Deployment Guide

# Next Generation Enterprise WAN

## IPv6 Migration Deployment Guide

September, 2011

NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE 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 NABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0809R).

Next Generation Enterprise WAN Video Deployment Guide.

# Contents

## Document Conventions

Command descriptions use these conventions:

| | |
|---|---|
| **boldface font** | Commands and keywords are in boldface. |
| *italic font* | Arguments for which you supply values are in italics. |
| [ ] | Elements in square brackets are optional. |
| [ x \| y \| z ] | Optional alternative keywords are grouped in brackets and separated by vertical bars. |

Screen examples use these conventions:

| | |
|---|---|
| screen font | Terminal sessions and information in the displays are in screen font. |
| **boldface screen font** | Information you must enter is in boldface screen font. |
| *italic screen font* | Nonprinting characters, such as passwords, are in angle brackets. |
| < > | Default responses to system prompts are in square brackets. |
| !, # | An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line. |

## NGEW Architecture Overview

**Figure 1.** Enterprise Network Architecture

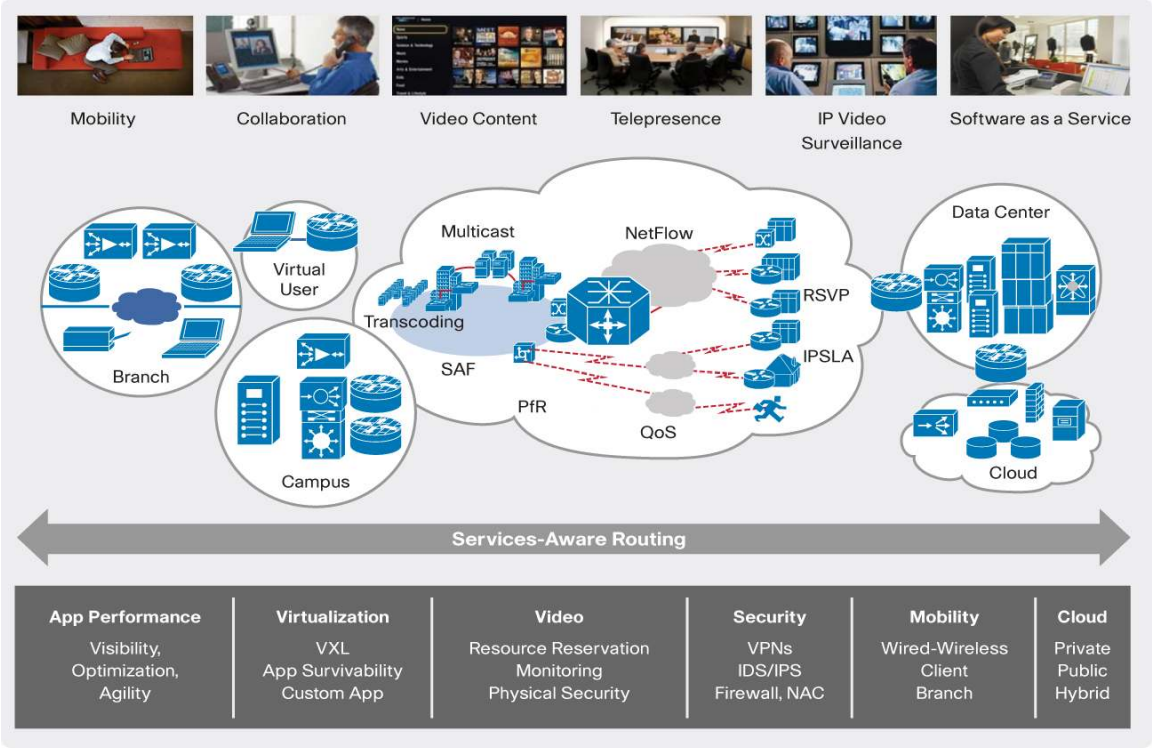


Enterprise networks must adapt to meet new and evolving business requirements. The introduction of cloud services (private, public, or hybrid) poses new challenges to current enterprise network designs. With a more distributed workforce, the proliferation of bandwidth-intensive video-enabled endpoints and the consolidation of servers into a few centralized locations require networks to carry more traffic with increased efficiencies while demanding the same or higher levels of performance and availability.

The Cisco® Next Generation Enterprise WAN (NGEW) is a Cisco end-to-end architecture that provides foundation building blocks for next-generation enterprise networks. The hierarchical design provides the scalability required by large enterprises and can be extended and replicated throughout multiple regions and theaters. This consistency leads to ease of deployment, maintenance, and troubleshooting (Figures 1 and 2).

**Figure 2.** NGEW Regional WAN Topology

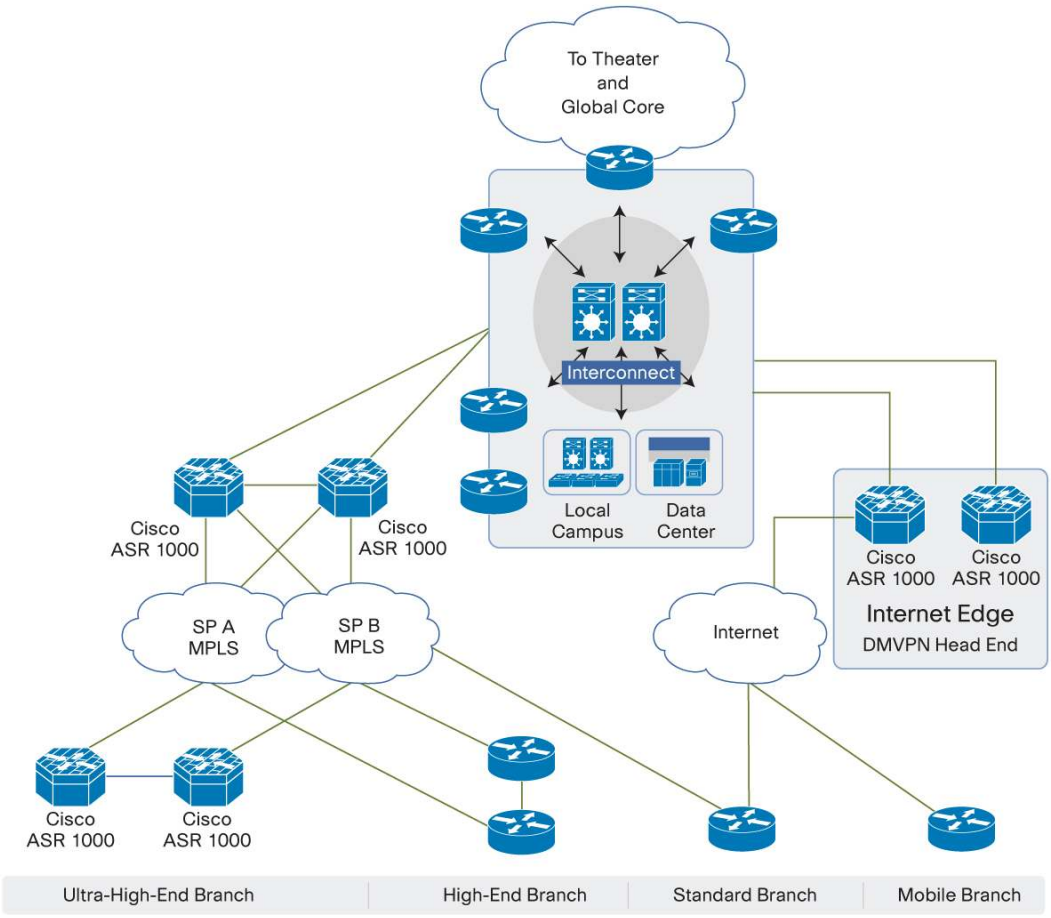


A logical starting point is the regional WAN (RWAN), where all branch-office locations connect to the highly scalable aggregation routers at the enterprise interconnect through various access technologies such as wireless (third and fourth generation [3G and 4G, respectively]), DSL, and Multiprotocol Label Switching (MPLS). The enterprise interconnect is the location where traffic from the RWAN is aggregated to the in-theater and global WAN cores. In addition, the enterprise interconnect links all the other components of the NGEW, including local data centers, the campus, and the enterprise edge, where the enterprise edge is the demarcation point between the enterprise network and any external network service such as Internet, cloud, and voice.

Four branch-office designs are defined within the NGEW RWAN module:

- Mobile branch office: Single tier, single WAN link, and mobile with minimal redundancy
- Standard branch office: Single tier with dual WAN links providing redundancy for link failures
- High-end branch office: Dual tier with dual WAN links providing maximum redundancy for both device and link failures
- Ultra-high-end branch office: Based on the high-end branch office with increased capacity and higher availability

In addition to providing advanced routing functions, one of the primary design goals of the NGEW is to build a network foundation that can reliably support new applications and services, including those in the Cisco Borderless Network - application velocity, medianet, IPv6, and mobility. Customers will benefit from investing in a Cisco network design that has gone through rigorous testing and that addresses there continuously evolving business requirements by supporting new applications and services.

## IPv6 Business Overview

**Figure 3.** Business Drivers

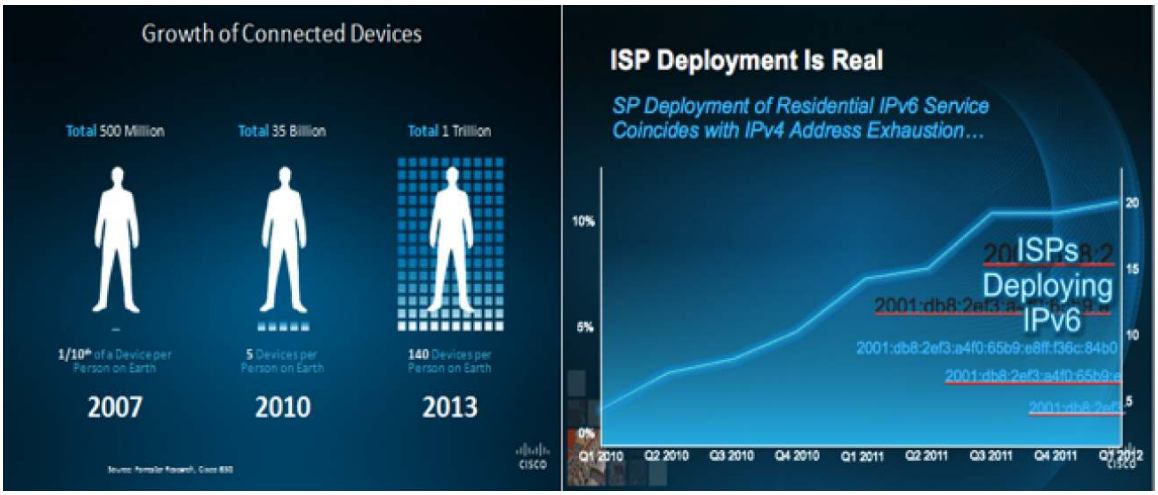


IPv6 is the next version of Internet Protocol that is designed to replace IPv4. According to the Internet Assigned Numbers Authority (IANA), the last /8 block of free IPv4 addresses was assigned in February 2011. By the end of 2011, there may be no more free IPv4 addresses. Although the anticipated IPv4 address exhaustion initially accelerated adoption of IPv6, IPv6 has other benefits, including:

- Simplified headers for routing efficiency
- Deeper hierarchy and policies for network architecture flexibility and efficient support for routing and route aggregation
- Server-less autoconfiguration, easier renumbering, and improved ready-to-use support
- Security with mandatory IP Security (IPsec) implementation for all IPv6 devices
- Improved support for Mobile IP and mobile computing devices (direct path)
- Enhanced multicast support with increased addresses and more efficient mechanisms

Cisco NGEW helps customer's transition to IPv6 by following three general approaches:

- Preserve: Continue to use existing IPv4 investments while determining exactly how and when to fully adopt IPv6 standards
- Prepare: Position the organization for growth with architectural designs, IPv6-ready technology solutions, and certifications
- Prosper: Expand innovation and productivity by smoothly transitioning to the global IPv6 customer base

## IPv6 Deployment Models Overview

Enterprises need to adopt IPv6 to meet new and evolving challenges of globalization. Globalization has necessitated the need to communicate with customers and branch offices in regions that had only IPv6 accessibility. Employee mobility also necessitates IPv6 to be enabled on the enterprise. Cisco Next Generation Enterprise WAN enabled IPv6 for the enterprise by enabling dual stack on the enterprise without affecting the experience and performance of users on existing IPv4 networks.

In the current NGEW phase, IPv6 deployment is limited to two major use cases - establishing an Internet presence in IPv6 and supporting IPv6 endpoints by connecting IPv6 islands using 6to4 tunnels.

### Internet Presence

In order to demonstrate the commitment to IPv6, having an IPv6 Internet presence is crucial. It is important to provide web content to IPv6 users while at the same time preserving the end-user experience. This paradigm allows the business to reach new subscribers and new geographies where IPv6 users can be found.

In order to provide an Internet presence, IPv6 Internet gateways are installed at the enterprise edge. IPv6 servers are accessible from the Internet either directly or through 6to4 tunnels to other IPv6 islands or hosts.

**Figure 4.** Enterprise Edge

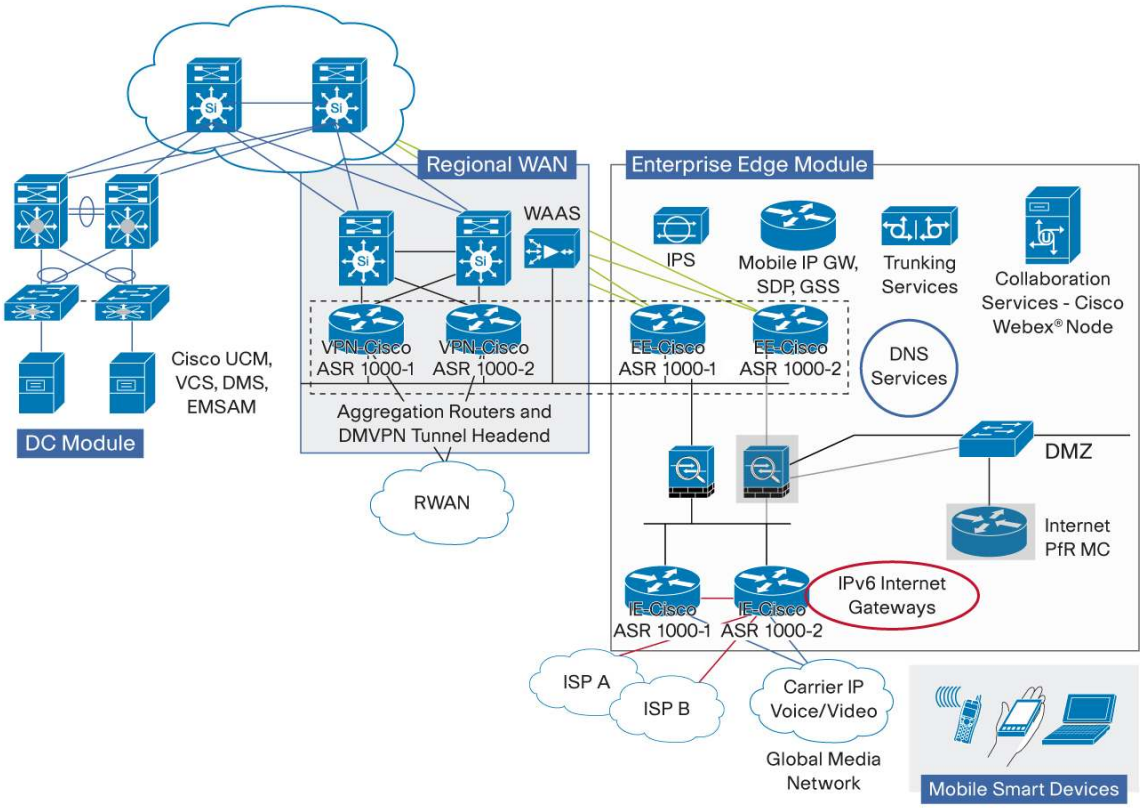

## IPv6 Hosts and Islands and Dual Stack

With the growing adoption of Windows 7 and other mobile devices such as Smartphone's, tablets, etc., it has become necessary to introduce a dual-stack network both in the enterprise core and at branch offices. Dual stack allows the network to support both IPv4 and IPv6 endpoints. IPv6-only clients at the branch office can access IPv4 content on the Internet using the Network Address Translation 64 (NAT64) services on the enterprise core. Details about NAT64 are available at ASR1000-NAT64-Deployment-Guide.pdf.

However, because most Internet service provider (ISP) connections (MPLS and Internet) today are IPv4-only, it is necessary to use 6to4 tunnels to interconnect the RWAN to other NGEW components and between RWAN branch offices.

**Figure 5.** NGEW IPv6 Regional WAN Topology

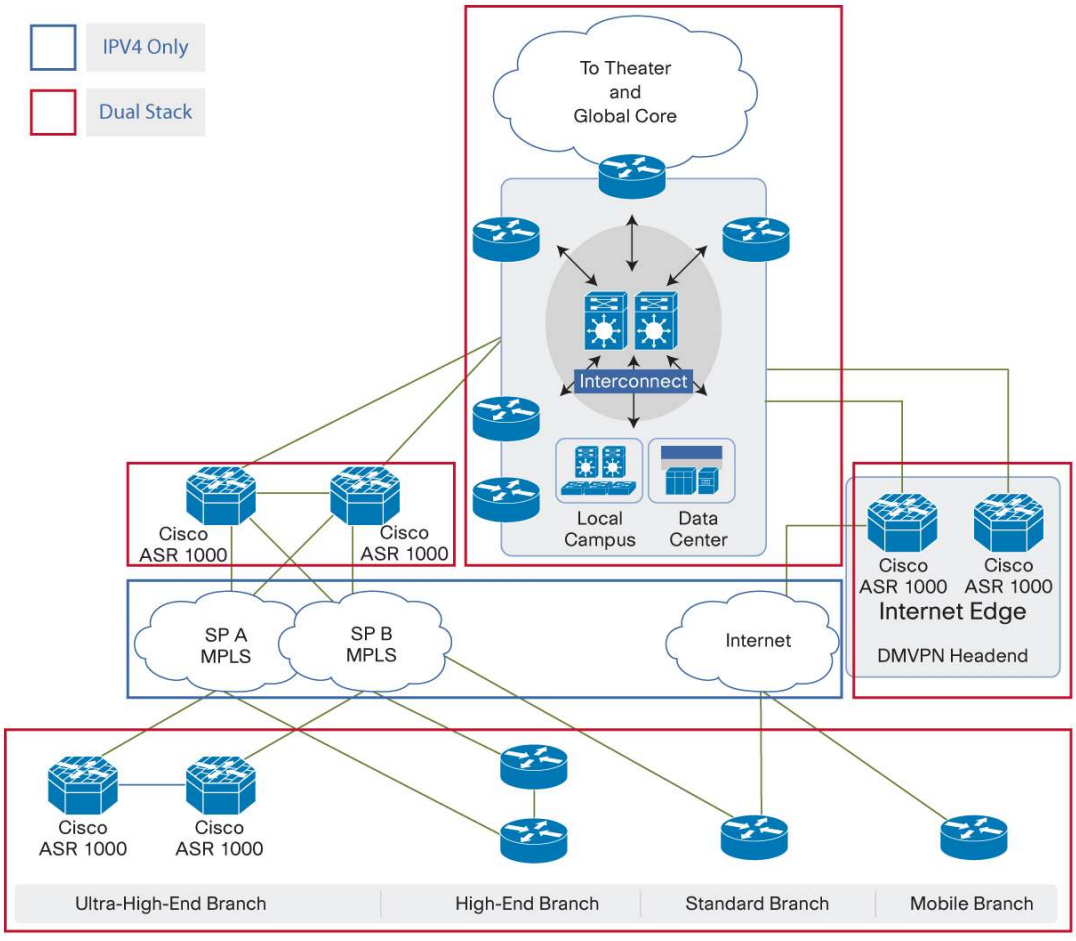

## Implementing IPv6 in NGEW

In order to support IPv6 endpoints in an enterprise branch office, a dual-stack network (IPv4 and IPv6) is enabled on the LAN while the WAN remains IPv4-only. An automatic 6to4 tunnel is then configured on the branch-office router, thus allowing users to access IPv6 servers in other parts of the enterprise network.

### Automatic 6to4 Tunnels

An automatic 6to4 tunnel allows connection of isolated IPv6 domains over an IPv4 network to remote IPv6 networks. The important difference between automatic 6to4 tunnels and manually configured tunnels is that the tunnel is not point-to-point - it is point-to-multipoint. In automatic 6to4 tunnels, routers are not configured in pairs because they treat the IPv4 infrastructure as a virtual nonbroadcast multi-access (NBMA) link. The IPv4 address embedded in the IPv6 address is used to find the other end of the automatic tunnel.

An automatic 6to4 tunnel may be configured on a border router in an isolated IPv6 network, creating a tunnel on a per-packet basis to a border router in another IPv6 network over an IPv4 infrastructure. The tunnel destination is determined by the IPv4 address of the border router extracted from the IPv6 address that starts with the prefix 2002/16, where the format is 2002:border-router-IPv4-address::/48. Following the embedded IPv4 address are 16 bits that can be used to number networks within the site.

The critical requirement for automatic 6to4 is that each site needs to have a globally unique IPv4 address. In NGEW RWAN branch-office design, a loopback address is used, enabling all four branch-office designs to support IPv6 endpoints with a similar implementation.

Use the following steps to configure automatic 6to4 tunnels on a branch-office router:

**Step 1.** Configure an IPv4 loopback interface. Be sure to advertise the interface in the IPv4 routing domain.

```
interface Loopback0
 ip address 192.168.0.10 255.255.255.255
```

**Step 2.** Configure an automatic 6to4 tunnel using the previously defined IPv4 loopback interface as the tunnel source. Note that a tunnel destination is not defined.

```
interface Tunnel200
 no ip address
 ipv6 address 2002:A03:1::1/16
 tunnel source Tunnel10
 tunnel mode ipv6ip 6to4
```

Figure 6 shows an example of configuring 6to4 tunnels.

**Figure 6.** Configuring 6to4 Tunnels

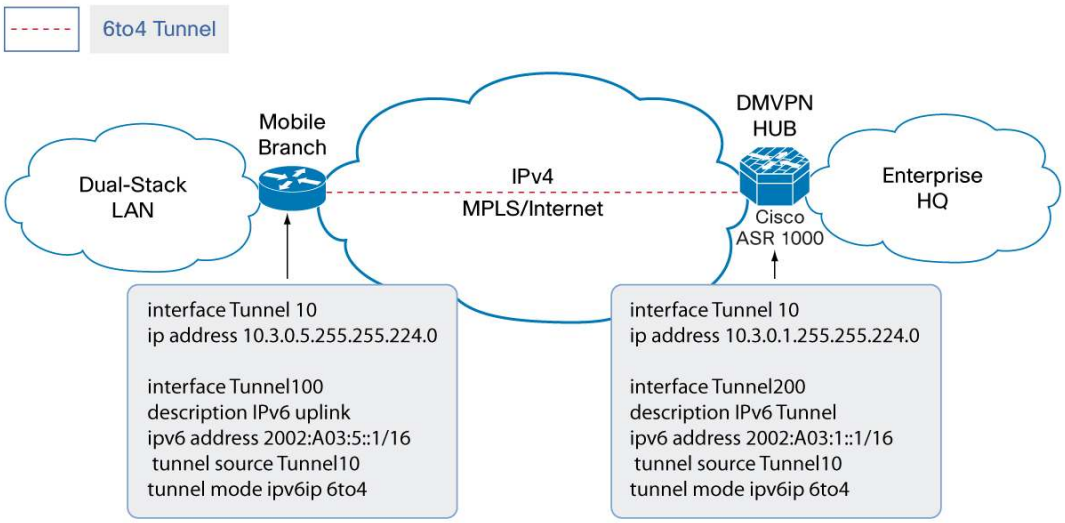


## IPv6 Routing

Both IPv4 and IPv6 can be enabled on the branch-office LAN network. Because multicast is not supported over 6to4 tunnels, Border Gateway Protocol (BGP) is the recommended IPv6 routing protocol. All branch-office routers should configure the VPN aggregation routers (Group Encrypted Transport VPN or Dynamic Multipoint VPN [DMVPN]) as their BGP neighbors, whereas the aggregation routers need to configure all branch-office routers as their BGP neighbors. This restriction is that branch office-to-branch office IPv6 traffic will have to travel through the aggregation point.

Use the following steps to configure IPv6 LAN and routing:

**Step 1.** Configure an IPv6 address on the LAN interface.

```
interface GigabitEthernet0/0.1
 ipv6 address 2001:DB8:1111::1/64
 ipv6 enable
```

**Step 2.** Configure BGP for IPv6.

```
router bgp 65010
 address-family ipv6
  network 2001:DB8:1111::/64
  neighbor 2002:C0A8:0001::1 remote-as 65001
  neighbor 2002:C0A8:0001::1 activate (Remote Tunnel address)
 exit-address-family
```

### IPv6 Host Address Allocation

You can configure a host address in IPv6 in four ways:

- Static configuration: The IPv6 address, mask, and gateway address are manually defined on the host.

This method is not recommended because it is not user-friendly.

```
interface GigabitEthernet 0/0
 ipv6 address 2002:A04:510A::1/64
 ipv6 enable
```

- Stateless address auto-configuration (SLAAC): With SLAAC, router solicitation messages are sent by booting nodes to request Router Advertisements (RAs). The host then autonomously configures its own address. A node on the link can automatically configure global IPv6 addresses by appending its interface identifier (64 bits) to the prefixes (64 bits) included in the RA messages. The resulting 128-bit IPv6 addresses configured by the node are then subjected to duplicate address detection to ensure their uniqueness on the link. If the prefixes advertised in the RA messages are globally unique, then the IPv6 addresses configured by the node are also guaranteed to be globally unique.

```
<no router side configuration is required>
```

- Stateful Dynamic Host Configuration Protocol Version 6 (DHCPv6): The host uses DHCP to get its IPv6 address. This behavior is similar to IPv4 behavior. However, DHCPv6 uses multicast for many of its messages. Initially, the client must first detect the presence of routers on the link using neighbor discovery messages. If a router is found, then the client examines the RA to determine if DHCP should be used. If the Managed flag is enabled in RA messages, the client then starts a DHCP solicitation phase to find a DHCP server.

This method is recommended for use in NGEW because it allows the most control by the administrator.

```
ipv6 dhcp pool DATA
 address prefix 2001:DB8:1111::0/64 lifetime infinite
 dns-server 2001:DB8:A:B::1
 domain-name cisco.com

interface GigabitEthernet 0/0.1
 ipv6 address 2001:DB8:1111::1/64
 ipv6 enable
 ipv6 address dhcp server DATA rapid-commit
```

- Stateless DHCPv6: Stateless DHCPv6 combines SLAAC for address assignment with DHCPv6 for all other configuration settings, such as a TFTP server, a Domain Name System (DNS) server, etc. A host builds its address by appending a host identifier to the /64 prefix received from the router and issues a DHCP solicit message to the DHCP server.

```
ipv6 dhcp pool DATA
 dns-server 2001:DB8:A:B::1
 domain-name cisco.com

interface GigabitEthernet 0/0.1
 ipv6 address 2001:DB8:1111::1/64
 ipv6 enable
 ipv6 dhcp server DATA
```

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.**



Printed in USA C07-685954-00 10/11