more effectively deploy and operate their networks, speed troubleshooting, and improve operational efficiency. Cisco Prime products offer many advantages, including these benefits:

- Converged management and troubleshooting of wired and wireless access, branch, and wide area networks
- Network lifecycle management: User access visibility, inventory, configuration management, "plug-and-play", radio frequency planning, and best practices reporting
- End-to-end application and service assurance visibility for isolation and troubleshooting of performance issues
- Simplified deployment and management of Cisco advanced technologies such as VPN, Zone-based Firewall, ScanSafe, and Application Visibility and Control (AVC)
- Configuration, provisioning, and troubleshooting of Carrier Ethernet, Unified MPLS, and converged packet-optical transport networks

The Cisco Prime portfolio of products is extensive. However, these key products, or suites of products, are several that may be used in the NGWAN:

- Cisco Prime for IT: This product suite simplifies network management, improves operations efficiency, reduces errors, and increases network predictability. Cisco Prime for IT includes Cisco Prime Infrastructure, Cisco Prime Collaboration, and Data Center.
- Cisco Prime for IP Next Generation Networks (NGN): A comprehensive suite of network management applications that helps simplify the design, provisioning, and management of carrier-grade networks. Some of the products in Cisco Prime for IP NGN that are applicable to the NGWAN are these:
  - Cisco Prime Provisioning: Provides service fulfillment across the entire network including the core, aggregation, access, and Internet/enterprise edge
  - Cisco Prime Optical: Provides scalable configuration provisioning and troubleshooting of converged IP and optical transport networks

Individual Cisco Prime products and their application to the NGWAN are discussed in further detail within the documentation provided for each of the NGWAN network modules.

SecureX

Security for enterprise networks, internal and external through the enterprise edge, is a critical requirement for the architecture. The solutions employed include the ability to protect data, create separate logical networks for privacy, and protect the infrastructure itself from attack.

Five security use cases are included with the architecture:

- Cloud (for example, the effect of a public or hybrid cloud at the enterprise edge)
- Infrastructure hardening
- Extranet and partner access
- Segmentation (such as closed user groups for regulatory compliance and mergers and acquisitions)
- Secure connectivity and data privacy, including IP Security (IPsec) encryption models and Cisco TrustSec support

MediaNet

Rich media applications are creating increased demands on the enterprise routing infrastructure. These technologies are being delivered within the enterprise routing architecture as necessary within each tier:

- Medianet (video): For ensuring quality of experience, monitoring, and deployment
  - Resource control
  - Monitoring
  - Auto-configuration
  - Media Services Interface (MSI)
  - Metadata
  - Media trace
- Hierarchical and consistent end-to-end QoS
- Cisco Performance Routing (Cisco PfR): Optimizing routes based on factors such as packet loss, latency, and jitter
- Video optimization: Providing the ability to scale video to branch-office locations
- Cisco Service Advertisement Framework (SAF): Scaling the Cisco Unified Communications environment

TrustSec

With “bring your own device” (BYOD), enterprise networks must now support a variety of devices including personal mobile devices. The TrustSec architecture provides an integrated solution that offers authentication, access control, and user policies to secure network connectivity and resources and authorization of all users on the network through the Cisco Identity Services Engine (ISE). Use cases include:

- Authentication and authorization
- Guest access management and enforcement

Application Visibility and Control

Cisco AVC is a powerful integrated solution for cloud services (public, private, and hybrid) with four primary use cases:

- Application visibility into difficult-to-classify applications obscured by HTTP or virtual desktop infrastructure (VDI) transports
- Application response time (ART) monitoring of application performance
- Media monitoring (MMon) to tracking the performance of voice and video applications
- Application optimization

The Cisco WAAS and WAAS Express technologies are used along with integrated Cisco IOS® Software technologies such as Cisco PfR, Next Generation Network-Based Application Recognition (NBAR2), Flexible NetFlow, quality of service (QoS), and Cisco IOS IP Service Level Agreements

(IP SLAs) to provide a full complement of functions including classification, optimization, monitoring, and troubleshooting.

Cloud Connectors

Cloud computing is a significant driver that is revolutionizing the enterprise and service provider industries. The routed WAN infrastructure provides a key role in these enterprise cloud strategies:

- Inter-cloud routing: The NGWAN metro and WAN cores provide reliable transport services between private clouds, public clouds, and virtual private clouds.
- Cloud security: Encryption for privacy protection between clouds and threat prevention, detection, and mitigation when public and virtual private cloud services are accessed.
- AVC: AVC is used to improve the monitoring and optimization of cloud applications to remote branch offices.

Cloud connectors improve the delivery and performance of cloud services. Some examples include Cisco ScanSafe, Hosted Collaboration Solution (HCS) connector, and WebEx cloud connectors. Find a complete list of available cloud connectors at this link:
http://www.cisco.com/en/US/prod/routers/cloud_connectors.html.

IPv6

The primary business case for IPv6 is being driven by the explosion of mobile devices and the fact that new IPv4 address space is very limited. The NGWAN architecture enables the use of IPv4 and IPv6 simultaneously. Customer scenarios are divided into four categories:

- Those who, through a government mandate or similar requirement, are being compelled to implement IPv6
- Those who want to deploy IPv6 to resolve problems such as IPv4 address exhaustion, globalization needs, mobile device adoption, or IT consumerization
- Thought leaders who are looking for differentiation or competitive advantage: Typical scenarios involve building smart grids, peer-to-peer applications, and mergers and acquisitions
- Those who are looking at the evolution of IPv6 and are concerned about investment protection

The use cases in this architecture are focused on the internal usage of IPv6 within IP and MPLS networks, as well as the interface to external partners and Internet applications. The focus is on supporting IPv6 and IPv4 in a dual-stack approach (preferred), while enabling a non-disruptive migration and transition strategy as more IPv6 hosts and applications become available.

Summary

The Cisco NGWAN provides a modular, hierarchical architecture for deploying routed WAN solutions that provide the services necessary to support the emerging technologies and user environment of the Borderless Network. This approach also allows the customer to adopt services incrementally and smoothly, rather than requiring large-scale redesigns and changes.

The architecture provides a blueprint for building services-led WAN architectures for branch, metro, core, and enterprise edge connectivity for enterprise networks that scale from a single branch network to a large global network that spans multiple continents.

The other documents in this series describe the specific use cases and deployment scenarios that provide the necessary guidance for success in deploying the Cisco NGWAN architecture.

For More Information

Read more about the Cisco NGWAN or contact your local account representative.