

Next Generation Enterprise WAN

Architecture Overview

February, 2012

Purpose

This document provides an Architecture Overview of the Next Generation Enterprise WAN, and takes an architectural approach to designing and deploying routed WAN solutions for Borderless Networks. It provides a breakdown of the architectural components as well as an overview of the Borderless Network Services integrated within the architecture.

This document is intended for the reader with any of the following:

- Large national or global Wide Area Networks
- Regional Branch networks with up to 15,000 sites
- MPLS VPN, Internet, and 3G/4G transport for branch connectivity
- Metro networks using Carrier Ethernet on self-deployed optical networks of up to 100 sites
- Enterprise Edge with collaboration (voice and video) or public cloud requirements
- Core networks within theater (country) and global, based on either 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 and/or countries

Related Documents

This document provides the initial details of the architecture, the component architectures, and the integration of Borderless Network services, as well as IPv6.

Several companion documents, listed as follows, provide more detailed definition of the component architectures, their designs, deployment and configuration guides, 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 Overview
- Regional WAN Deployment Guide
- Video Deployment Guide
- Regional WAN Remote Access VPN Overview
- Regional WAN Remote Access VPN Deployment Guide
- PfR Implementation Guide
- DMVPN-to-GET VPN Migration Guide
- IPv6 Migration Guide
- Network Management Implementation Guide
- QoS Supplemental Guide
- Core Solution Overview
- Metro Network Solution Overview
- Enterprise Edge Solution Overview

Contents

Purpose	2
Related Documents	3
Introduction	5
Architecture Benefits	5
Architecture Overview	5
Architecture Modules	6
Regional WAN Module	7
Metro Module	8
Enterprise Edge Module	9
Enterprise Interconnect Module	11
In-Theater and Global Core Modules	12
Borderless Network Services and IPv6	12
Security	12
Mobility	13
Application Velocity	13
Rich Media	13
IPv6	13
Summary	14

Introduction

The Cisco Next Generation Enterprise WAN (NGEW) is a comprehensive architecture developed for large enterprise and public sector entities, targeting large-scale routed WAN deployments. The architecture encompasses branch and metro connectivity, in-theater and global core backbones, as well as newly emerging enterprise edge functions.

Architecture Benefits


The Cisco Next Generation Enterprise WAN is a comprehensive architectural approach to defining the enterprise routed WAN. You can realize several benefits using this architectural approach:

- Reduced operating costs
- Increased scalability
- Improved security and management of risk
- Improved application performance
- Maximum network availability
- Flexibility and “services on demand”
- Reduced overall complexity
- Faster time to market for new IT services
- Accelerated merger and acquisition ROI
- Regulatory compliance

Architecture Overview

The Next Generation Enterprise WAN (NGEW) Architecture is modular and hierarchical in nature, providing a scalable solution across the enterprise customer segments (Figure 1).

Figure 1. High Level Architecture Abstract



The Architecture is comprised of 5 core modules:

- The Regional WAN: Used to connect branch offices and aggregate remote locations
- The Regional MAN: Used to connect remote offices and Data Centers across metro transports
- The WAN Core: Used to interconnect Regional networks and Data Centers within a country or theater or globally
- The Enterprise Edge: Used to connect the Enterprise network to other external networks and services
- The Enterprise Interconnect: Used as an interconnect and aggregation point for all modules

This modular approach allows for the design of a Next Generation Enterprise WAN utilizing basic building blocks and provides the capability to build a regional WAN, regional MAN, and WAN core and interconnect them as required. Tying these modules together is accomplished using the Enterprise Interconnect, acting as an enterprise distribution network and interconnected by the In-theater level core network. In the largest cases, a top-tier global core backbone can interconnect such theater networks.

The architecture also focuses on the Enterprise Edge, encompassing emerging strategies such as connectivity to cloud (Private, Public, Hybrid) and collaboration services, as well as business-to-business rich-media capabilities such as TelePresence.

In addition to providing 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 incorporation of Medianet functionality, Security services, Application Performance, IPv6 transition, and Mobility functionality within the routing architecture is an important consideration when considering a WAN solution. This approach developed within the architecture allows for incremental additions of services such as these without the need for large scale replacement of equipment or redesign.

Finally, the architecture also looks closely at the linkages to the Cisco Virtualization, Collaboration, and Service Provider Architectures. By including linkages to the Enterprise Virtualization and Collaboration architectures, the Next Generation Enterprise WAN builds on the ability to provide additional functionality and true end-to-end value.


Architecture Modules

The architecture comprises multiple tiers - Regional and Metro Access, In-Theater Core, and Global Core. An additional access tier provides external access through the Enterprise Edge (next generation Internet edge). By creating this hierarchical structure, the architecture can be separated into the elements that are relevant to your environment. When a global footprint is required, all aspects of the architecture will likely apply, whereas a footprint that is solely within a single theater will not require the Global Core, but can be added when there is a requirement to expand into other theaters.

The architecture also introduces a new concept of the Enterprise Interconnect (Figure 2). The Enterprise Interconnect is a part of any location that brings together a combination of the Regional WAN and Metro networks, Enterprise Edge, and Core for the purpose of integrating them into a single architecture. Also within the interconnect are optional elements for the local data center and campus that may be co-resident at the location.

Figure 2 depicts the architecture and the individual components.

Figure 2. Hierarchical Architecture Overview



To fully understand the architecture and the ability to provide Borderless Network services, the following sections will provide more details about what each module encompasses.

Regional WAN Module

The Regional WAN is the combination of branch access and WAN aggregation routing, along with the transport interconnecting the branch to the regional aggregation router. It also includes the services resident within the branch office and aggregation. Branch locations in this context are remote locations that are provided connectivity to the central campus or data center locations:

- **Standard branch:** A standard branch requires application availability and performance, but 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 will have requirements for collaboration and video services.
- **High-end branch:** A high-end branch is defined as a branch where the customer demands the latest technology for collaboration and application availability, regardless of size. This branch model also encompasses more services in support of critical applications such as telepresence, as well as IP video surveillance, digital signage, etc. Given the criticality of this branch in terms of the applications and performance requirements, providing high availability and application performance capabilities is crucial.


- **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 with very high bandwidth and encryption requirements, usually at speeds up to OC-12 or line-rate Gigabit Ethernet. Because of the performance and availability requirements, the ultra-high-end branch 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. Mobility of the entire branch is the primary requirement, but application performance is also important over the high-latency links of cellular 3G, 4G, or satellite.

Metro Module

The Metro networks are divided into two separate “build-vs.-buy” models. A typical deployment generally uses an Ethernet hand-off from the service provider. Self-deployments take advantage of privately owned or leased fiber assets to deliver a richer set of features as part of an enterprise-owned service offering.


- The Metro Access model provides connectivity to remote offices over Layer 2 Metro Ethernet (Metro-E) service offerings. The Metro-E service offering could be provided by a service provider or by the enterprise-owned metro network. The Metro Access branch and aggregation design is very similar to the Regional WAN design except that Metro-E services are typically layer 2 and offer higher speeds. Figure 3 shows these connectivity models.

Figure 3. Metro Network Access Models



- The standard architectural model uses carrier-supported metro ethernet that uses a service such as Virtual Private LAN Service (VPLS), and requires just that the enterprise create access to the carrier. The typical deployment model employs point-to-point configurations, with redundancy to a second data center or Enterprise Interconnect in an active/standby model. Service-level agreements (SLAs) from the provider typically focus solely on uptime, with best-effort delivery within the carrier core. An alternative configuration consists of a multipoint service, where the customer controls routing and the provider SLAs are strictly on “access” information rates into the Ethernet cloud. This model offers active/active connections but also normally rate-limits multicast and broadcast transfers. Figure 4 depicts the standard offering.

Figure 4. Standard Metro Network Architecture




- The self-deployed model is more robust in redundancy because the remote site has dual routers and more robust optical redundancy and resiliency features. The Metro Access is derived from the standard architecture as defined in the previous paragraph, but the reliance is no longer on a carrier transport; it is, rather, one that the customer deploys.
- The Enterprise MPLS/Metro-E deployment model includes an enterprise-owned and -deployed 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 and time-division multiplexing (TDM) circuit services.
- The Enterprise MPLS + Optical model combines both a self-deployed optical and MPLS metro network for the customers who need the maximum availability, capacity, virtualization, and transport flexibility.

Enterprise Edge Module

The changing requirements for IP connectivity outside of the enterprise require enterprises to consider the 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, introduce new security concerns, and have more instances of service provider connectivity. This situation has created a need to evolve the Internet edge and define 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. Users are employees, partners, customers, guests, etc. Resources are Internet access, business-to-business connectivity, collaboration resources, hosted services, or hosted applications. The Enterprise Edge is also a place where services such as security, collaboration acceleration, etc. reside (Figure 5).

Figure 5. Enterprise Edge




The Enterprise Edge is divided into five sub-modules:

- Web Services
- Mobility Services
- Communications Services
- Cloud Services
- Business-to-Business Services

These sub-modules are optionally installed at each location where the Enterprise Edge connectivity is required - central sites, regional hubs, or branch offices (Figure 6).

Figure 6. Enterprise Edge Component View




Enterprise Interconnect Module

The Enterprise Interconnect is the location where the Core, Regional WAN, Metro, and Enterprise Edge modules are interconnected. Within this location there could also be instances of a local data center or campus.

The Interconnect may consist of any or all of the architecture modules, depending upon the requirements and overall enterprise architecture (Figure 7).

Figure 7. Enterprise Interconnect Framework



The aggregation for the Regional WAN and Metro networks provides the connectivity for remote sites into the Interconnect. The WAN Core modules provide the interconnectivity between Interconnect locations within a theater and global connectivity between Interconnects across multiple theaters.

The Enterprise Interconnect infrastructure is responsible for delivering services that integrate the other component parts of the architecture, such as Inter-Virtual Route Forwarding (VRF) routing, as well as linking the end-to-end integration of Borderless Network services such as medianet. The Interconnect supports services that are required of the NGEW architecture by any co-located data center or campus architectures, but those component architectures (campus and data center) are well-defined elsewhere and are outside the scope of this document.

In-Theater and Global Core Modules

The Enterprise Interconnect provides the connectivity between the regional WANs, MANs, data centers, Enterprise Edge, and campus networks. The WAN Core networks provide connectivity between regional Enterprise Interconnects within a theater as well as globally between theaters. The WAN Core networks are sometimes referred to as the WAN backbone. Figure 2 earlier in this document highlights the positioning of the global and in-theater cores within the architecture.

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, allowing 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, while balancing 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, you can deploy multiple planes (multiplanar cores) to enable availability of up to 6 nine's for both in-theater and global cores.

Borderless Network Services and IPv6

The architecture has four core Borderless Network services, along with IPv6: Security, Application Velocity, Rich Media, and Mobility. The following sections describe the use cases and features that are included within the enterprise routing architecture.

Security

Security for enterprise networks, both internal as well as 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 (closed user groups for regulatory compliance, mergers and acquisitions, etc.)
- Secure connectivity and data privacy, including IP Security (IPsec) encryption models and Cisco TrustSec[®] support

Mobility

Mobility in the context of the Next Generation Enterprise WAN includes the use case of device mobility, especially considering the effect on location services and Enterprise Edge requirements. Also included is the capability of using 3G or 4G cellular and satellite connectivity for mobile branch offices within the Regional WAN. In addition, services are provided for the mobility of virtual machines between data centers (VM Mobility) across both the Metro and Core architectures.

Application Velocity

Application velocity has four primary use cases:

- Virtual desktop infrastructure (VDI)
- Data center-to-data center (for example, Cisco Data Center Interconnect [DCI] transport and VM mobility)
- Application optimization
- Cloud services (public, private, and hybrid)

The Cisco Wide Area Application Services (WAAS) and WAAS Express technologies are used along with integrated Cisco IOS® Software technologies such as Performance Routing (PfR), Network-Based Application Recognition (NBAR), quality of service (QoS), and IP SLA, providing a full complement of functions including classification, optimization, monitoring, and troubleshooting.

Rich Media

Rich Media applications are creating increased demands on the enterprise routing infrastructure. The following 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
- QoS
- 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 Unified Communications environment

IPv6

The primary business use case for IPv6 is to enable the transition to IPv6 while also supporting IPv4. 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 looking for differentiation or competitive advantage: Typical scenarios involve building smart grids, peer-to-peer applications, and mergers and acquisitions
 - Those looking at the evolution to 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 peering interface to external partners and Internet applications. The focus will be on supporting both 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 Next Generation Enterprise WAN 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 necessitating 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 spanning multiple theaters and continents.

The remaining documents in this series outline the specific use cases and deployment scenarios that will provide the necessary guidance for success in deploying the Next Generation Enterprise WAN architecture.



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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)