IP6FD

IPv6 Fundamentals, Design, and Deployment

Volume 3

Version 3.0

Student Guide

Text Part Number: 97-2947-01

https://t.me/learningnets

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 of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at 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. (1005R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS." CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

© 2010 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Table of Contents

Volume 3

Deploying IPv6 ***8-1***

Overview 8-1
Module Objectives 8-1

Examining IPv6 Address Allocation **8-3**

Overview 8-3
Objectives 8-3
IPv6 Internet 8-4
IPv6 Address Allocation 8-6
Connecting to the IPv6 Internet 8-15
Summary 8-23
Resources 8-23

Understanding the IPv6 Multihoming Issue **8-25**

Overview 8-25
Objectives 8-25
IPv6 Multihoming Aspects and Issues 8-26
Provider-Independent vs. Provider-Allocated 8-28
PA IPv6 Addressing 8-30
PI IPv6 Addressing 8-32
PA versus PI 8-33
Internet Routing Table Growth 8-34
Multihoming Requirements 8-35
Multihoming and IPv6 8-37
Site Multihoming Issues 8-38
IPv6 Multihoming Status 8-39
Routing Example 8-43
Protocol-Based Solutions 8-46
SHIM6 8-46
LISP 8-47
Summary 8-49
Resources 8-49

Identifying IPv6 Enterprise Deployment Strategies **8-51**

Overview 8-51
Objectives 8-51
Enterprise Networks 8-52
Security of IPv6 Tunnels 8-56
Impacts of Network Services 8-62
WAN Networks 8-68
Dual Stack: Advantages and Disadvantages 8-69
Tunneling: Advantages and Disadvantages 8-70
Translation: Advantages and Disadvantages 8-71
Application Layer Gateways 8-72
Summary 8-73
Module Summary 8-75
Module Self-Check 8-77
Module Self-Check Answer Key 8-79

https://t.me/learningnets

IPv6 and Service Providers 9-1

Overview 9-1
Module Objectives 9-1

Identifying IPv6 Service Provider Deployment 9-3

Overview 9-3
Objectives 9-3
IPv6 Service Provider Deployment 9-4
Upgrade the Entire Network to Dual Stack 9-4
Deploy a New Parallel, IPv6-Only Network 9-5
Upgrade the Network Edge 9-6
Dual-Stack Deployment 9-7
IPv6-Only Deployment 9-9
Encapsulation 9-11
IPv6 Services 9-13
Key Service Provider Strategies 9-15
Service Layer Address Allocation 9-16
Encapsulation Support 9-24
Selection Guidelines 9-25
Summary 9-27

Understanding Support for IPv6 in MPLS 9-29

Overview 9-29
Objectives 9-29
MPLS Operations 9-30
IPv6 over MPLS Deployment Scenarios 9-36
IPv6 Tunnels Configured on CE Routers 9-37
Pros 9-38
Cons 9-38
IPv6 over Layer 2 MPLS VPN 9-39
Pros 9-40
Cons 9-40
Description of Pros and Cons 9-40
Cisco 6PE 9-41
Multiprotocol Extensions for BGP4 9-42
Life of a Packet (Control Plane) 9-44
Life of a Packet (Data Plane) 9-47
How to Deploy Cisco 6PE on MPLS Networks 9-54
Summary 9-56

Understanding 6VPE 9-57

Overview 9-57
Objectives 9-57
Cisco 6VPE 9-58
Standardization 9-59
Configuring 6VPE 9-61
Summary 9-67
References 9-67

Understanding IPv6 Broadband Access Services 9-69

Overview 9-69
Objectives 9-69
IPv6 Rapid Deployment 9-70
Customer Link Encapsulations 9-75
FTTH Access Architecture 9-76
Cable Access Architecture 9-77
Wireless Access Architecture 9-78
DSL Access Architecture 9-79
Summary 9-80
References 9-80
Module Summary 9-81

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check 9-83
Module Self-Check Answer Key 9-85

IPv6 Case Studies ***10-1***

Overview 10-1
Module Objectives 10-1

Planning and Implementing IPv6 in Enterprise Networks **10-3**

Overview 10-3
Objectives 10-3
Enterprise Network Definition 10-4
Implementing IPv6 in an Enterprise Campus Network 10-9
IPv6 in an Enterprise WAN Network 10-17
Summary 10-27
Resources 10-27

Planning and Implementing IPv6 in Service Provider Networks **10-29**

Overview 10-29
Objectives 10-29
Service Provider Network Design 10-30
Native IPv6 Deployment in Service Provider Access Networks 10-35
Native IPv6 Deployment in the Service Provider Core Network 10-38
6PE Deployment in the Service Provider Core Network 10-42
Summary 10-45

Planning and Implementing IPv6 in Branch Networks **10-47**

Overview 10-47
Objectives 10-47
Branch Deployment Overview 10-48
Branch Deployment Profiles: Single-Tier Profile Implementation 10-58
Branch Deployment Profiles: Dual-Tier and Multitier Profile Implementations 10-62
Summary 10-67
Resources 10-67
Module Summary 10-69

https://t.me/learningnets

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Module 8

Deploying IPv6

Overview

IP version 6 (IPv6), like IP version 4 (IPv4), is provisioned through standards bodies that define its operations, structure, and deployment. Understanding how IPv6 addresses are allocated, how organizations connect to the Internet, and how different network structures require different approaches to IPv6 integration will ensure that your deployment is consistent with industry standards and best practices. This module describes the standards bodies that define IPv6 address allocation, in addition to one of the leading IPv6 deployment issues, which is multihoming. The module also describes IPv6 deployment practices for enterprise and service provider networks.

Module Objectives

Upon completing this module, you will be able to describe the standards bodies that define IPv6 address allocation, in addition to one of the leading IPv6 deployment issues—multihoming. This ability includes being able to meet these objectives:

- Describe IPv6 address allocation in the IPv6 Internet
- Understand multihoming issues and potential solutions
- Identify IPv6 enterprise integration strategies

https://t.me/learningnets

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1

Examining IPv6 Address Allocation

Overview

The requirements for obtaining and allocating IP version 6 (IPv6) addresses are similar to IP version 4 (IPv4), with some specific differences. Allocation methods differ significantly from IPv4 allocation methods because of the extensive size of the IPv6 address space. This lesson describes the IPv6 networking environments in use today, the process of becoming an IPv6 Internet service provider (ISPs), address allocation policies and organizations, and strategies for connecting to the IPv6 Internet.

Objectives

Upon completing this lesson, you will be able to explain IPv6 address allocation in the IPv6 Internet. This ability includes being able to meet these objectives:

- Describe the IPv6 Internet
- Explain how IPv6 addresses are allocated
- Describe how to connect a network to the IPv6 Internet

https://t.me/learningnets

IPv6 Internet

This topic describes the 6bone network.

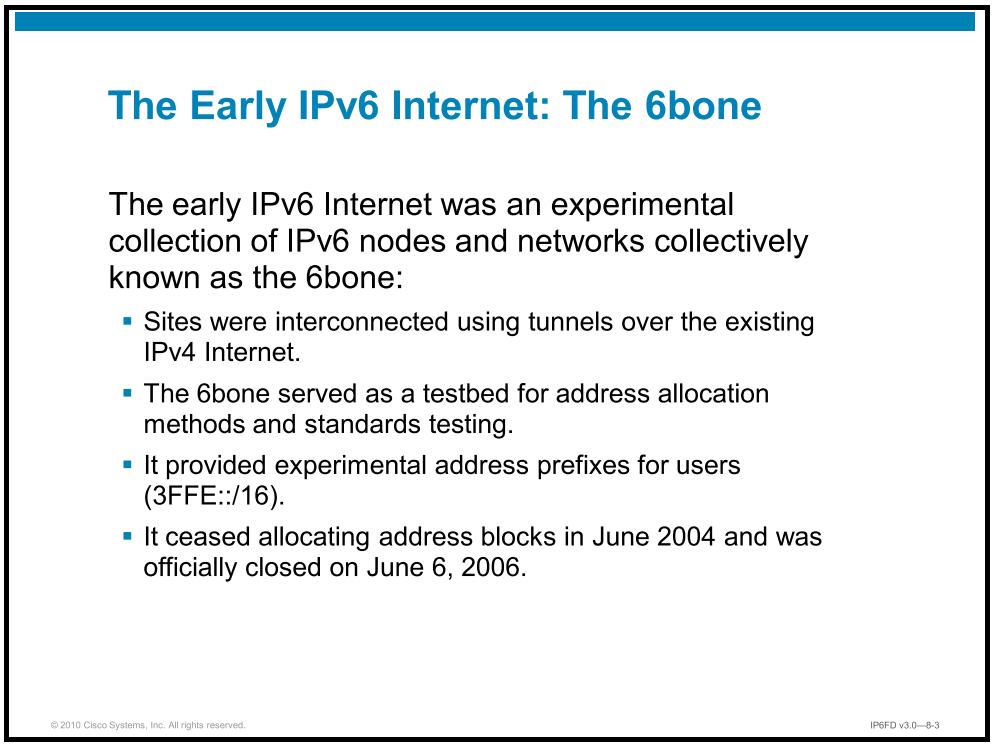

In the early days of IPv6, little opportunity for early adopters to communicate over the Internet existed. The Internet is, basically, a collection of interconnected networks and nodes that share a common communications language. With few IPv6-capable nodes and no well-structured system in place for addressing, there was no general capability for an IPv6 Internet. To overcome this issue, the Internet Engineering Task Force (IETF) created the 6bone.

The 6bone was started as a tunneled network over the IPv4 Internet. Past experience with similar tunneled networks, such as the multicast backbone, suggests that this setup was temporary until all links were converted to native links.

The 6bone was a testbed for trying new protocols, implementations, transition mechanisms, and operational procedures, to give feedback to developers and protocol designers. It was also the primary distributor of IPv6 addresses for early experimenters. The 6bone used the 3FFE::/16 address pool to provide address allocations, with the understanding that these were temporary assignments.

By design, the 6bone was structured to gracefully terminate operations and close. In June of 2004, the 6bone ceased distribution of 3FFE:: addresses and instead encouraged requestors to seek production addresses. On June 6, 2006, the 6bone terminated operations.

Note You can find historical information about the 6bone at http://www.6bone.net. The process for terminating 6bone operations is described in IETF RFC 3701, *6bone (IPv6 Testing Address Allocation) Phaseout.*

https://t.me/learningnets

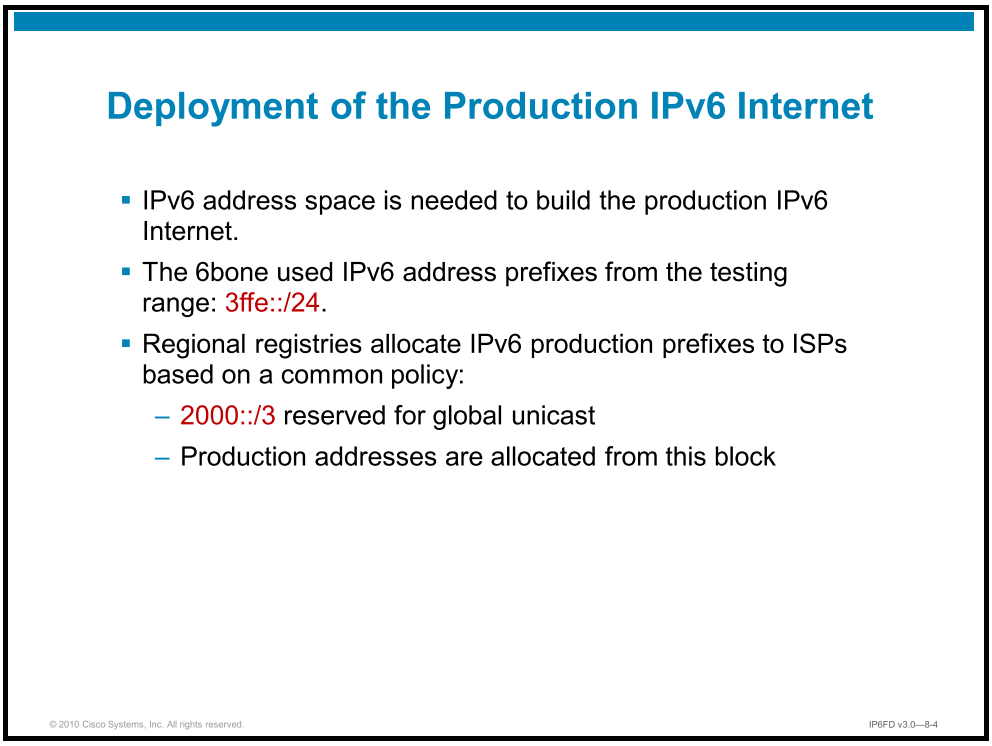

The 6bone was an experimental network, built for interoperability testing in a global context. The production IPv6 network will be very long-lived and will require much addressing space to accommodate the large number of devices.

The 6bone used the **3ffe::/24** prefix to allocate IPv6 addresses.

Within each geographical region, Regional Internet Registries (RIRs) assign addresses to ISPs. A common policy is used to make these assignments, so that each central authority treats each region evenly.

IANA assigns RIRs **/12** address blocks from the **2000::/3** range.

https://t.me/learningnets

IPv6 Address Allocation

This topic describes the regional address registries that are used for IPv6.

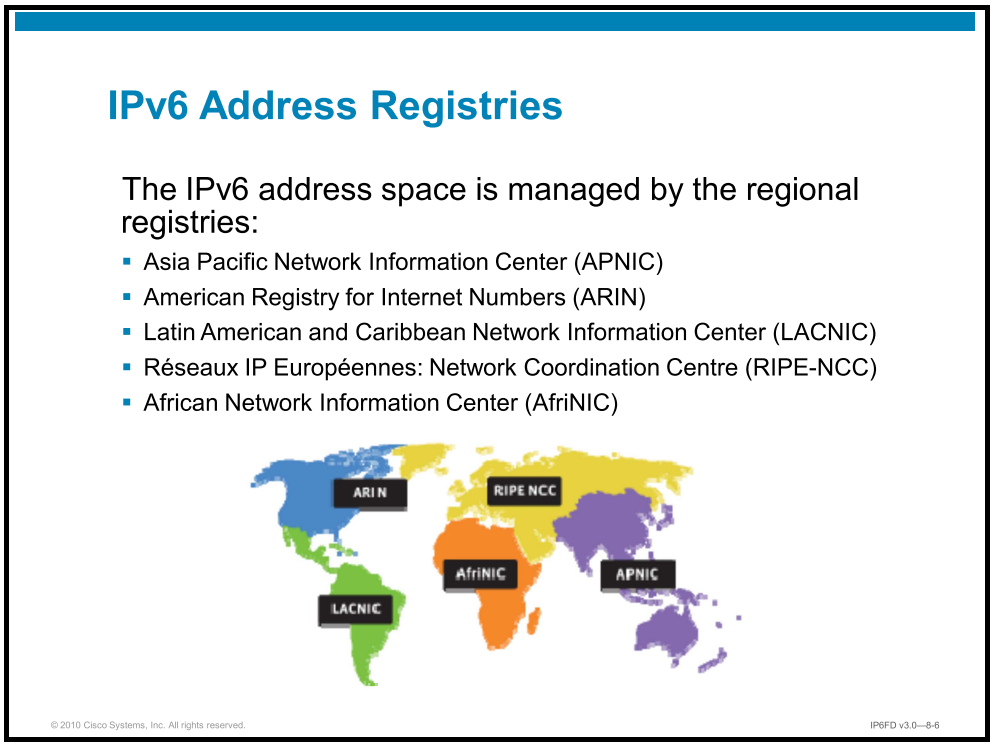

The five registries[1] are located in different regions of the world: Asia, North America, South America, Europe, and Africa, as follows:

- **Asia Pacific Network Information Center (APNIC):** Asia and the Pacific
- **American Registry for Internet Numbers (ARIN):** North America
- **Latin American and Caribbean Network Information Center (LACNIC):** Latin America (Central and South America) and the Caribbean
- **Réseaux IP Européennes: Network Coordination Centre (RIPE-NCC):** Europe and the Middle East
- **African Network Information Center (AfriNIC)**: Africa

In both IPv4 and IPv6, ISPs receive their address space from those registries, which generally charge for the service.

Providers should apply to the registry that is in their region.

[1] Image Copyright © IANA

https://t.me/learningnets

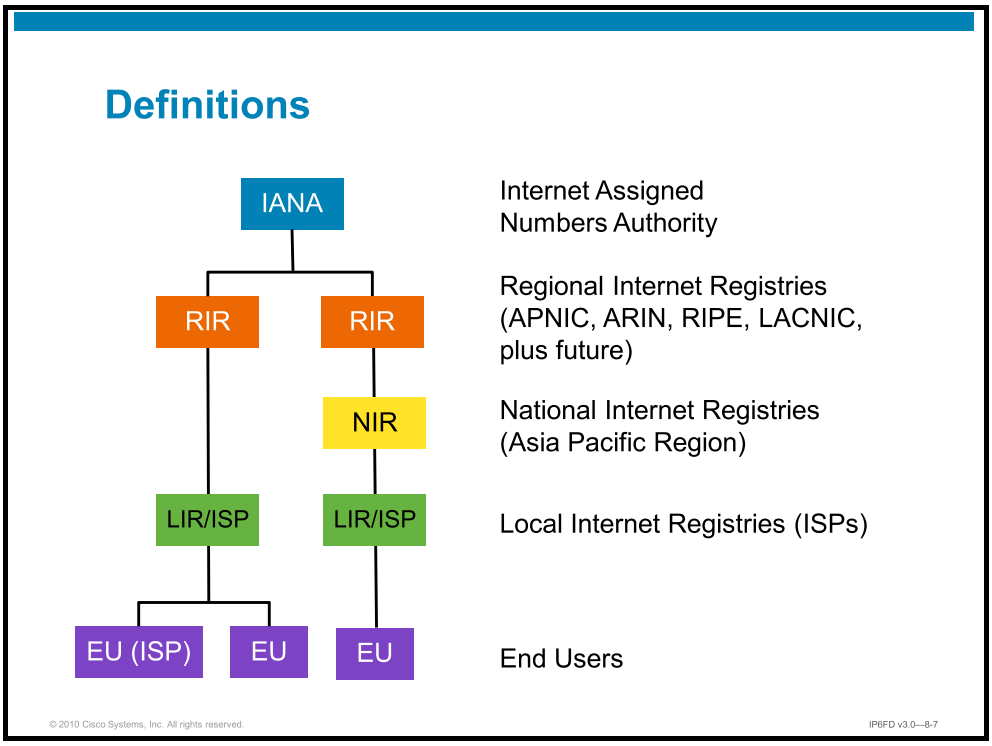

Some of the definitions that are used in address delegation are listed as follows:

- **Internet Assigned Numbers Authority (IANA):** IANA serves as a central coordinating organization for Internet numbering systems. In addition to being the top-level entity for assigning IP addresses to the RIRs, IANA also provides key Domain Name Services (DNS).
- **Internet Registry (IR):** An organization that is responsible for distributing IP address space to its members or customers and for registering those distributions. IRs are classified according to their primary function and territorial scope within the hierarchical structure that is depicted in the illustration.
- **Regional Internet Registry (RIR):** Regional communities establish and authorize RIRs, which are recognized by the IANA, to serve and represent large geographical regions. The primary role of RIRs is to manage and distribute public Internet address space within their respective regions.
- **National Internet Registry (NIR):** An NIR primarily allocates address space to its members or constituents, which are generally local Internet registries that are organized at a national level. NIRs exist mostly in the Asia Pacific region.
- **Local Internet Registry (LIR):** An LIR primarily assigns address space to the users of the network services that it provides. LIRs are generally ISPs whose customers are primarily end users or possibly other ISPs.

https://t.me/learningnets

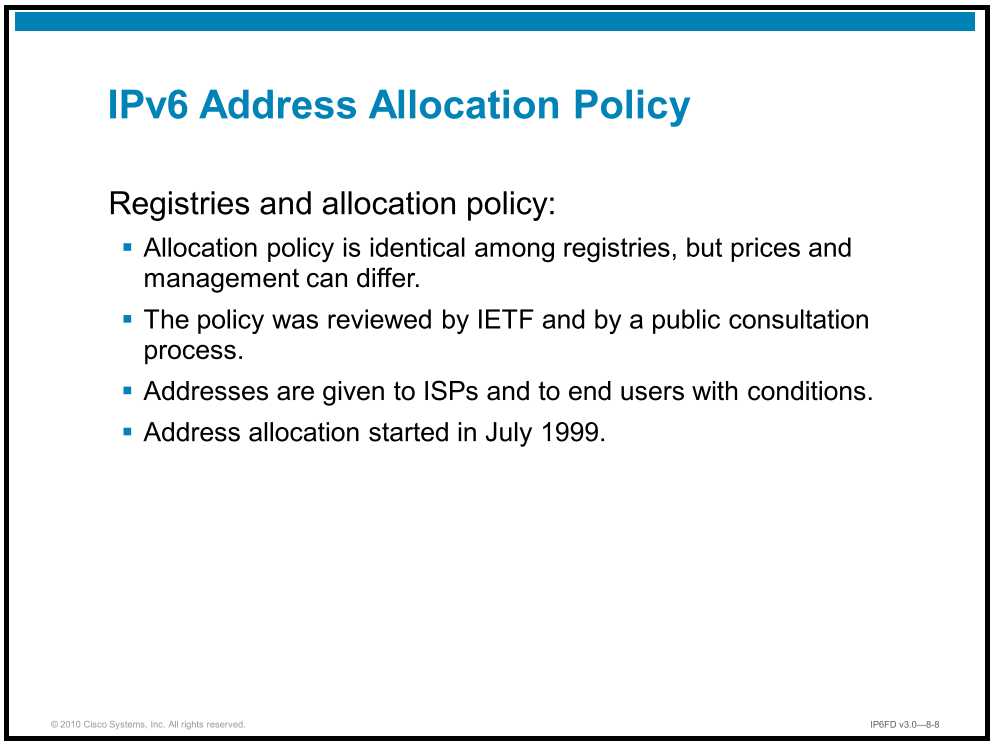

All registries generally share the same allocation policy to discourage ISPs from attempting to obtain address space from a favorite registry. Address space is not sold, but fees are charged for services and these fees can vary by region.

These organizations assign addresses with guidance provided by the IETF and input from the individual RIR membership. The allocations were first given in July 1999 with formal adoption of the policy in July 2002.

In the past, addresses were given only to ISPs to maintain routing aggregation for a large network. The changing requirements of IPv6 adopters are helping to establish new allocation policies.

As a result of growing demand from businesses, RIRs began supplying users with “Direct Assignment” addresses, enabling the use of provider-independent IPv6 address blocks. Until 2009, only provider-aggregatable address blocks were possible, giving the user a provider-dependent IPv6 address block. Users that are multihomed for implementation of high availability for their e-commerce applications, and so on, requested provider-independent addressing.

Note Directly assigned address space users have to meet certain conditions and pay a fee to the RIR because there is additional maintenance of provider-independent addresses. To avoid the high costs of becoming a direct assignment user, use your LIR to obtain provider-independent address space.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

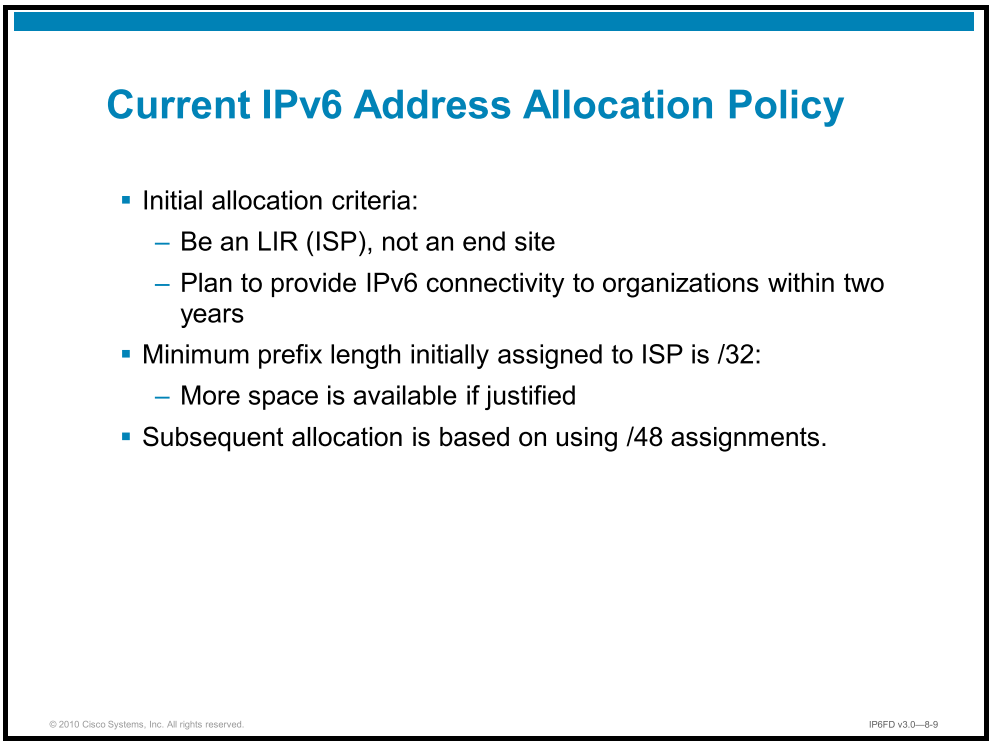

The current IPv6 address allocation policy states that, to qualify for an initial allocation, an organization must be an LIR (ISP). End sites are not eligible.

The ISP must also plan to provide IPv6 connectivity to organizations or clients. This connectivity includes advertising the reachability of these organizations through the ISP-aggregated address allocation and having a plan to assign a /48 prefix to at least 200 organizations within two years.

Organizations that meet the initial allocation criteria are eligible to receive a minimum allocation of /32. Organizations that hold a /35 IPv6 prefix allocation from the previous allocation policy are automatically entitled to receive a /32 prefix. This new prefix contains the already allocated /35 prefix.

As stated in the *IPv6 Address Allocation and Assignment Policy* (RIPE-NCC, September 2009) document, “Organizations may qualify for an initial allocation greater than /32 by submitting documentation that reasonably justifies the request. If so, the allocation size will be based on the number of existing users and the extent of the organization's infrastructure.” Numerous precedents exist for a larger allocation, including some /20 blocks.

The subsequent allocation criteria are based on using the /48 assignments that are done by the ISP. The Host-Density ratio (RFC 3194, *The Host-Density Ratio for Address Assignment Efficiency: An update on the H ratio*) is used to determine the thresholds that justify the additional address space required.

https://t.me/learningnets

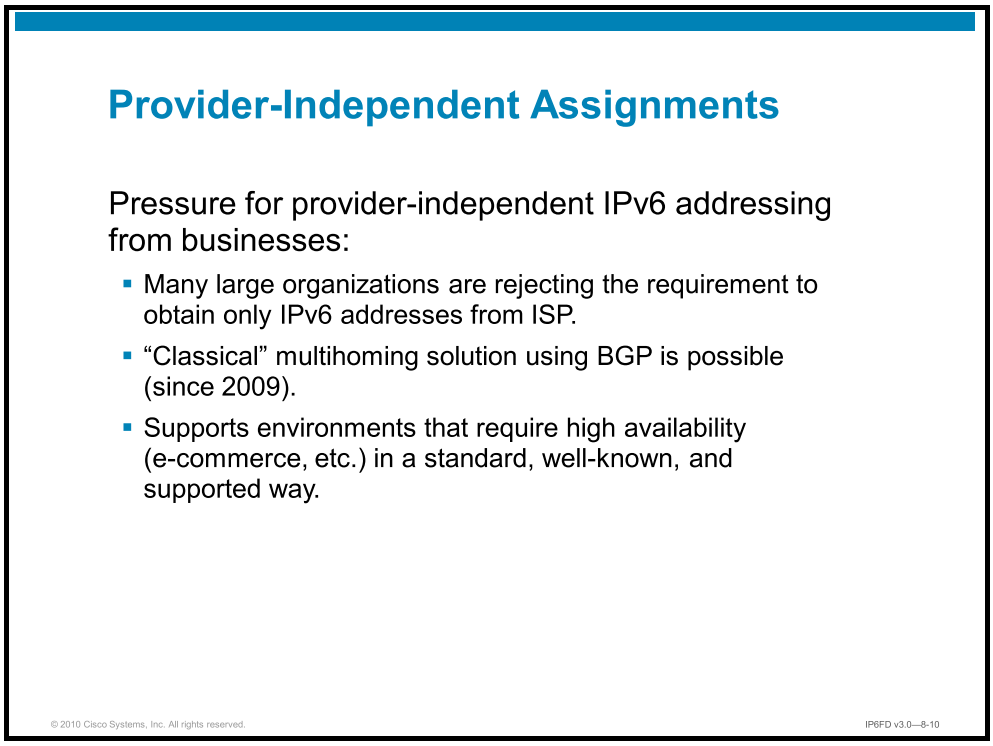

Provider-independent addresses supply one of the possible ways to make your network multihomed.

IANA, with the Internet Registrars, introduced the option for networks to obtain provider-independent IPv6 address allocations and assignments. Pressure increased from business clients that demanded multihoming, and the solutions that are offered in IPv6 were inadequate.

The reasoning for such implementations was that high availability is business-critical. The enterprises wanted a standard (or “usual”) implementation using Border Gateway Protocol (BGP) and not using IPv6-specifically offered solutions, such as Site Multihoming by IPv6 Intermediation (SHIM6) and Stream Control Transmission Protocol (SCTP).

The client should request a provider-independent address block through its service provider, or LIR. The prefix for the client will be announced in the global IPv6 Internet routing table.

The main downside of this approach is that it negates the scalability of the IPv6 routing table that has been planned since the beginning. It reintroduces many relatively small blocks in the global routing table, thereby “polluting” it.

Because IPv6 routes use more memory and bandwidth that will be exchanged among routers, wide use of provider-independent addressing will introduce BGP scalability challenges.

Note You can read more about multihoming topics in the dedicated lesson in this course.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

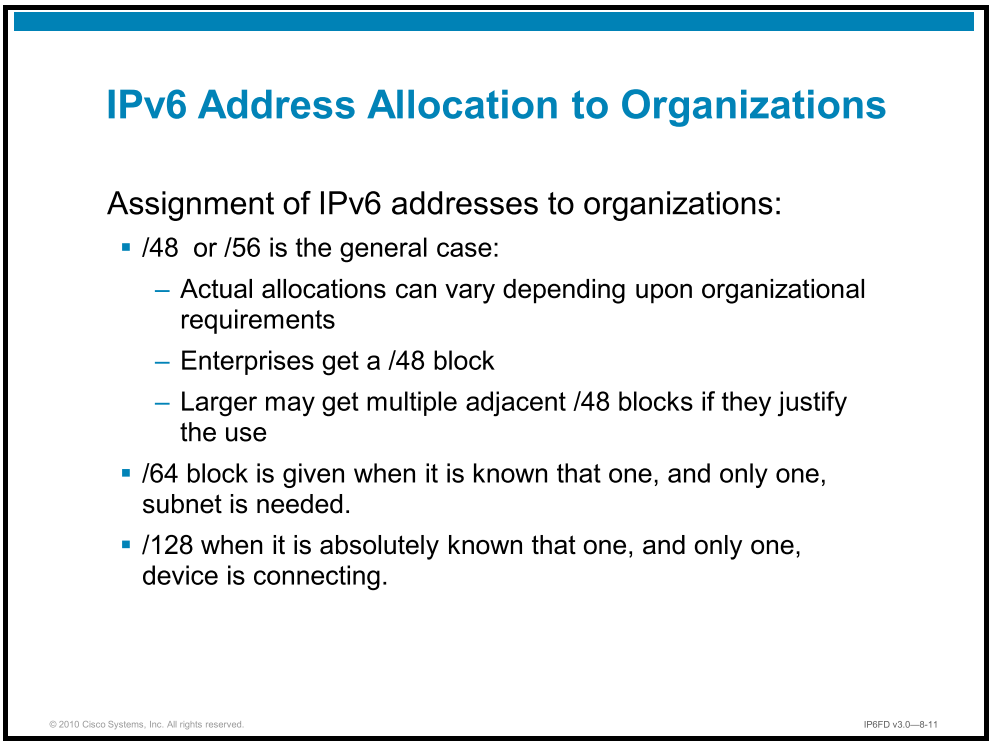

The *IPv6 Address Allocation and Assignment Policy* document also provides guidelines for address space assignment, as follows:

- For an ISP, the size of the IPv6 address assignment to the end-customer should be a /48. In early 2006, a new policy created greater allocation flexibility.
- Enterprises with smaller requirements may obtain /56 address space. If the ISP has enough /48 spaces that are allocated, it can bundle /56 classes. Therefore, an enterprise can scale while it is growing and maintain a summarized route (prefixes from /55 to /49).
- A /64 prefix size may be assigned, but only when it is known that one, and only one, subnet is needed. Such allocation can be used for home (residential) users.
- A /128 can be assigned, but only if it is absolutely known that one, and only one, device is connecting, for example, a dialup user or a mobile endpoint. This assignment is used very rarely.

These rules are based on the guidelines in RFC 3177, *IAB/IESG Recommendations on IPv6 Address Allocations to Sites* (http://www.ietf.org/rfc/rfc3177.txt).

Note ARIN proposal 2005-8 incorporated additional flexibility for allocations to end sites. In recognition of continued IP address conservation and the need for requirements larger than a /48, 2005-8 moved away from the strict allocation sizes that are prescribed in RFC 3177. This move will give ISPs the ability to better respond to customer demands. More information is available at http://www.arin.net/policy/proposals/2005_8.html.

https://t.me/learningnets

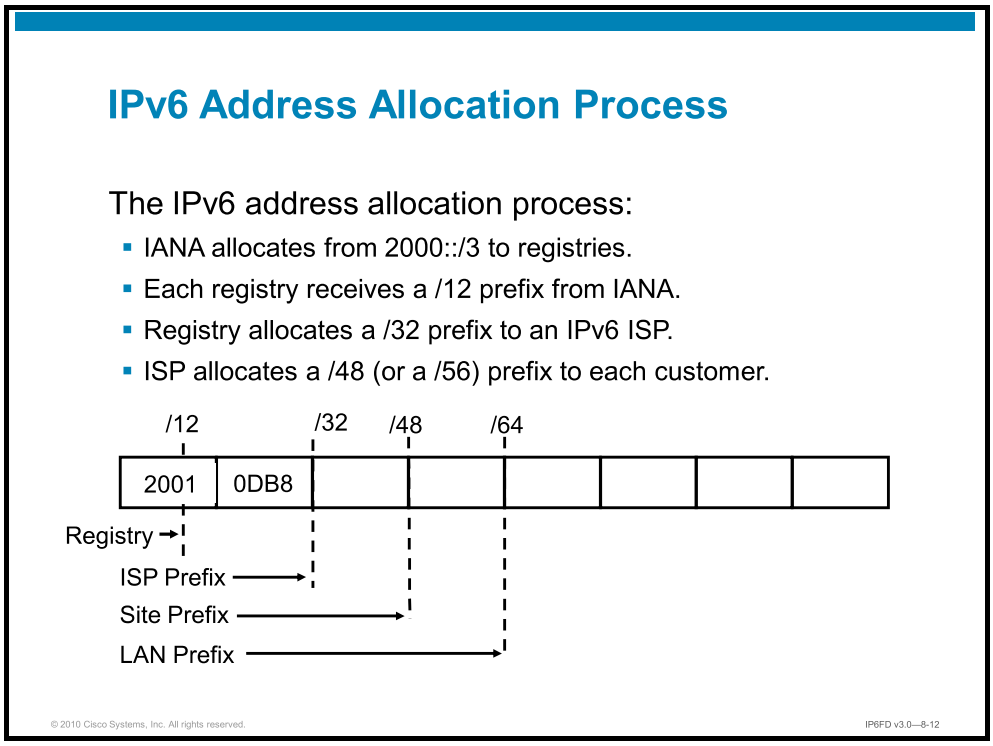

The generic allocation process is as follows:

- IANA allocates addresses to the registries from the 2000::/3 block.
- Each registry should receive a /12 from IANA.
- The registry allocates an initial /32 prefix to a new IPv6 ISP. A larger prefix can be allocated if justified by the ISP.
- The ISP allocates a /48 or a /56 prefix (out of the /32) to each customer, depending on the size of the customer network and its needs. Actual allocations may vary depending upon the requirements of the requesting organization.

Note In June 2005, a new policy was created for IPv6 address allocation from IANA to the registries. Demand for allocations larger than /32 had strained the abilities of the registries to meet their customer demands. Originally, the registries received a /23 from IANA. However, allocation requests for blocks in the /20 to /16 range required special provisioning. To meet growing demand and simplify the allocation process, a new policy was established that provided a /12 to the registries. No /12 has yet been allocated, although larger blocks, including a /16, have been allocated.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

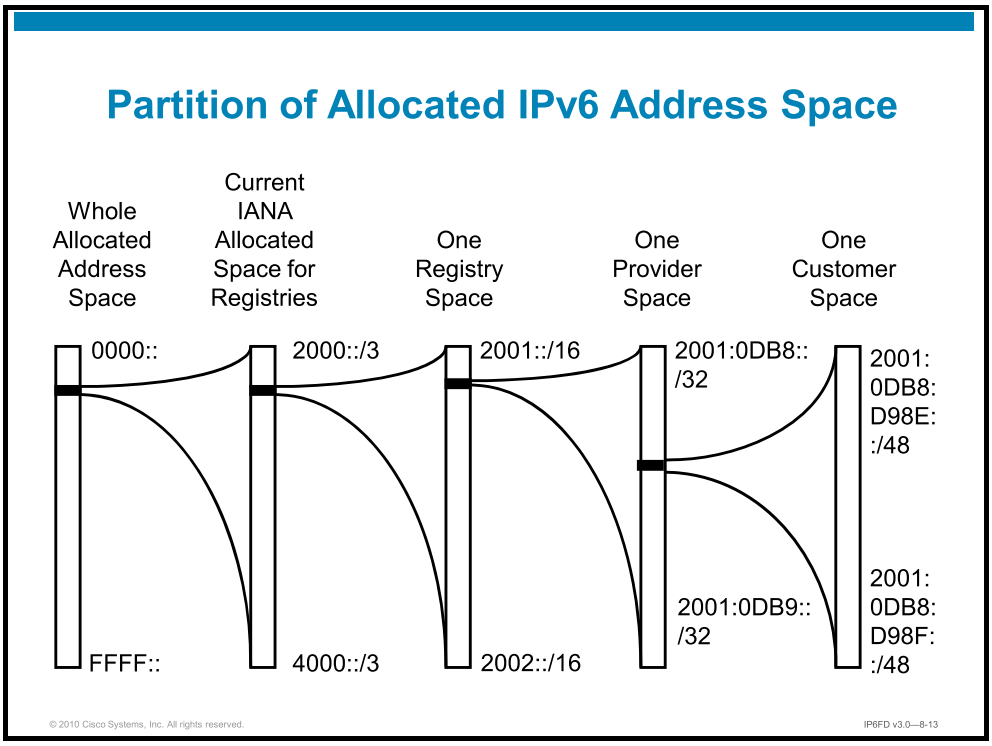

The figure shows an example of how allocated IPv6 address space is partitioned.

From the left:

- Within the whole address space, IANA currently allocates only the 2000::/3 portion to the registries. This space is quite large.
- From this block, IANA allocates /12 blocks to the registries.
- Within their own registry space (2001::/12 on the schema), registries allocate /32 to each provider, for example, 2001:0DB8::/32.
- A provider then assigns /48 to its end customers.
- The customer receives 2001:0DB8:D98E::/48.

https://t.me/learningnets

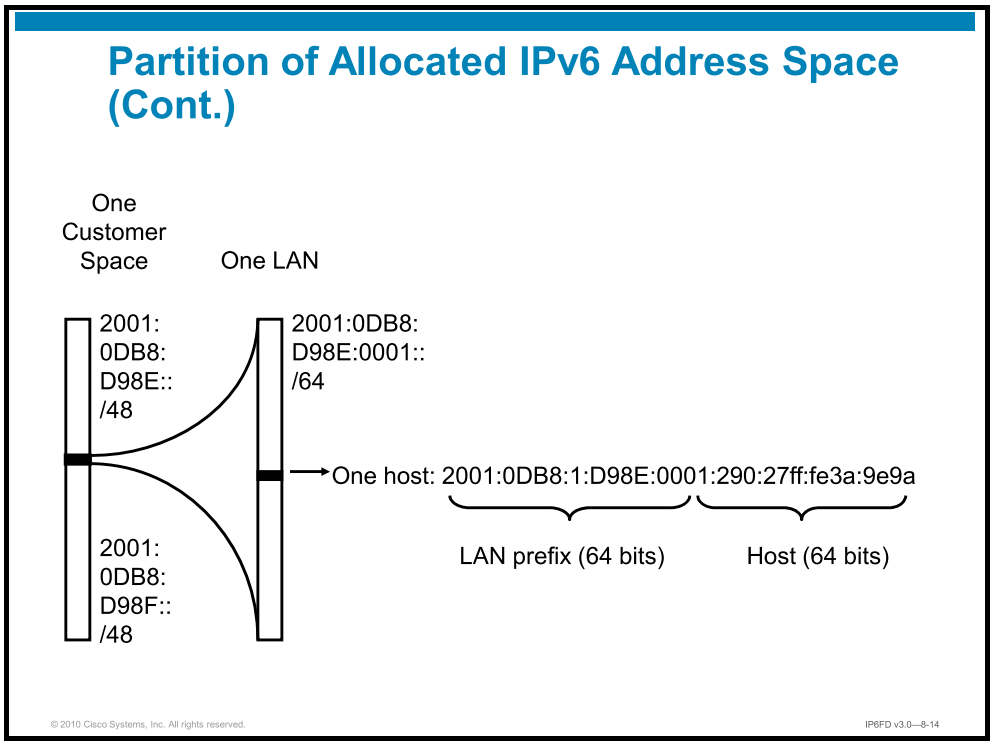

- A site receives a /48 prefix that is further allocated to LANs. Each LAN is a /64 prefix, for a maximum of 264 hosts in each LAN.

Therefore, 16 bits (64 to 48) are available to subnet the site, for a maximum of 65,536 LANs per site. A site should have an address plan before beginning to allocate its /48 space.

In the example, each LAN receives a /64 prefix from the customer space. The LAN receives 2001:0DB8:D98E::/64. The host part of the address is the 64 bits on the far right. For autoconfiguration, the host part is derived from the link interface address and assigned automatically. In the example, the host part is :0290:27ff:fe3a:9e9a. The complete address of a host is the concatenation of the LAN prefix and the host part.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Connecting to the IPv6 Internet

This topic describes how IPv6 ISPs connect to each other.

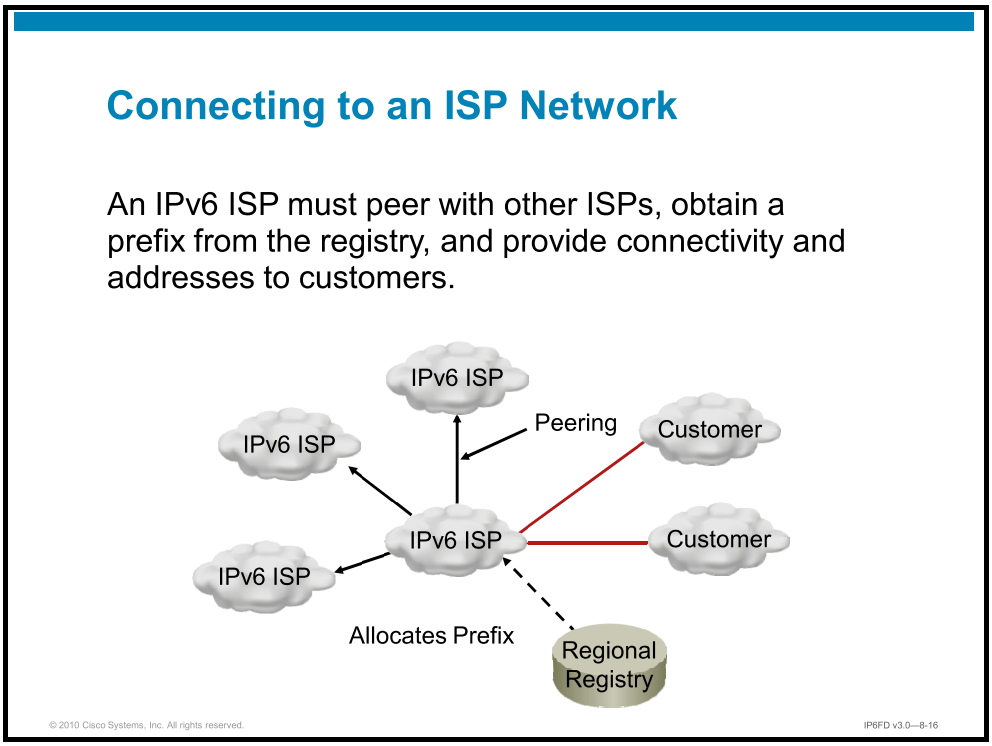

Here is a simplified sequence of steps that are required for an ISP to connect to the IPv6 production Internet:

Step 1 Peer with other ISPs.

Step 2 Obtain address prefix from a registry.

Step 3 Allocate address prefix to customers.

Step 4 Aggregate the routing entries from the customers so that only the allocated prefix is announced.

As in the current IPv4 Internet, a small ISP can be connected to a larger ISP that is connected to an even larger ISP. The number of intermediate levels is not fixed. IPv6 does not change this behavior. However, the current IPv6 address allocation forces the aggregation of routing entries, which should limit the number of entries in the global routing table. This situation may change somewhat with the introduction of IPv6 provider-independent addressing.

https://t.me/learningnets

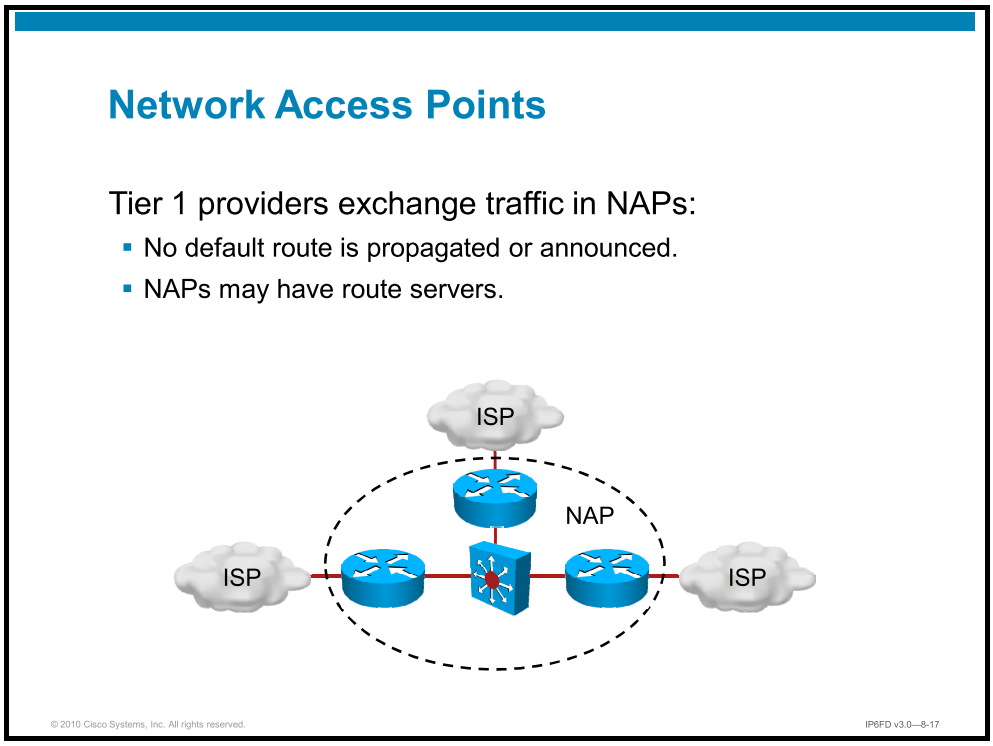

Tier 1 providers exchange traffic in locations that are called network access points, as follows:

- They have a high-speed LAN to which providers come with links and routers.
- Routes are exchanged.
- No default route is propagated or announced.
- Some network access points (NAPs) have route servers to optimize the peering and the route policing.

Because NAPs are generally Layer 2 exchanges, it is easy for them to enable native IPv6 on their infrastructures. Agreements among the peers and the exchange do not need to be changed because IPv6 does not really introduce new issues. However, if a route server is needed, the route server must be IPv6-enabled.

Because of the transition from IPv4 to IPv6, some NAPs are offering transition mechanisms because they fit well in those locations. A good example of an additional NAP service is a 6to4 relay.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Connecting an Enterprise Network

To connect an enterprise network to the IPv6 Internet, perform the following steps:

- Connect the enterprise network to an IPv6 ISP.
- Obtain an end-site prefix from the ISP.

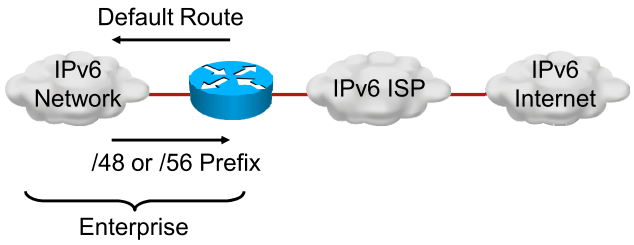

© 2010 Cisco Systems, Inc. All rights reserved.

Follow these steps to connect an enterprise network:

Step 1 Find an IPv6 ISP.

Step 2 Connect the enterprise network to the ISP.

Step 3 Obtain an end-site prefix from the ISP.

Step 4 Design an address plan for the enterprise network.

Step 5 Allocate /64 prefixes to each LAN.

Multihoming is possible with IPv6. The IETF is working to enable IPv6 multihoming while still enforcing the aggregation of routing entries.

One methodology is carried out through a process that is called address selection, in which the host chooses the source address prefix from one of its connected providers.

Other techniques involve building multiple connections to providers, working with long-haul providers for redundancy at Layer 2, or using Mobile IPv6 techniques.

The third approach is most recent and is accomplished by using provider-independent IPv6 address blocks and routing using BGP to implement path redundancy.

https://t.me/learningnets

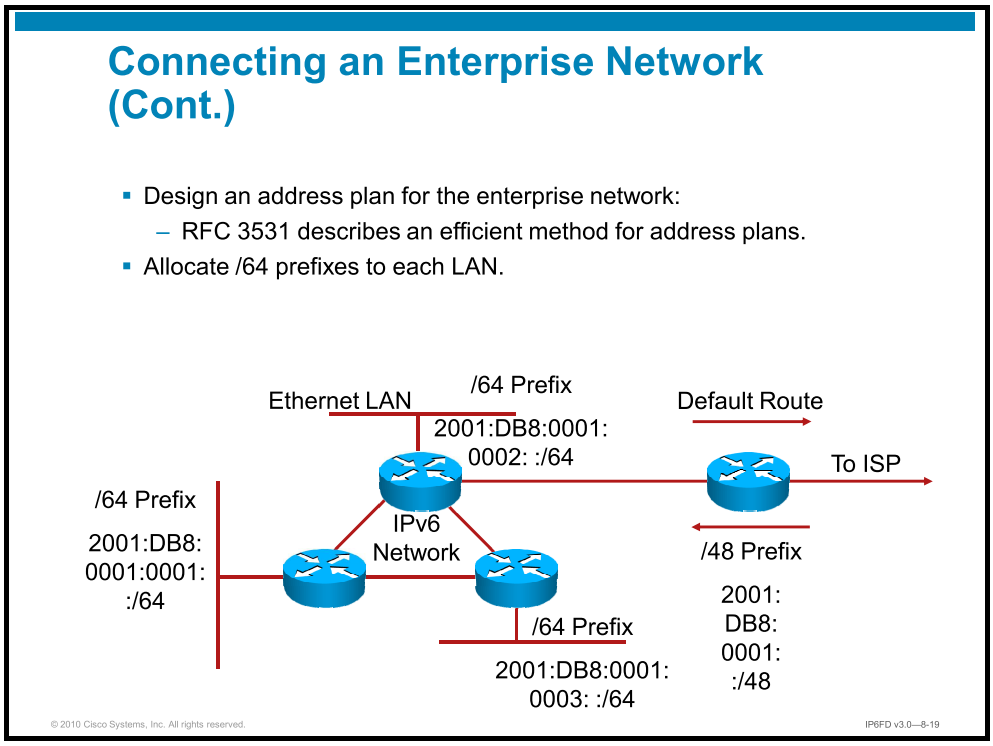

The figure shows the allocation of prefixes to LANs inside an enterprise network. The site receives the 2001:0DB8:0001::/48 from its upstream provider. Each LAN has a /64 prefix from the site prefix. Three LANs are shown:

- Left: 2001:0DB8:0001:0001::/64
- Center, top: 2001:0DB8:0001:0002::/64
- Right, bottom: 2001:0DB8:0001:0003::/64

IETF RFC 3531, *A Flexible Method for Managing the Assignment of Bits of an IPv6 Address Block*, describes different methods for developing an allocation address plan.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

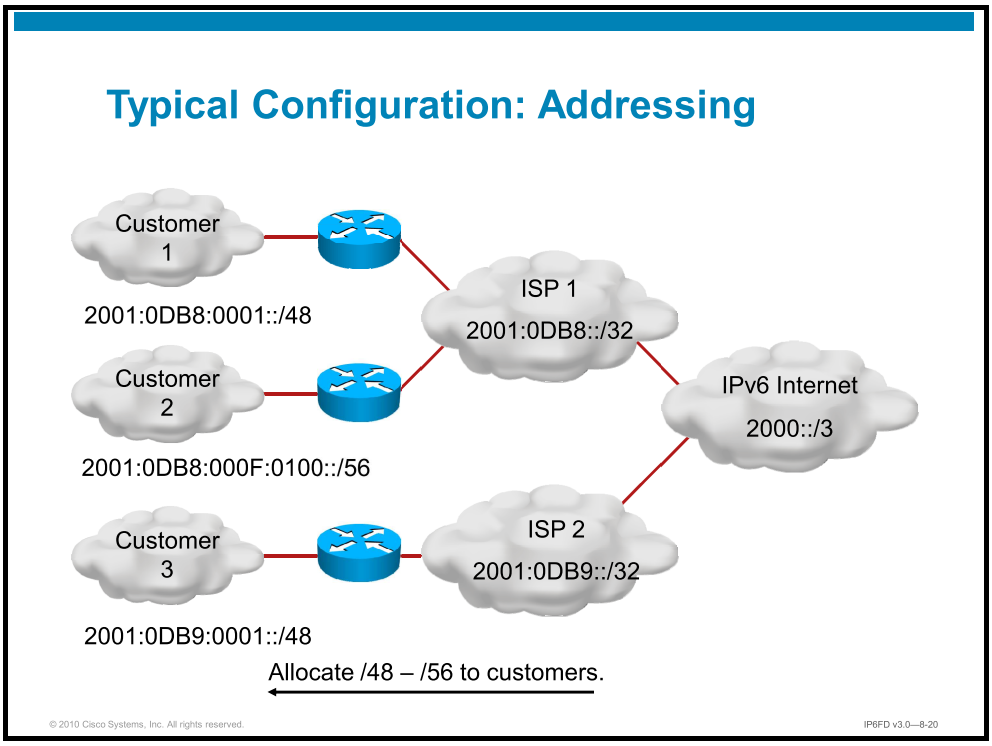

The figure shows a typical address allocation of three customer sites from two different ISPs.

ISP 1 has allocated a 2001:0DB8::/32 prefix from its regional registry. It assigns /48 to its customers (Customers 1 and 2) from its own prefix. A simplified algorithm for allocation is to start from the first /48 prefix available and increase.

Customer 1 is allocated 2001:0DB8:0001::/48.

Customer 2 has a smaller network and is allocated a 2001:0DB8:000F:0100::/56 prefix.

ISP 2 is allocated a 2001:0DB9::/32 prefix from its regional registry. It assigns 2001:0DB9:0001::/48 to its customer, Customer 3.

https://t.me/learningnets

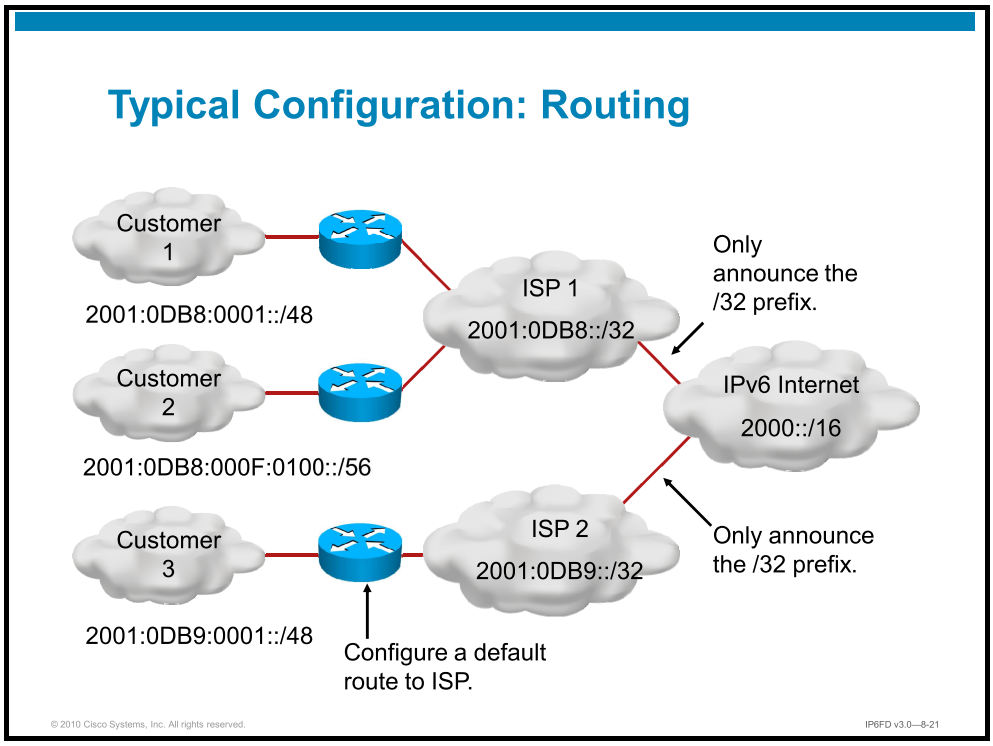

The figure shows a typical routing configuration for three customer sites with two providers. The sites configure a default route pointing to their upstream providers. Providers aggregate their customer routes and announce only their own /32 prefixes.

https://t.me/learningnets

Connecting as a Host Using a Tunnel Server

A tunnel server can be used to connect a single host or a small network over the IPv4 Internet.

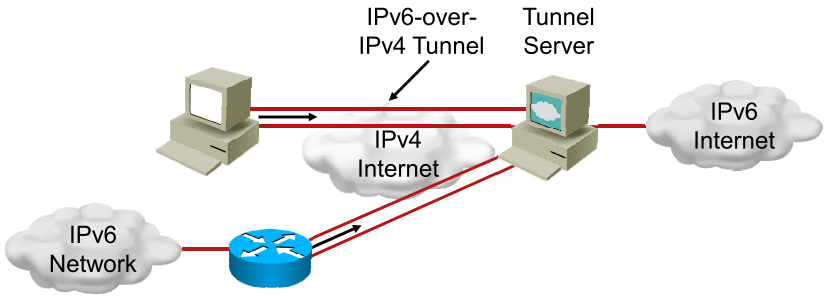

© 2010 Cisco Systems, Inc. All rights reserved.

An individual host that is located on the current IPv4 Internet can link to the IPv6 Internet by using a transition mechanism. One of the most common mechanisms is an IPv6 over IPv4 tunnel.

Of the various tunneling approaches, the tunnel broker (or tunnel server) approach is the most common and the easiest for individual hosts and small networks to use. The figure shows a host that is connected to a tunnel server by using an IPv6 over IPv4 tunnel up to the server. It also shows a network that is connected to a tunnel server. The server de-encapsulates the packets and routes them to the IPv6 Internet.

In most cases, the host must have a globally routable IPv4 address to serve as one of the tunnel endpoints. It cannot be an RFC 1918 address behind a Network Address Translation (NAT). However, certain tunnel services support NAT traversal, which would overcome the requirement for a unicast IPv4 address.

Many free tunnel servers offer IPv6 tunnels over the Internet, enabling hosts and networks to connect freely with the IPv6 Internet using the current IPv4 Internet as an overlay transport layer.

https://t.me/learningnets

Connecting as a Host Using 6to4

The 6to4 tunneling mechanism can be used to connect a single host or a network over the IPv4 Internet.

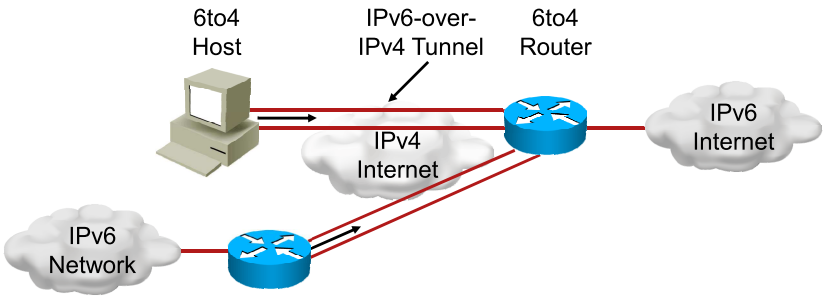

© 2010 Cisco Systems, Inc. All rights reserved.

Another tunneling mechanism to connect an individual host or network to the IPv6 Internet is the 6to4 mechanism. The figure shows both the network case and the host case.

Both the host and the router must have a 6to4 IPv6 stack to enable this scenario.

In this case, no server is needed. All 6to4 devices can communicate over the IPv4 Internet. One of them should provide the 6to4 relay functionality to enable routing to the regular IPv6.

https://t.me/learningnets

Summary

This topic summarizes the key points that are discussed in this lesson.

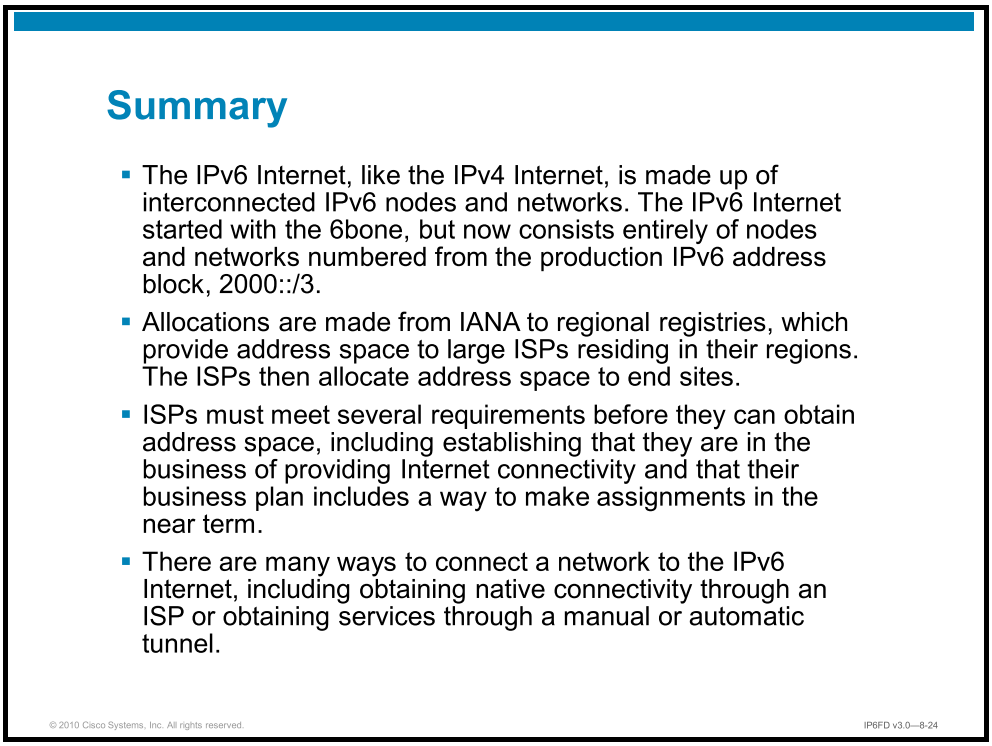

Resources

To learn more about IPv6 deployment, refer to the following resources:

- **IANA Number Resources** at http://www.iana.org/numbers/
- **RIPE-NCC Internet Resources** at http://www.ripe.net/rs/index.html
- **RIPE-NCC: Requesting Independent Resources** at http://www.ripe.net/rs/independent-resources.html
- **ARIN Number Resources** at https://www.arin.net/resources/index.html
- **IPv6 Address Allocation and Assignment Policy** at http://www.ripe.net/docs/ipv6policy.html

https://t.me/learningnets

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Understanding the IPv6 Multihoming Issue

Overview

A key aspect of designing the architecture for an IPv6 network is identifying any IPv6 multihoming issues. It is important that you understand IPv6 multihoming solutions to provide your client with a robust network. This lesson describes the issues about IPv6 multihoming and the proposed solutions.

Objectives

Upon completing this lesson, you will be able to identify an IPv6 multihoming issue and prescribe a potential solution. This includes being able to meet these objectives:

- Outline the multihoming issues in IPv6
- Review the current IPv6 multihoming solutions

https://t.me/learningnets

IPv6 Multihoming Aspects and Issues

This topic describes the multihoming issues in IPv6 networks.

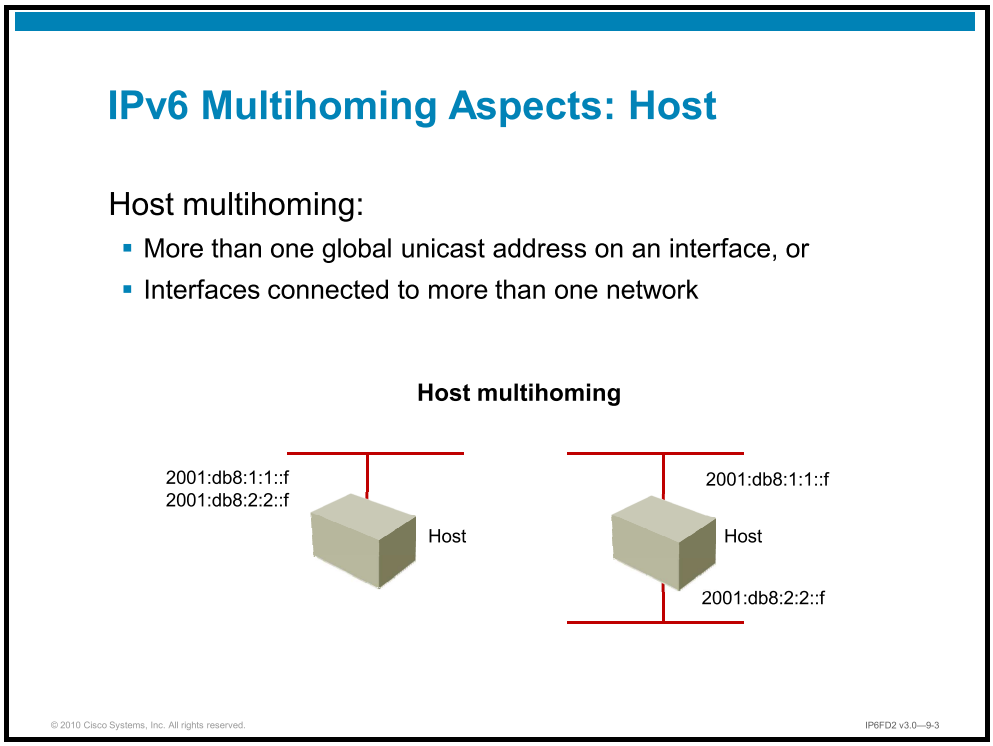

There are several aspects of multihoming being examined: host multihoming, site multihoming, and provider independence.

Generally, multihoming is the practice of joining a large site network to multiple upstream Internet service providers (ISPs). It provides redundancy against router failure, link failure, and ISP failure.

Host Multihoming

With host multihoming, each host has multiple addresses that are drawn from the prefixes assigned by the multiple ISPs serving the site. This technique is used to ensure that all packets sourced from site hosts can exit from the site network via the ISP network (or networks). Each ISP requires packets to be sourced from the /48 prefix that they assign for use within the site; therefore, host multihoming for Internet access is always combined with site multihoming.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

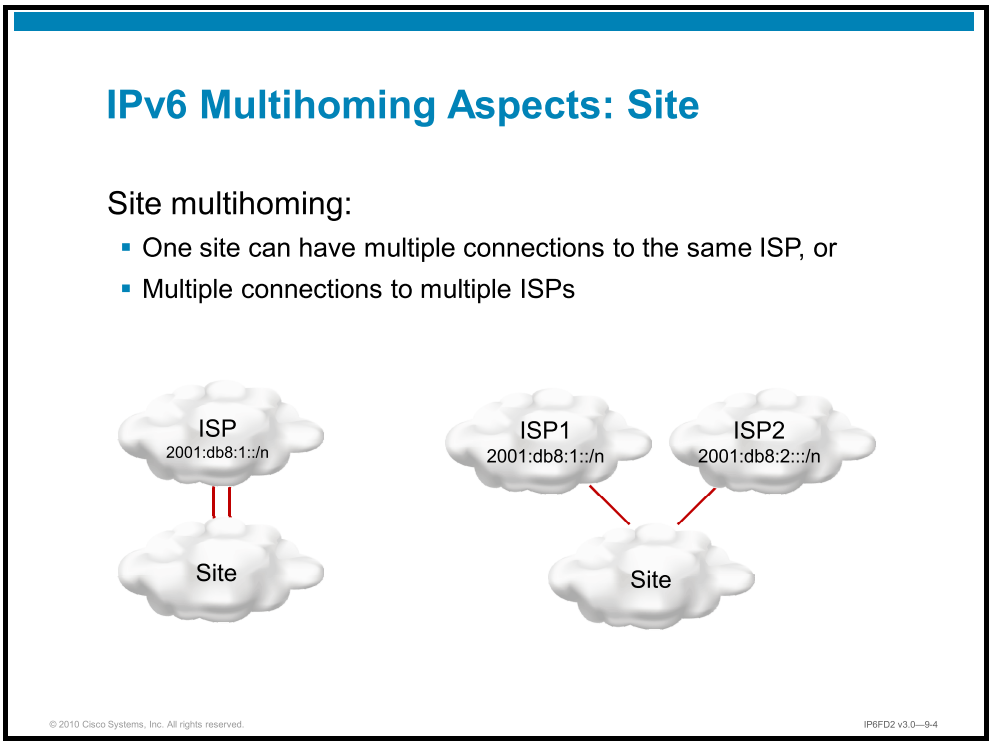

Site Multihoming

Site multihoming (connecting to more than one ISP) is an essential service for enterprise customers and small or mid-sized companies. Current IPv4 site multihoming practices have been added on to the classless interdomain routing (CIDR) architecture, which assumes that routing table entries can be aggregated based upon a hierarchy of customers and service providers.

There has been an enormous growth in the number of multihomed sites. For purposes of redundancy and load sharing, the multihomed address blocks are introduced into the global table, even if they are covered by a provider aggregate. This contributes to the rapidly increasing size of both the global routing table (and the turbulence exhibited within it) and places stress on the interprovider routing system.

https://t.me/learningnets

Provider-Independent vs. Provider-Allocated

This subtopic describes the *IPv4* provider independence.

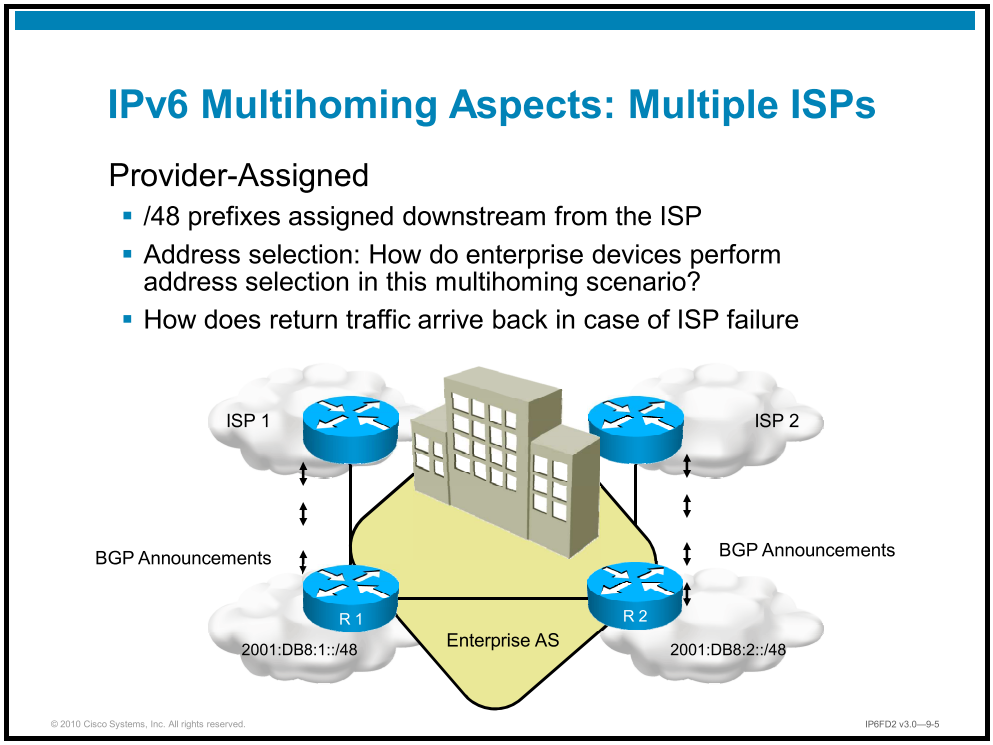

For most of its history, IPv4 used the concept of portable addresses. Enterprises obtained blocks of addresses that they would use for any purpose, and these address blocks could appear anywhere on the Internet. For example, if a multinational company received an address block 192.0.2.0/16, it could deploy Class C addresses from that block anywhere in the world. In this situation, the global routing table had to have a route to each of those networks because there was no way to summarize (or aggregate) the routes.

IPv6 provider-independent addressing is another multihoming solution for IPv6 networks which has been available only recently (since 2009).

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

IPv6 Multihoming Aspects: Multiple ISPs (Cont.)

Provider Independence (IPv4 and IPv6)

- No need for renumbering on provider change
- Explosion of the routing table
- Massive loss of aggregation

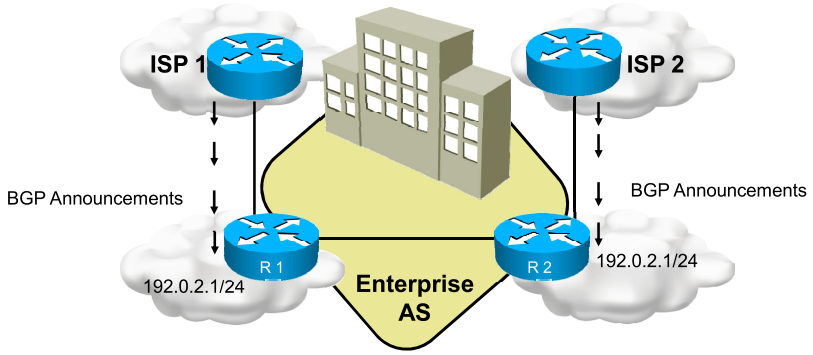

© 2010 Cisco Systems, Inc. All rights reserved.

IPv6 initial design is to use provider-based addressing, which means that the enterprise receives its addresses from its ISP. If applied to the example, the multinational company would receive a /48 prefix from each provider in an area where that network will be deployed. This allows the routes to be aggregated or summarized by region or logical topology, which greatly reduces the size of the backbone routing tables.

However, the enterprise in this example is prohibited from advertising the prefix from ISP 1 up to ISP 2 and vice versa—this is known as “forced aggregation” and ends up preventing IPv4-style site multihoming from being an option in IPv6.

The IPv6 design is great for reducing the size of the backbone routing tables through summarization, but it makes the problem of multihoming a network without advertising the prefixes more difficult.

https://t.me/learningnets

PA IPv6 Addressing

This subtopic describes the provider-assigned address space characteristics.

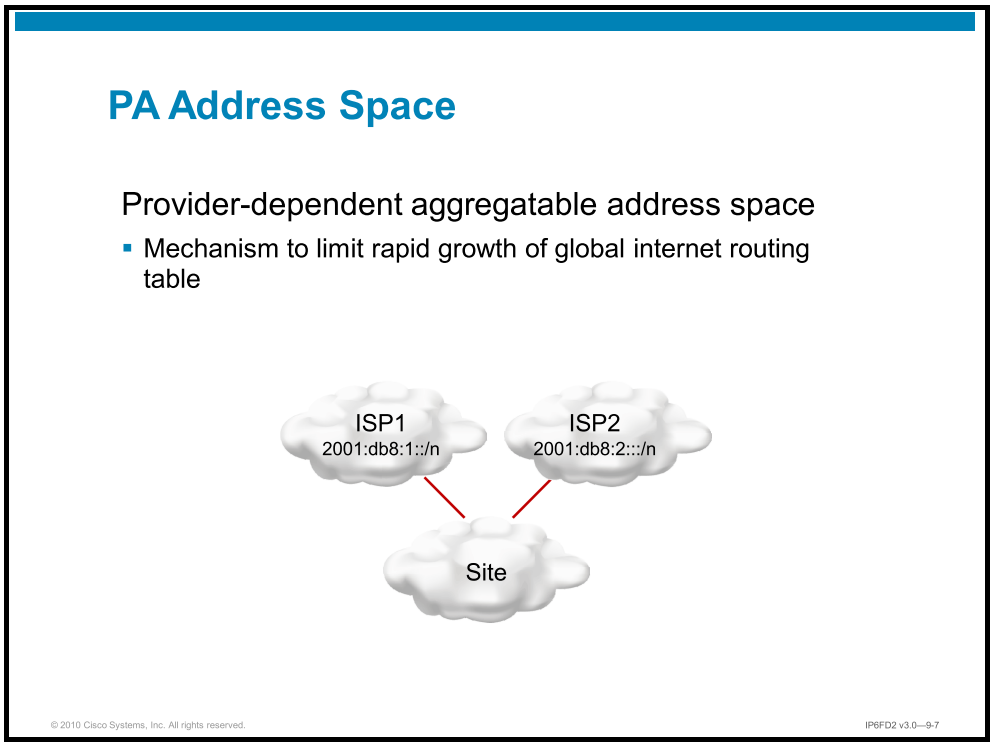

Due to the growth of the BGP routing tables, IPv6 multihoming is required to preserve the scalability of the interdomain routing system. One of the methods to preserve the table size is assigning provider dependent aggregatable address space to customer sites.

PA address space is typically assigned by the ISP. The ISP assigns the address space in a fashion which enables the ISP to aggregate the address space and keeps the number of routing table entries at an acceptable level. As a result the total number of aggregated routes is much lower when compared to propagating each customer routes through the whole interdomain routing system.

PA space is assigned by the ISP to the customer and needs to be returned after the customer decides to terminate the contract. The assigned address space is therefore returned to the ISP and the customer will be required to reconfigure devices. ISP which assigned the address space will be able to reuse the returned block.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

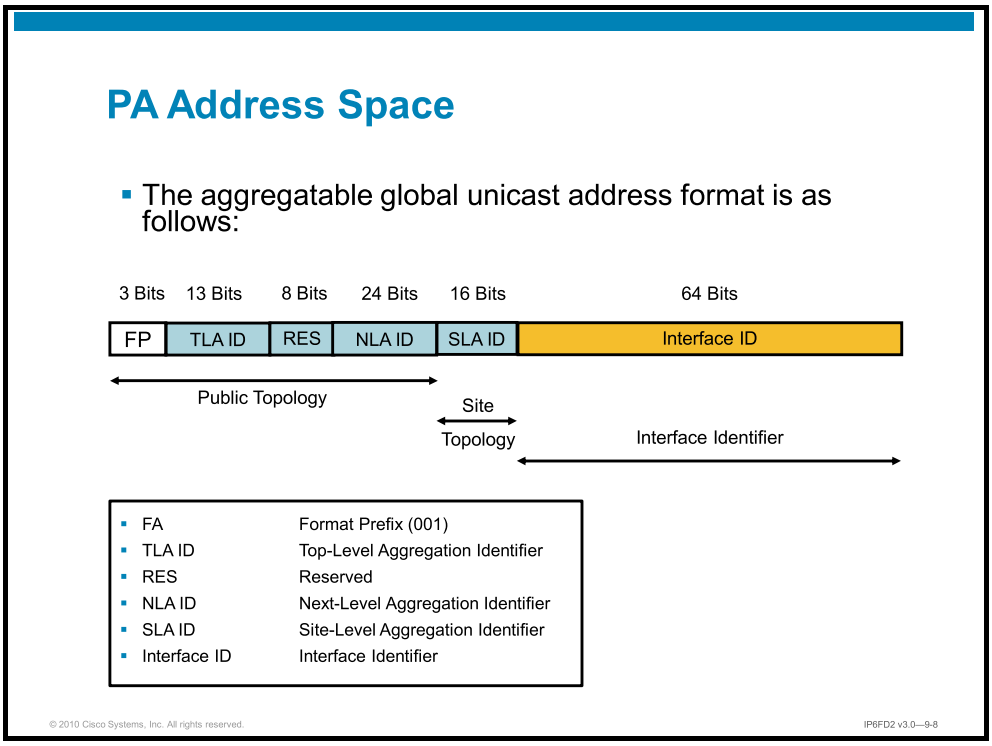

PA address space consists from the following fields:

- **Top-Level Aggregation Identifiers**: TLA ID is the top level in the routing hierarchy. Default-free routers must have a routing table entry for every active TLA ID and will have additional entries providing routing information for the TLA ID in which they are located. They may have additional entries in order to optimize routing for their specific topology, but the routing topology at all levels must be designed to minimize the number of additional entries fed into the default free routing tables.
- **Reserved**: The Reserved field is reserved for future use and must be set to zero. The Reserved field allows for future growth of the TLA and NLA fields as appropriate.
- **Next-Level Aggregation Identifier**: Next-Level Aggregation Identifier's are used by organizations assigned a TLA ID to create an addressing hierarchy and to identify sites. The organization can assign the top part of the NLA ID in a manner to create an addressing hierarchy appropriate to its network.
- **Site-Level Aggregation Identifier:** The SLA ID field is used by an individual organization to create its own local addressing hierarchy and to identify subnets. This is analogous to subnets in IPv4 except that each organization has a much greater number of subnets. The 16 bit SLA ID field supports 65,535 individual subnets.
- **Interface Identifier:** Interface identifiers are used to identify interfaces on a link. They are required to be unique on that link. More information can be found in the previous lessons.

https://t.me/learningnets

PI IPv6 Addressing

This subtopic describes the provider-independent address space characteristics.

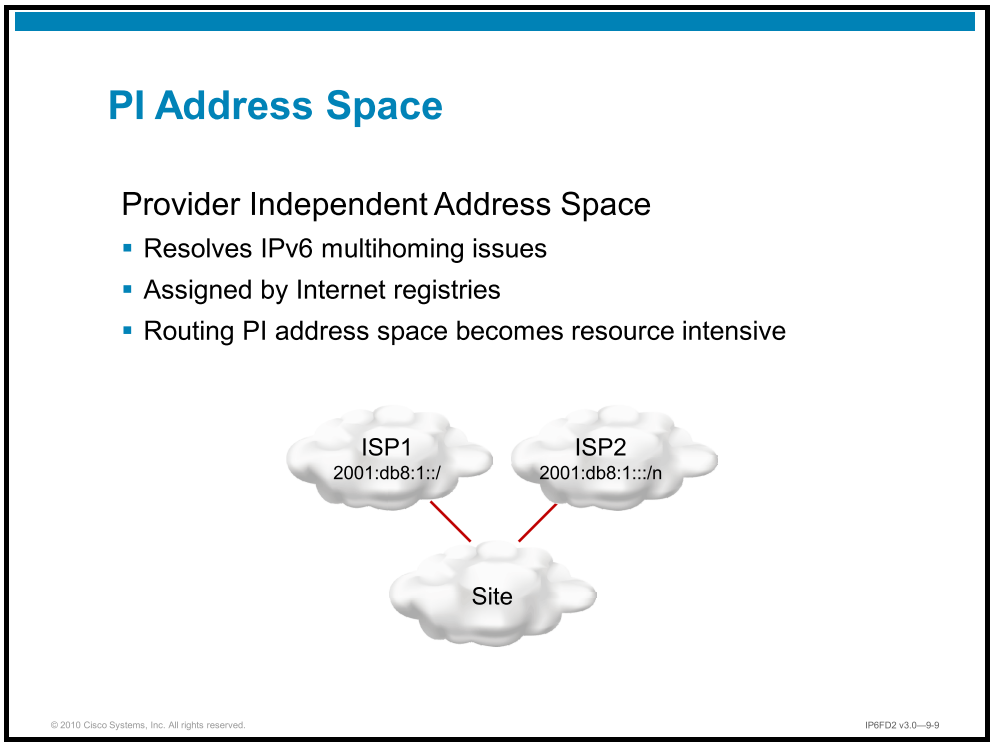

PI address space gets assigned by Regional Internet Registries (RIR) which are governed by IANA and the end customer will be assigned an independent address space which can remain in use at the customer site even after the termination of the ISP contract. Reconfiguration of network devices will not be required with PI address space.

PI address space also resolves any multihoming issues which an IPv6 enabled end user would be faced with using PA address space. The end user will be able to advertise only one block of globally unique address space to its ISP.

Current concern with PI address space is that aggregation is not possible in the same manner as with PA address space. Immediate issues include an expected growth of the global routing table and higher resource requirements for routing PI addresses.

There are also other IPv6 protocol based multihoming solutions:

- SHIM6
- MAST
- SCTP
- HIP
- LISP

Note Frequently the end customer will have to provide or demonstrate multihoming requirements to be assigned PI address space. The minimum size of the assignment will be /48. Organizations which require larger assignments will have to provide justification for additional subnets to the RIRs.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

PA versus PI

This subtopic shows an example of provider-assigned and provider-independent address space and their differences.

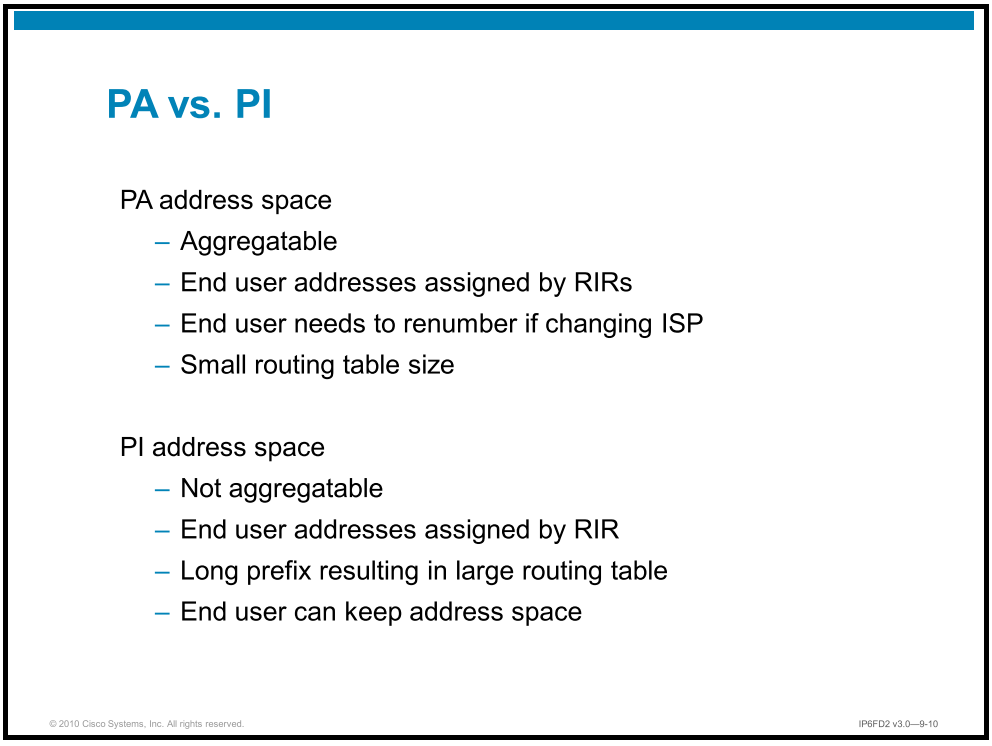

In comparison PA adress space offer aggregation capabilites, however end users will be forced to renumber or reconfigure networking devices after a change of ISP.

On the contrary PI address space will be assigned by a RIR and offer many advantages especially in resolving multihoming issues. End user will also be able to keep the assigned address space even after a change of ISPs, so no renumbering will be needed. The negative side is that PI address space requires the block to be globaly routabel which can influence the explosive growth of the routing table.

As a current policy by IANA and RIRs end users will be able to request both types of adress space from either ISP ove f RIR. In case of a PI request the end user needs to provide additional documentation describing multihoming requirements.

https://t.me/learningnets

Internet Routing Table Growth

This subtopic describes current IPv6 announced prefixes and expected global internet routing table growth.

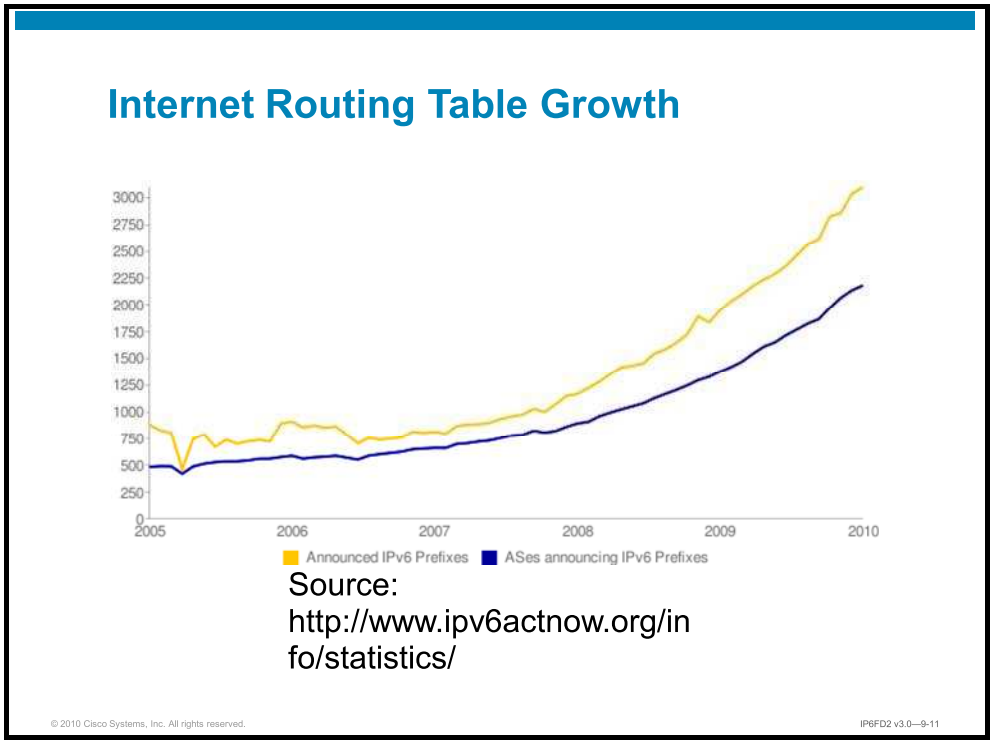

The figure displays a graphical representation of IPv6 address space on the internet as seen by RIPE NCC. At the moment of writing there are approximately 3000 IPv6 prefixes announced worldwide by 2000 ASes and the trend is growing. New allocation requests are mostly made by countries which cannot request new IPv4 address space from internet registries.

In comparison the IPv4 global routing table currently holds approximately 320 000 prefixes and should gradually stop increasing since IANA has declared IPv4 address space depletion to happen in 2011.

Note To view current statistics several ISPs and educational organizations have set up looking glass routers which hold the full BGP table for statistical purposes for both IPv4 and IPv6 prefixes.

There are some resource considerations involved in deploying an IPv6 capable router. Additional memory resources will be required as an IPv6 address is four times the size of an IPv4 address.

PA address space is the main method with which IANA would like to mitigate explosion of the IPv6 global routing table size, since it can be aggregated by an ISP to minimize number of prefixes which need to be announced to neighboring ASes.

Note Source of the figure is http://www.ipv6actnow.org/info/statistics/ which includes near real time generated graphs.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Multihoming Requirements

This subtopic describes multihoming goals.

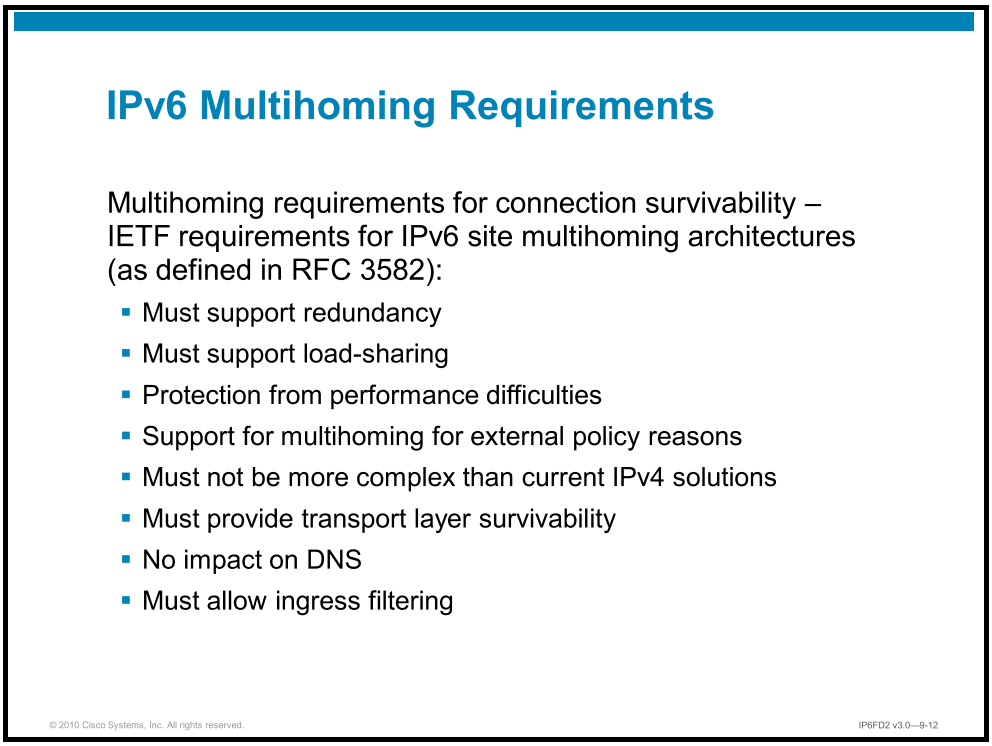

The multihoming architecture, as it is defined in RFC 3582, should accommodate continuity of connectivity to meet the following criteria:

- Physical failure, such as a fiber cut or router failure
- Logical link failure, such as a misbehaving router interface
- Routing protocol failure, such as a Border Gateway Protocol (BGP) peer reset
- Transit provider failure, such as a backbone-wide Interior Gateway Protocol (IGP) failure
- Exchange failure, such as a BGP reset on an interprovider peering

The IETF requirements for IPv6 site multihoming architectures include the following:

- **Load-sharing:** Load-sharing a site by multihoming should enable the distribution of both inbound and outbound traffic between multiple transit providers. The goal of load-sharing is concurrent use of multiple transit providers, not just the usage of one provider over one interval of time and another provider over a different interval of time.
- **Protection from performance difficulties:** By multihoming, a site should be able to protect itself from performance difficulties between the transit providers of the site. For example, suppose site E obtains service from transit providers T1 and T2, and there is long-term congestion between T1 and the Internet. The multihoming architecture should allow site E to ensure that, in normal operation, none of its traffic is carried over the congested interconnection via T1.

https://t.me/learningnets

- **Policy:** A customer may choose to use multihoming for a variety of policy reasons beyond the technical scope (for example, cost, acceptable use conditions, and so on). For example, customer C, who is homed to ISP A, may wish to shift traffic of a certain class or application, Network News Transfer Protocol (NNTP), for example, to ISP B as a matter of policy. A new IPv6 multihoming proposal should provide support for site multihoming for external policy reasons.
- **Simplicity:** Because any proposed multihoming solution must be deployed in real networks with real customers, simplicity is paramount. The current IPv4-based multihoming solution is quite straightforward to deploy and maintain. A new IPv6 multihoming solution should not be substantially more complex to deploy and operate for multihomed sites or for the rest of the Internet.
- **Transport layer survivability:** Multihoming solutions should provide rehoming transparency for transport layer sessions (that is exchange of data between devices on the multihomed site), and devices elsewhere on the Internet may proceed with no greater interruption than that associated with the transient packet loss during the rehoming event. New transport layer sessions should be able to be created following a rehoming event. Transport layer sessions include those involving transport layer protocols, such as TCP, User Datagram Protocol (UDP), and Stream Control Transmission Protocol (SCTP) over IP.
- **No Impact on DNS:** Multihoming solutions either should be compatible with the observed dynamics of the current Domain Name System (DNS), or the solutions should demonstrate that the modified name resolution system required to support them is readily deployable.
- **No Limitations on Ingress Filtering:** Multihoming solutions should not preclude filtering packets with forged or otherwise inappropriate source IP addresses at the administrative boundary of the multihomed site or at the administrative boundaries of any site in the Internet.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Multihoming and IPv6

This subtopic describes advanced features of IPv6 and multihoming.

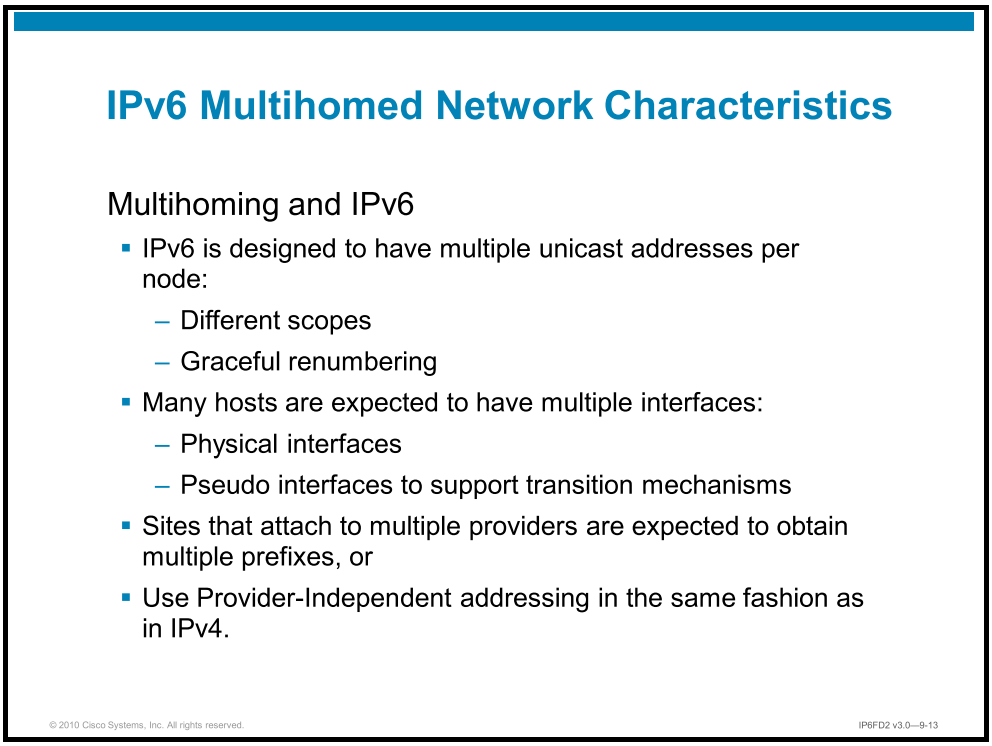

IPv6 multihoming is based on some advanced features of IPv6, as follows:

- IPv6 nodes are expected to have multiple addresses as a matter of course and use a sophisticated source selection algorithm to choose among them when sending a packet. The addresses will have multiple scopes (interface local, unique IPv6 local unicast address, global), and renumbering is built into IPv6.
- Many nodes will have multiple interfaces, even nodes that are virtual tunnel interfaces, such as those used for multicast or mobility functions.
- The basis for the multihoming work that is being done today is that sites are expected to obtain their prefixes from their upstream providers; if they connect through multiple providers, they are expected to obtain a prefix from each one. Only packets sourced appropriately to the specific provider are allowed outbound via the provider network.
- Finally, you can choose to implement multihoming using BGP and provider-independent addressing in the same way as in IPv4 – by advertising the same prefix to both service providers.

https://t.me/learningnets

Site Multihoming Issues

This subtopic describes various site multihoming issues.

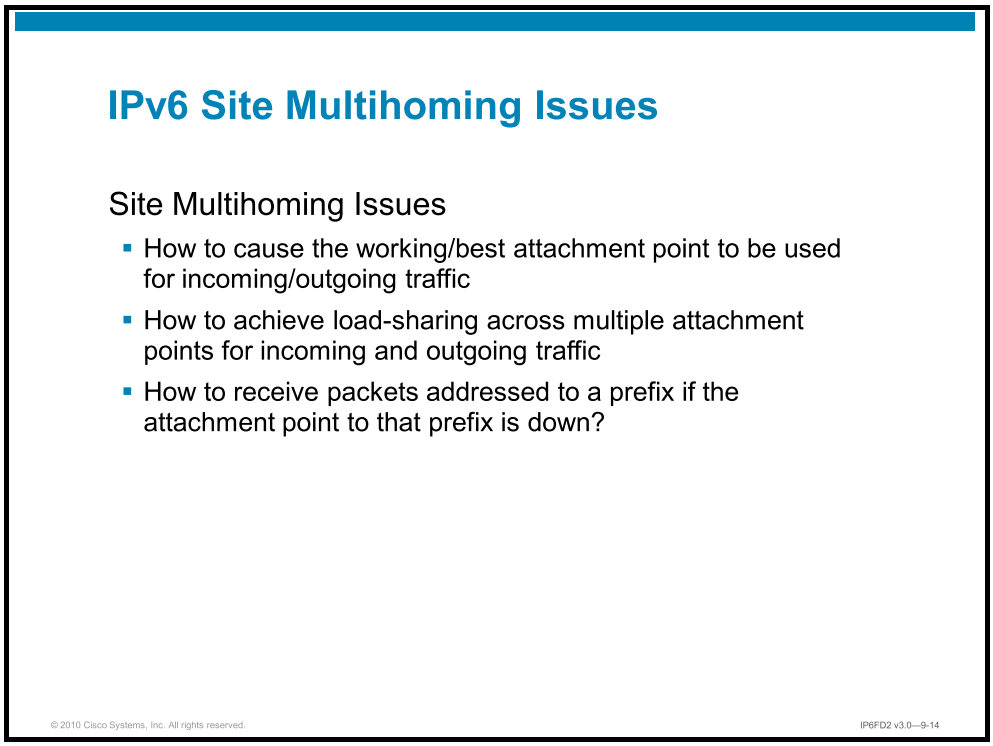

This is a high-level list of site multihoming issues.

- How to cause the working or best attachment point to be used for incoming or outgoing traffic
  - How do hosts know how to choose their source addresses so that their packets leave the enterprise via the best path? If hosts do not have to choose source addresses, how can the enterprise edge devices make source address choice transparent to the internal hosts, without using Network Address Translation (NAT)?
  - How does a local problem with the upstream ISP for a destination affect traffic to the site? If ISP A is down and IPS B is up, how do senders obtain that information to address packets appropriately?
- How to achieve load-sharing across multiple attachment points for incoming and outgoing traffic
  - How is this accomplished for simple load-sharing if both ISPs are active?
- How to receive packets addressed to a prefix if the attachment point to that prefix is down
  - If packets are delivered to ISP A when the attachment of ISP A to the enterprise is down, how can the packets be passed to the enterprise via ISP B?

Note There are many multihoming issues. This is not a complete list; it only includes the most important issues.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

IPv6 Multihoming Status

This topic describes the status of IPv6 multihoming solutions.

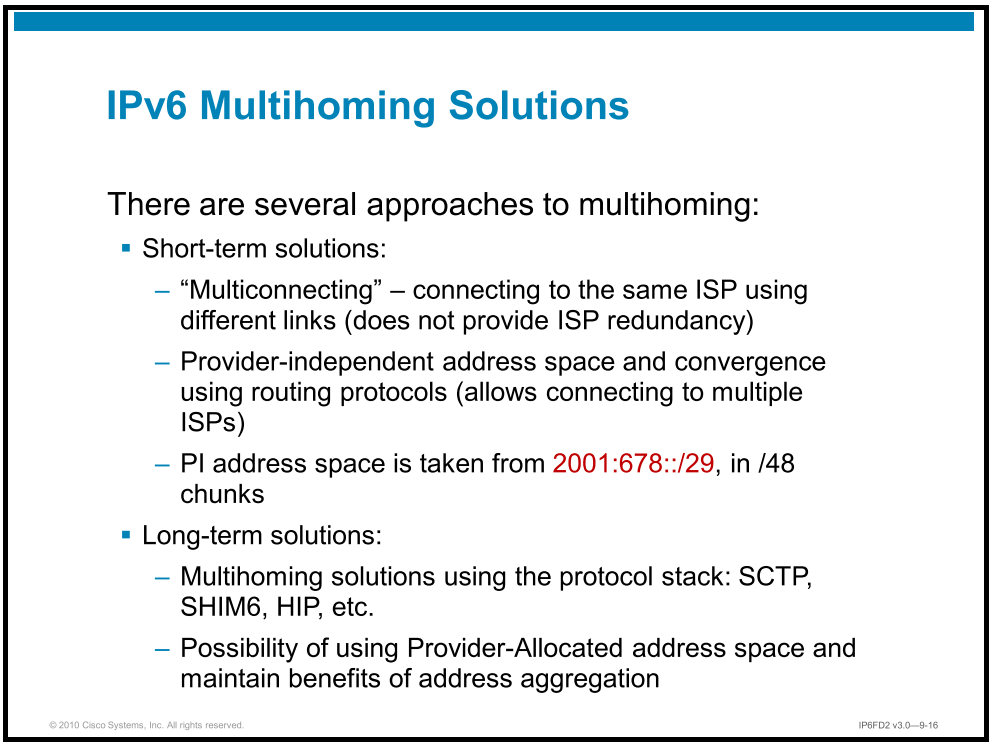

Viable Multihoming Mechanisms

First, it is important to consider and analyze different solutions to see which mechanisms could be viable and determine how they could be positioned. These approaches are grouped into three categories: immediate, short-term, and long-term. The definitions of the latter two are intentionally vague, but short-term solutions should not take more than one to three years to implement and deploy. These different categories of approaches are as follows:

- **Immediate approaches:** Multi-connecting seems to be the only obvious way to work around multihoming issues. Multi-connecting means connecting multiple times to a single ISP. Multi-connecting is *not* generally considered a multihoming mechanism. Multi-connecting is typically achieved by having multiple site border routers and connecting each of them to separate routers at the ISP, usually in different locations. Host-centric multihoming and multihoming at site exit routers would also be applicable (to an extent) immediately, if no ISP implements ingress filtering.
- **Short-term approaches:** These short-term approaches are under development:
  - Host-centric IPv6 multihoming and multihoming at site exit routers (fully fleshed out) are both short-term solutions that need to be further developed.
  - Provider-independent addressing based on autonomous system numbers is likely to be a short-term solution, based on routing table growth estimates. The same applies to advertising more specific routes, if done from specific allocations to make them distinguishable. Provider-independent address space is addressed by the American Registry for Internet Numbers (ARIN).

https://t.me/learningnets

Note	Provider-independent address space is taken from the **2001:678::/29** IPv6 address range, and is given in subnets of /48. There are concerns about the disadvantages of this approach, since the global routing table will become populated and saturated with such prefixes over time. On the other hand, BGP is infinitely scalable, limited only with physical memory and bandwidth, which is becoming cheaper through time.

- **Long-term approaches:** All transport solutions are only generally applicable in the long term, if at all. The following new identifier and locator separation (“two-space”) solutions are long-term solutions:
  - **Site Multihoming by IPv6 Intermediation (SHIM6):** SHIM6 is an IPv6-based site multihoming solution that inserts a new sublayer (shim) into the IP stack of end-system hosts. It will enable hosts on multihomed sites to use a set of provider-assigned IP address prefixes and switch between them without upsetting transport protocols or applications. While it is still in draft form, SHIM6 is currently the leading site multihoming proposal.
  - **Locator Identifier Separation Protocol (LISP):** The basic idea behind the Loc/ID split is that the current Internet routing and addressing architecture combines two functions: Routing Locators (RLOCs), which describe how a device is attached to the network, and Endpoint Identifiers (EIDs), which define "who" the device is, in a single numbering space, the IP address.

Routing-Based Solutions

RFC 3178, *IPv6 Multihoming Support at Site Exit Routers*, describes a mechanism for basic IPv6 multihoming support and its operational requirements. Unlike IPv4 multihoming, this mechanism does not affect the worldwide routing table size or IGP routing table size in upstream ISPs. The mechanism can be combined with more sophisticated (or complex) multihoming support mechanisms, and it can be used as a foundation for other mechanisms.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

PA-Based Solution Example

This example depicts a solution by advertising prefixes from the provider-aggregatable address space to both providers.

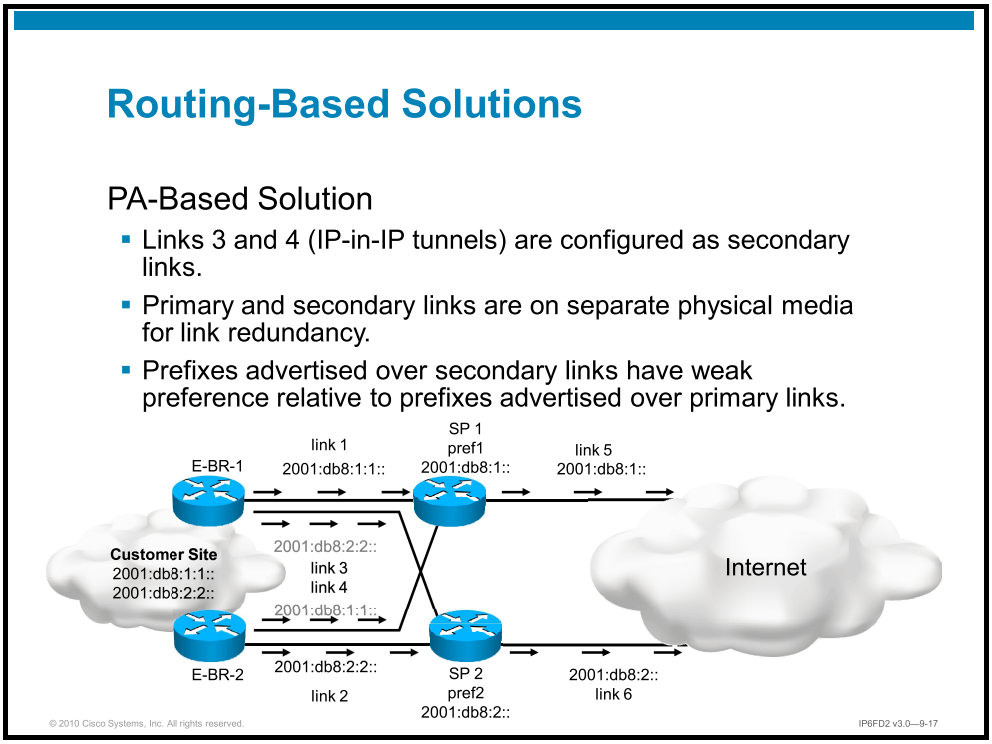

This example shows a solution with provider-aggregatable prefixes only. Prefixes are advertised to both providers, and path selection is manipulated by setting preference.

Note	Each provider must agree to the exception that it will accept advertised prefixes from the range of the other provider.

The site is connected to two ISPs, SP1 and SP2. IP address prefixes are assigned, pref1 and pref2 from SP1 and SP2, respectively. Hosts near SP1 will get an address from pref1 and hosts near SP2 will get an address from pref2. Within the site, routes are exchanged locally for pref1 and pref2, so that hosts in the site can communicate with each other without using an external link. SP1 and the customer site are connected by a primary link between the ISP router at SP1 and router E-BR-1—the customer router facing SP1. SP2 and the customer site are connected by a primary link between the SP2 edge router and router E-BR-2—the customer edge router facing SP2.

A secondary link is established between SP1 and E-BR-2, and between SP2 and E-BR-1, respectively. The secondary link usually is an IPv6-in-IPv6 tunnel. Under this configuration, both inbound and outbound packets can survive link failure on either side. Routing information with weak preference will be available as backup for both inbound and outbound cases.

There are several advantages and disadvantages associated with this example. The most serious disadvantage is that no fault tolerance is provided if an ISP has problems. The only failure mode that results in proper ongoing network function is a problem inside the enterprise at one of the edge routers.

https://t.me/learningnets

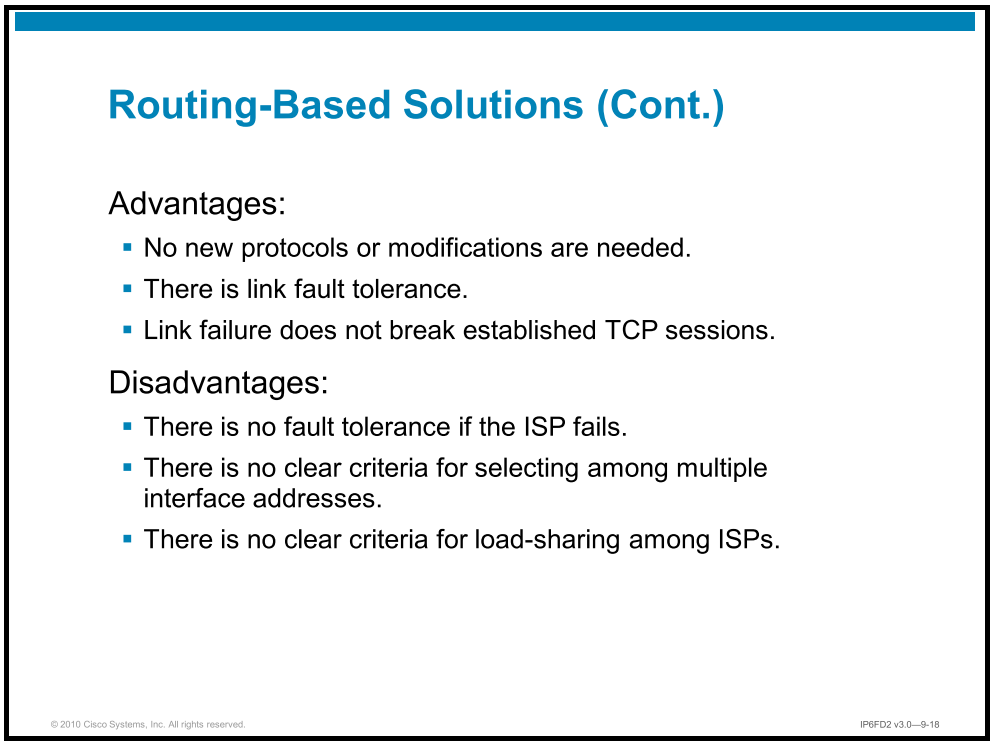

Here are the advantages associated with this example:

- No new protocols or modifications are needed.
- There is link fault tolerance.
- Link failure does not break established TCP sessions.

Here are the disadvantages associated with this example:

- There is no fault tolerance if the ISP fails.
- There are no clear criteria for selecting among multiple interface addresses.
- There are no clear criteria for load-sharing among ISPs.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Routing Example

This subtopic shows a routing example of IPv6 multihoming status.

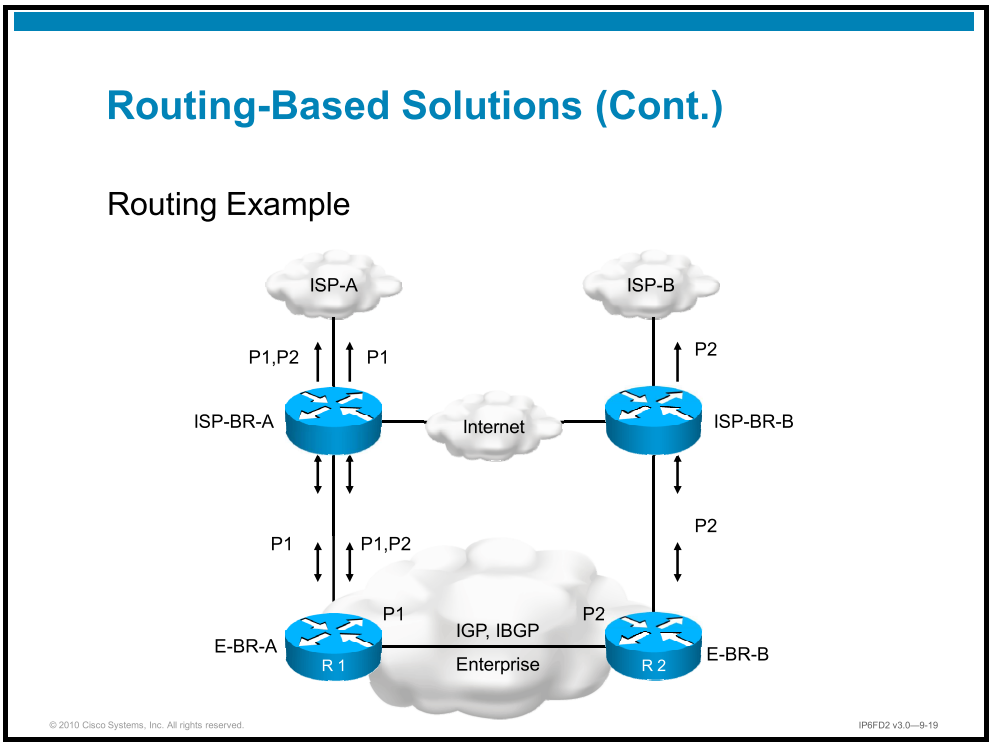

The figure shows an enterprise connected to two ISPs, ISP-A and ISP-B, as follows:

- The enterprise border router connects the enterprise to ISP-A as E-BR-A, and the ISP-A border router is connected to E-BR-A as ISP-BR-A. The enterprise border router connects the enterprise to ISP-B as E-BR-B, and the ISP-B border router is connected to E-BR-B as ISP-BR-B.
- P1 is the address prefix that ISP-A allocated to the enterprise, and P2 is the address prefix that ISP-B allocated to the enterprise.
- E-BR-A maintains direct external Border Gateway Protocol (eBGP) peering with ISP-BR-A and advertises reachability to P1 over that peering. E-BR-A also maintains an indirect eBGP peering with ISP-BR-B and advertises reachability to P2 over that peering.
- E-BR-B maintains direct eBGP peering with ISP-BR-B and advertises reachability to P2 over that peering. E-BR-B also maintains an indirect eBGP peering with ISP-BR-A and advertises reachability to P1 over that peering.

When connectivity between the enterprise and both of its ISPs is up, traffic destined to hosts whose addresses were assigned out of P1 would flow through ISP-A to ISP-BR-A to E-BR-A and then into the enterprise. Likewise, traffic destined to hosts whose addresses were assigned out of P2 would flow through ISP-B to ISP-BR-B to E-BR-B and then into the enterprise. Now, consider what would happen when connectivity between ISP-BR-B and E-BR-B goes down. In that case, traffic to hosts whose addresses were assigned out of P1 would be handled as before. But traffic to hosts whose addresses were assigned out of P2 would flow through ISP-B to ISP-BR-B; ISP-BR-B would encapsulate this traffic and send it to E- BR-A, where the traffic would get decapsulated and then be sent into the enterprise.

https://t.me/learningnets

In this scheme, there is no additional routing information due to multihomed enterprises that have to be carried in the “default-free” zone of the Internet. In addition, this scheme does not degrade in the presence of ISPs that filter out routes based on the length of their address prefixes. Note that the set of routers within an ISP that maintains indirect peering with the border routers within an enterprise does not have to be restricted to the border routers of the ISP that have direct peering with the border routers of the enterprise. The indirect peering could be maintained with any router within the ISP. This could improve the overall robustness in the presence of failures within the ISP.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

PI-Based Solution Example

This example depicts a solution by advertising prefixes from the provider-independent address space to both providers (“IPv4 fashion”).

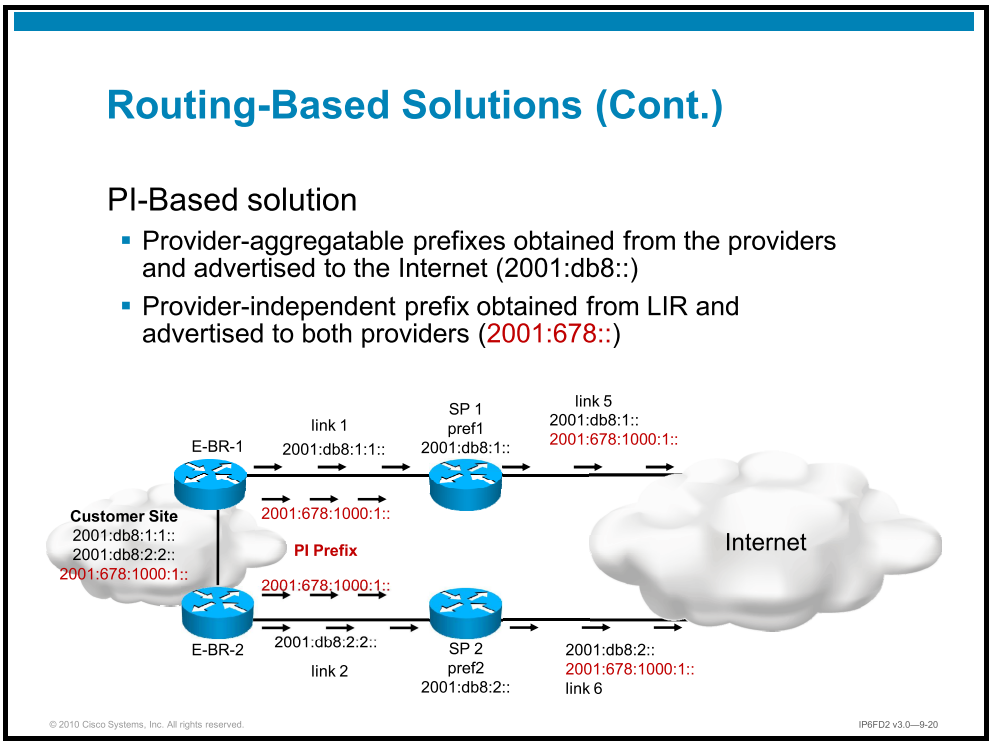

In this example, the customer is using provider-independent addressing.

The edge routers (E-BR-1 and E-BR-2) maintain external BGP sessions with SP1 / SP2 routers, and an internal BGP session between each other. This is done in the same way as multi-homing in IPv4.

Both providers advertise the common PI prefix to the Internet, along with a PA prefix from a particular provider assignment range.

The devices that need to be accessible to the clients from the Internet are additionally configured with IPv6 addresses using the PI prefix, along with the PA prefix. A web server in a DMZ can have three (3) IPv6 addresses: 2001:db8:1:1::80 [first ISP], 2001:db8:2:2::80 [second ISP], and 2001:678:1000:1::80 [PI address].

This approach is available since 2009 (proposed in 2006), and will be the most common as it is the easiest to implement, in the same way as multihoming is currently implemented in IPv4 networks.

There are concerns about routing table overpopulation with PI prefixes, however, currently there will be no more than the current IPv4 prefixes. Memory and bandwidth come at a lower price through time, which should address long-term scalability concerns. Another issue is that RIRs control the allocation of PI prefixes in a more strict way than PA prefixes.

https://t.me/learningnets

Protocol-Based Solutions

This topic describes multihoming solutions that are embedded in the protocol stack, and not using routing techniques.

SHIM6

This subtopic describes Site Multihoming by IPv6 Intermediation (SHIM6).

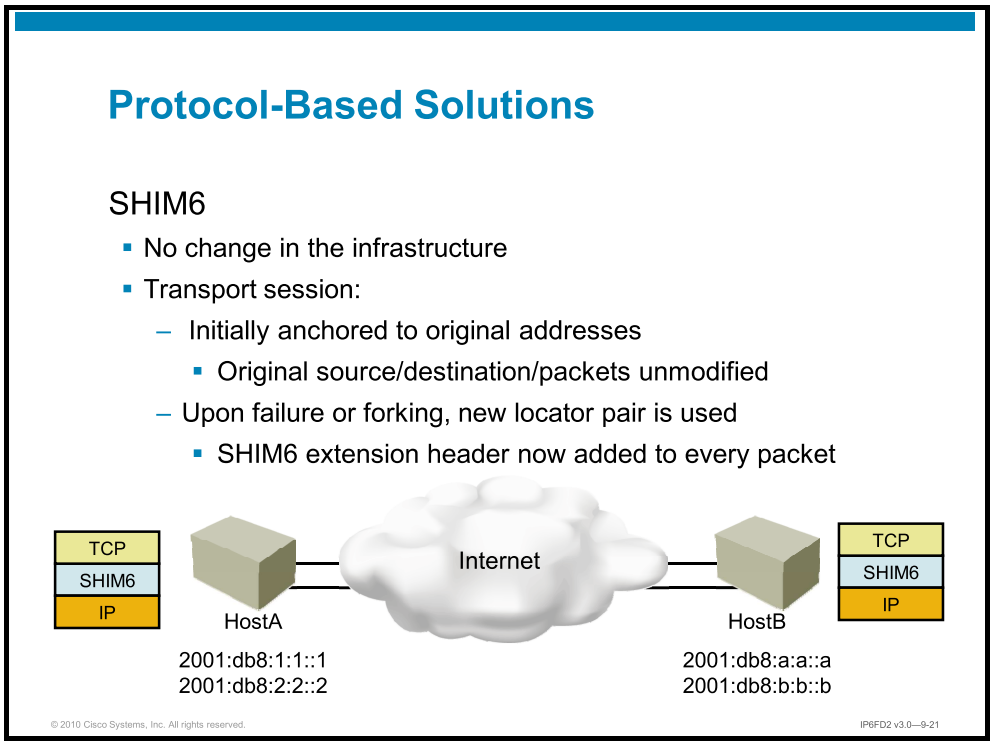

SHIM6 provides locator agility via a Layer 3 SHIM protocol, below the transport layer and protocols, so that multihoming can be achieved for IPv6 nodes, providing failover and load-sharing properties without assuming that a multihomed site will have provider-independent IPv6 addresses. A locator is a 128-bit value, which will initially be seen by the upper layer protocols as the IPv6 addresses used in the initial communication.

After a SHIM6-enabled, multihomed host (HostA) has established standard IPv6 communication with another host (HostB), which may or may not be multihomed itself, the local SHIM module on HostA will attempt to establish a shared state. A four-way exchange is attempted, which, if successful, results in HostA and HostB knowing a list of locators for each other. If the four-way exchange fails, the assumption is that HostB is not SHIM6-enabled and standard IPv6 communication proceeds.

Once a shared state is reached, communication continues without any overhead within the packets themselves, until something fails, at which point HostA or HostB probes their alternate locator pairs until a working pair is found. From this point on, the packets will have a SHIM6 extension header added to every packet, and while the upper-layer protocols still use the original IPv6 address pairs, the SHIM6 protocol will modify the packets to use the new locator pair.

Additionally, SHIM6 supports the notion of "context forking," by which an application will be able to choose a specific locator pair to be used. Implementation of this would be via an API which is not currently defined.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Note	For more information on SHIM6, see the documents on http://www.shim6.org.

LISP

This subtopic describes Locator/ID Separation Protocol (LISP).

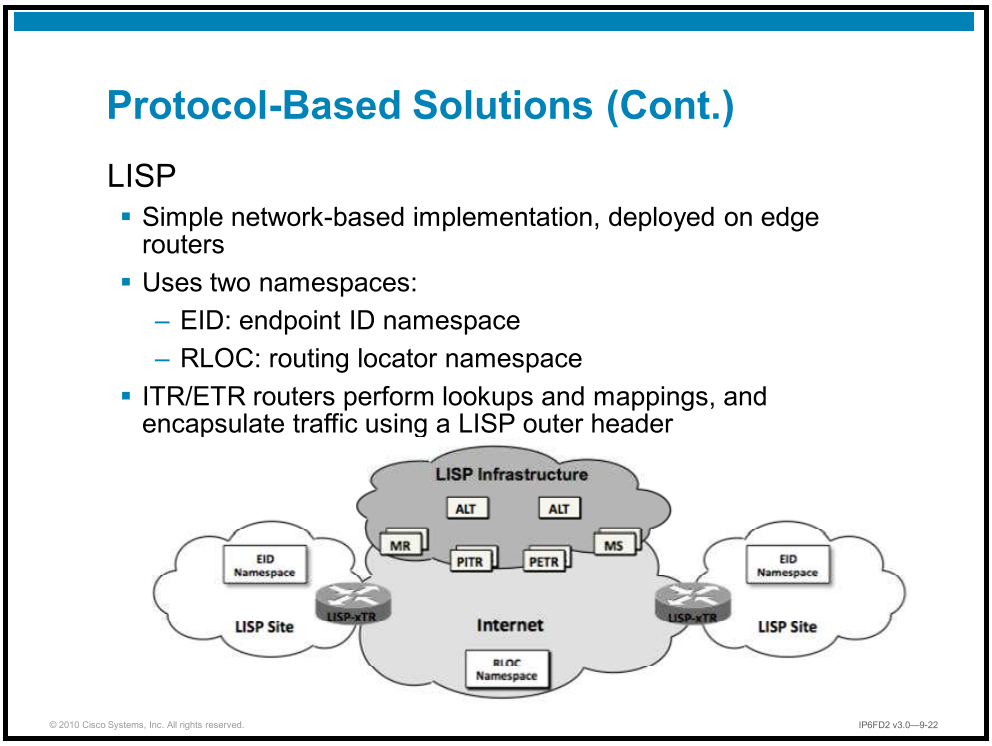

LISP is a network-based protocol that enables separation of IP addresses into two new numbering spaces: Endpoint Identifiers (EIDs) and Routing Locators (RLOCs). No changes are required to either host protocol stacks or to the "core" of the Internet infrastructure.

Routing Locators (RLOCs) are topologically assigned to network attachment points and are used for routing and forwarding of packets through the network. Endpoint Identifiers (EIDs) are assigned independently from the network topology and are used for numbering devices.

LISP then defines functions for mapping between the two numbering spaces and for encapsulating traffic originated by devices using non-routeable EIDs for transport across a network infrastructure that routes and forwards using RLOCs.

Note	Both RLOCs and EIDs are syntactically identical to IP addresses; it is the semantics of how they are used that is different.

Another key LISP concept is the "Tunnel Router". A tunnel router prepends LISP headers on host-originated packets and strip them prior to final delivery to their destination. The IP addresses in this "outer header" are RLOCs. During end-to-end packet exchange between two Internet hosts, an ingress tunnel router prepends a new LISP header to each packet and an egress tunnel router strips the new header. The ITR performs EID-to-RLOC lookups to determine the routing path to the ETR, which has the RLOC as one of its IP addresses.

https://t.me/learningnets

For example, when the client host1.abc.com wants to communicate with the server host2.xyz.com:

1. host1.abc.com wants to open a TCP connection to host2.xyz.com.
   It does a DNS lookup on host2.xyz.com. An A/AAAA record is returned. This address is the destination EID. The locally assigned address of host1.abc.com is used as the source EID. An IPv6 packet is built and forwarded through the LISP site as a normal IP packet until it reaches a LISP ingress tunnel router.

2. The LISP ITR must be able to map the EID destination to an RLOC of one of the egress tunnel routers at the destination site.

3. The egress tunnel router, when getting the packet, looks at the destination EID of the Map-Request and matches it against the prefixes in the ETR's configured EID-to-RLOC mapping database.
   This is the list of EID-prefixes the ETR is supporting for the site it resides in. If there is no match, the Map-Request is dropped. Otherwise, a LISP Map-Reply is returned to the ITR.

4. Subsequent packets from host1.abc.com to host2.xyz.com will have a LISP header prepended by the ITR using the appropriate RLOC as the LISP header destination address learned from the ETR.

5. The ETR receives these packets directly (since the destination address is one of its assigned IP addresses), strips the LISP header and forwards the packets to the attached destination host.

Note More information about LISP can be found at http://tools.ietf.org/wg/lisp/draft-ietf-lisp/. For Cisco-specific information please refer to the "Cisco LISP Configuration Guide".

LISP is supported in Cisco IOS, IOS-XE, and NX-OS.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

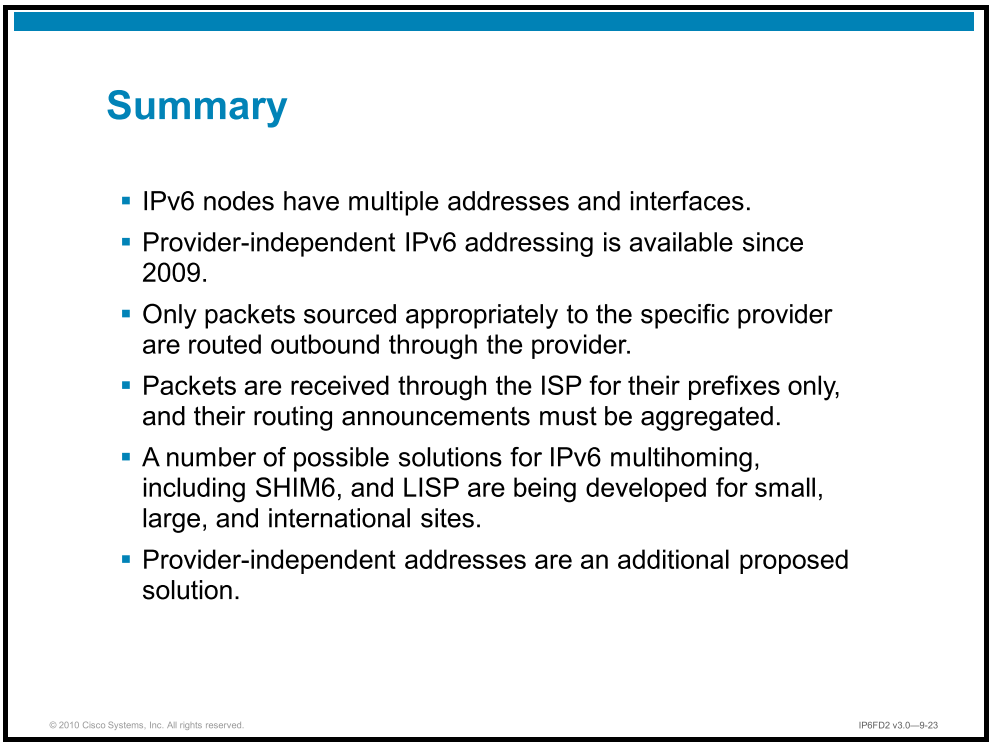

Resources

To learn more about IPv6 addressing and multihoming issues, refer to the following documents:

- **Address Space Managed by the RIPE NCC**
  https://www.ripe.net/ripe/docs/ripe-ncc-managed-address-space.html#2e
- **Understanding IP Addressing**
  http://www.ripe.net/info/info-services/addressing.html
- **IPv6 Multihoming Summary**
  http://en.wikipedia.org/wiki/Multihoming#IPv6 multihoming
- **SHIM6 Multihoming**
  http://www.shim6.org/
- **LISP**
  http://tools.ietf.org/wg/lisp/draft-ietf-lisp/
- **Cisco LISP Configuration Guide**
  http://www.cisco.com/en/US/docs/ios/lisp/configuration/guide/LISP configuration guide.pdf

https://t.me/learningnets

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3

Identifying IPv6 Enterprise Deployment Strategies

Overview

No single pathway exists for deploying IP version 6 (IPv6) in an enterprise environment. Differing requirements for security, routing, management, and infrastructure upgrades will dictate various approaches to IPv6 integration. Understanding the pros and cons of different transition mechanisms will help you develop successful transition plans. This lesson describes several IPv6 enterprise deployment strategies.

Objectives

Upon completing this lesson, you will be able to identify IPv6 enterprise integration strategies. This knowledge includes being able to meet these objectives:

- Define the enterprise network
- Identify specific concerns with campus networks that use IPv6
- Identify specific concerns with WAN networks that use IPv6
- Identify the advantages and disadvantages of dual-stack integration in IPv6 enterprise networks
- Identify the advantages and disadvantages of using encapsulation in IPv6 enterprise networks
- Identify the advantages and disadvantages of translation in IPv6 enterprise networks

https://t.me/learningnets

Enterprise Networks

This topic describes specific concerns about enterprise networks that use IPv6.

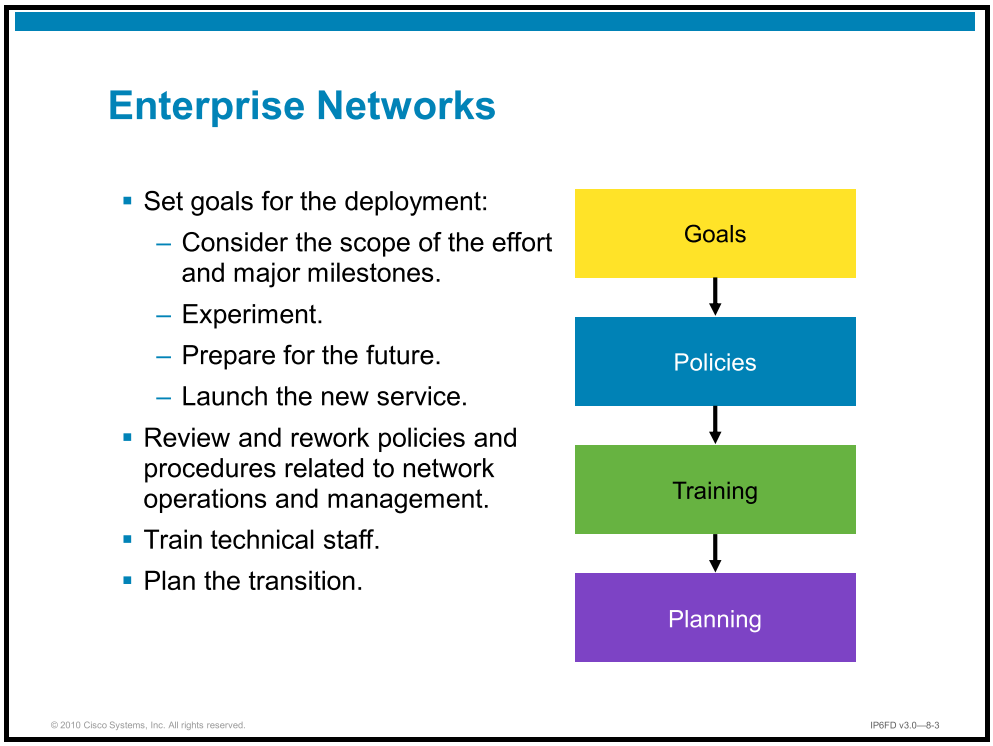

When you design an IPv6 enterprise network, keep in mind the specific needs and requirements of the company by considering the following points:

- Consider the goals for the deployment, in terms of outcome and process, and choose major milestones.
- Review existing network policies to determine how IPv6 will affect them. What policy constraints can be lifted, and what new policies are needed?
- Any new technology project must consider the skills and skill-development needs of the people responsible for the job. Consider the staffing plan needed to get the work done, and begin preparing early.
- Most deployments will happen on existing networks, either as end-to-end upgrades or selected upgrades serving specific users. Any project that affects the core network must be planned carefully to avoid outages.

The three major deployment scenarios to consider are as follows:

- **Experimental use of IPv6:** In this case, an IPv6 network is built with a relatively small number of devices to evaluate and to learn the use of IPv6. Because this scenario is an experimental setup, it will use the existing network infrastructure as it is. The primary goal will be to gain experience with IPv6 in a safe environment.
- **Preparation for future use of IPv6:** In this case, a corporate network infrastructure is IPv6-enabled for future applications. IPv6 will be operating side-by-side with IPv4. Because this scenario introduces a production IPv6 on the live enterprise network, it must be carefully coordinated.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

- **Building an IPv6 network for specific, advanced new services:** New business systems are built with IPv6. Such a network can be separate from the current IPv4 network. A Virtual Private Network (VPN) may be used to build an overlay network to support certain advanced services (such as multicast).

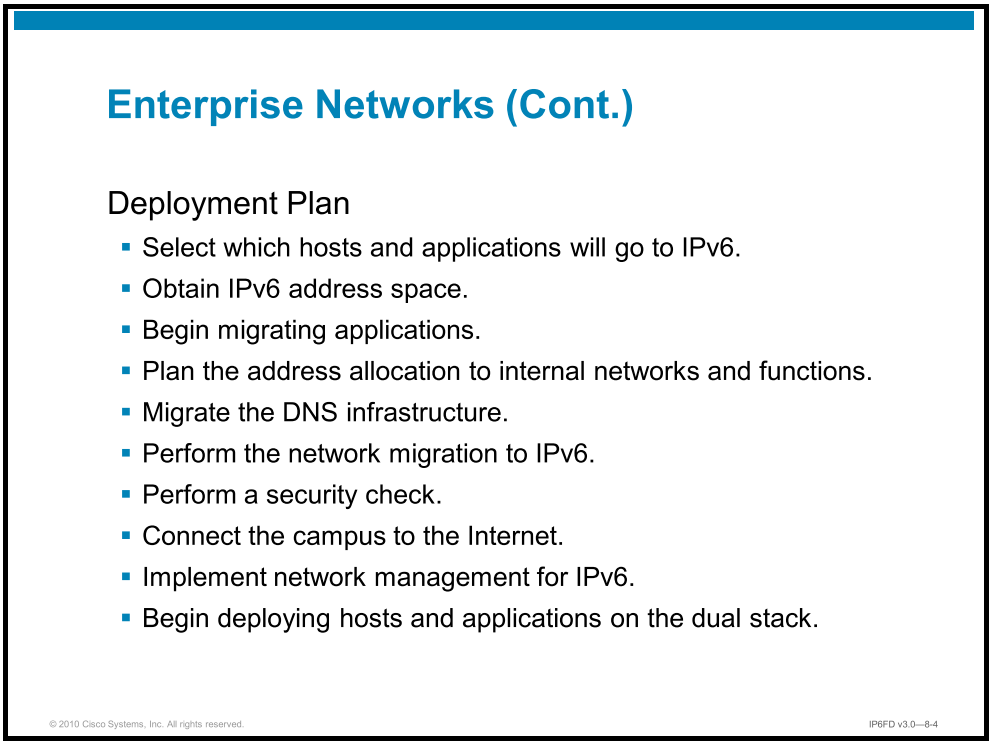

The goal of the enterprise deployment strategy is to plan ahead to minimize downtime on existing networks and to ensure a smooth transition.

The enterprise deployment strategy for most campus networks will follow these basic steps (Some enterprises may perform these steps in a different order.):

Step 1 Select which hosts and applications will go to IPv6.

Coordinate the network and application transition, and consider host operating systems and support systems.

Step 2 Obtain an IPv6 address space.

Step 3 Complete the Domain Name System (DNS) infrastructure migration.

This migration is a key step to complete early on because it consists of updating your DNS environment to support using IPv6 records in addition to IPv4 records. Transport can remain on IPv4, but IPv6 transport is now available on most platforms.

Step 4 Perform the network migration or IPv6-enable the existing IPv4 network.

This is the primary task during the migration and can be completed in small stages, as long as all production hosts remain on IPv4 and services are not interrupted.

Step 5 Perform a security check.

Make sure that site security is ready before connecting the IPv6 portion of the network to the Internet.

https://t.me/learningnets

Step 6 Connect the campus to the Internet.

Connect via the ISP from which you obtained your /48 prefix allocation. Complete performance and security testing.

Step 7 Implement network management for IPv6.

Update existing tools or obtain new tools to monitor the IPv6 status of your network. Management traffic can remain on IPv4, if needed, for as long as the network remains dual-stacked.

Step 8 Begin deploying hosts and applications on the dual stack.

Begin migrating hosts and applications to dual stack, or to IPv6, if this network is a new IPv6-only network.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

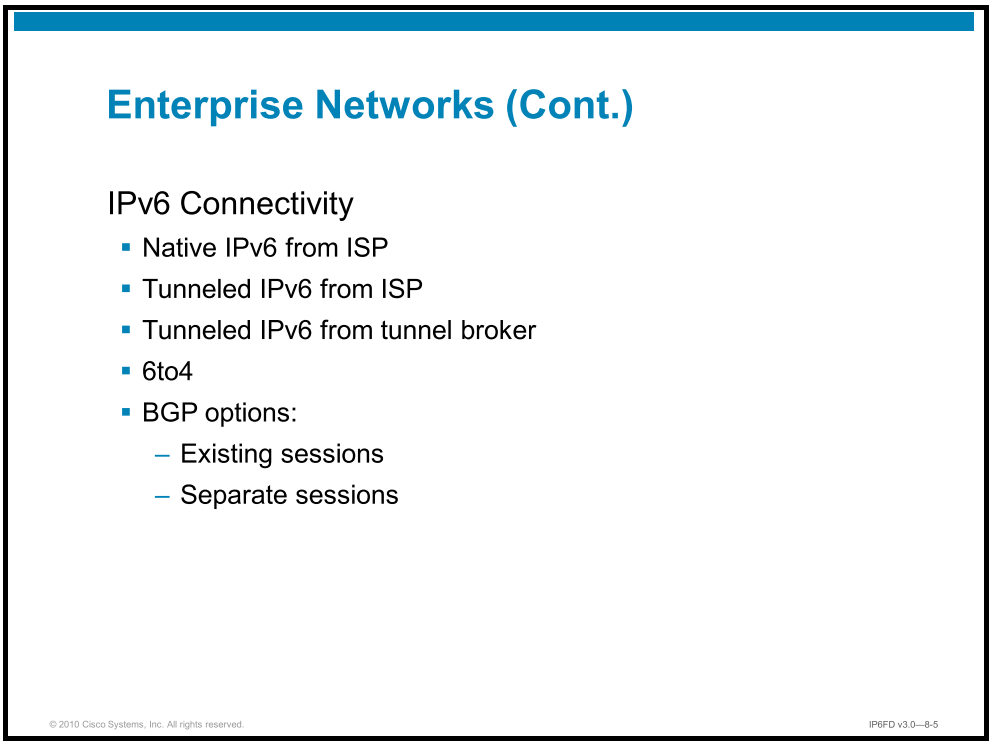

The best option for Internet connectivity is native IPv6, which can run either over the same physical link as IPv4 or a separate link.

If native IPv6 external connectivity is not available, use a transition mechanism. Here are the choices:

- **Set up an IPv6 tunnel to the ISP:** Sometimes the ISP cannot offer a native IPv6 link, but can offer a tunneling service. In this case, you have to use encapsulation but your Internet exit point is the same service provider for both protocols.
- **Set up an IPv6 tunnel to a virtual ISP:** If your ISP does not offer IPv6 connectivity, you need to obtain access from a different service provider. In this case, the tunnel is built over the IPv4 Internet to reach the remote IPv6 tunnel endpoint. The drawback to this approach can be considerable latency of the IPv6 tunnel if the virtual ISP is geographically distant to your network.
- **Use 6to4:** This transition mechanism is automatic, but requires a relay to the native IPv6 world. In addition to the security considerations for automatic tunneling, you cannot use assigned IPv6 address space with 6to4. You are restricted to the 6to4 2002::/16 prefix. Another point against 6to4 is dependence on public 6to4 relays.
- **Border Gateway Protocol (BGP):** If you are a multihomed enterprise network, you are required to use BGP. In a mixed IPv4 and IPv6 environment, you need to advertise and receive both types of routing information. You can receive this information over one BGP neighbor-relationship session or over two independent sessions—one for each protocol—as described in the routing section.

https://t.me/learningnets

Security of IPv6 Tunnels

Comprehensive network security is required for any enterprise or organizational network, whether the network is IPv4-based, IPv6-based, or both. When using transition mechanisms, you must control the external IPv6 entry points of your enterprise network, like external IPv4 entry points are controlled. Potential security threats have been identified and documented with mechanisms such as IPv4-mapped addresses or 6to4, but you can use IPv4 security mechanisms to protect these tunnels. For example, IPv4 IPsec can be configured to protect any IPv6-over-IPv4 tunnels.

You must control IPv6 traffic crossing your IPv4 firewall or edge-security router device. You may enable protocol 41 (IPv6 in IP tunnels – manual tunnels, ISATAP, etc.) only between a pair of well-known IPv4 routers, inside and outside of the corporate network. Generally, IPv4 firewalls do not inspect the content of IP-over-IP traffic; this situation is also true in IPv6. Terminating IPv6 tunnels directly on an IPv6 firewall-capable device, either stateless or stateful, is essential for inspecting and protecting your IPv6 network.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

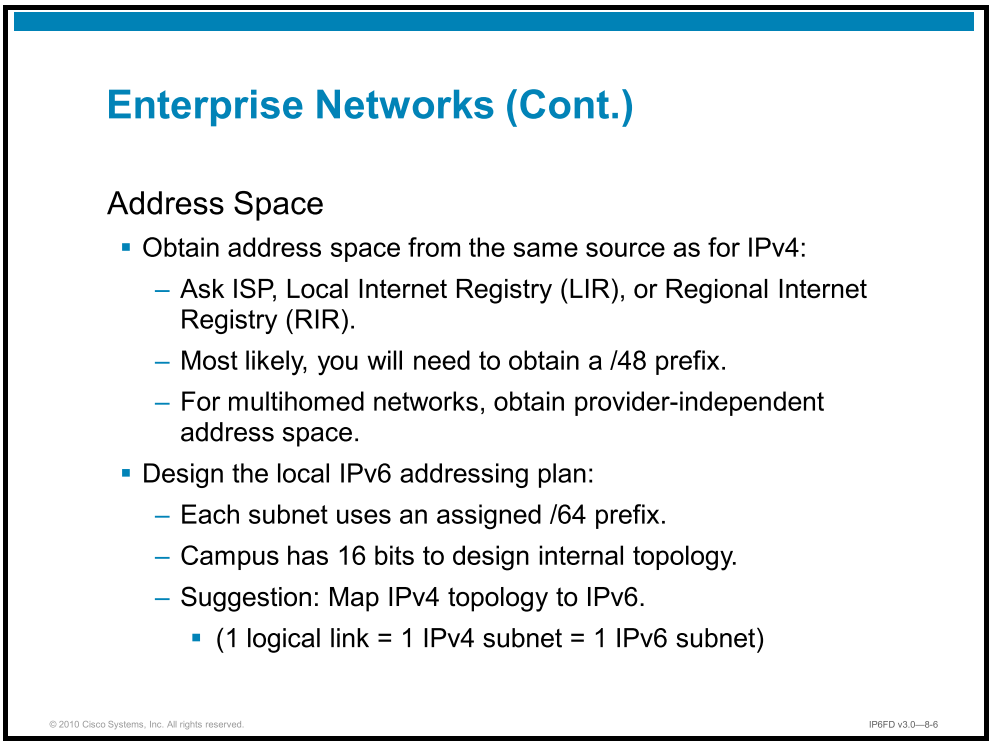

Addressing provides the foundation for the campus deployment. In most cases, because of IPv6 address aggregation policies, addresses must be obtained from your upstream ISP. If your enterprise is large enough, you may qualify as an ISP yourself (or for a provider-independent prefix) and obtain an allocation directly from the Internet registries. This type of situation is rare and requires a documented plan to have a certain number of customers within a given period. Most enterprises will likely receive a /48 prefix. Smaller enterprises may receive /52 or /56.

Multihomed networks can request provider-independent (PI) address space. In such cases, /48 is assigned.

All IPv6 subnets that use stateless autoconfiguration need a /64 prefix. In fact, the current recommendation is to use /64 networks even for point-to-point connections. With a /48 overall allocation, this setup allows 16 bits for internal network assignment or the ability to deploy roughly 64,000 /64 networks inside the enterprise.

The local addressing plan should map to the existing network, if one exists. Topologically, that will make overlaying the IPv6 network more intuitive and will ease transition to an all-IPv6 network.

Note For more information regarding addressing, see **RFC 5375: IPv6 Unicast Address Assignment Considerations**.

https://t.me/learningnets

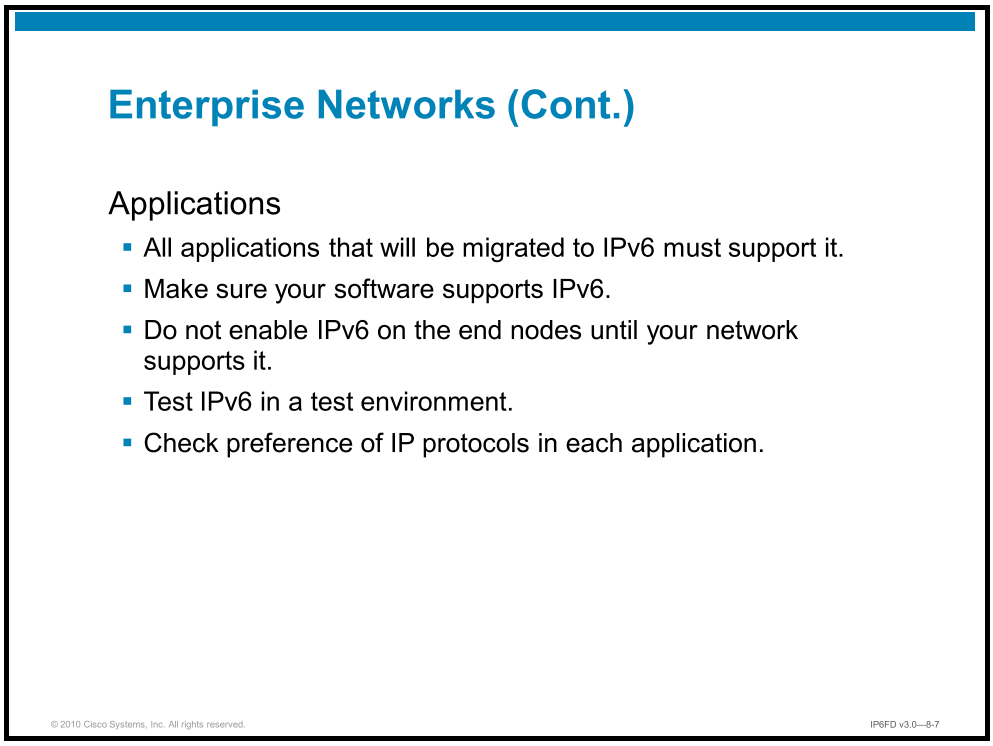

Migration of an application to support IPv6 is mandatory for successful deployment of IPv6. Not all applications need to support IPv6, as discussed later in this lesson, but it makes sense for the network to support at least one service.

Software departments should be given enough time to perform upgrades, so that all installations are IPv6-capable. On the other hand, IPv6 should not be enabled until the entire network is up—most importantly, the Internet access if end nodes will be using them. This step can easily be achieved by disabling IPv6 support on LAN interfaces on LAN gateways.

Before deploying IPv6 in a production environment, it is prudent to test the same setup in a test environment. Check the application software for protocol preference. While most applications will prefer IPv6 connections when both protocols are available, some applications will stick with IPv4. IPv6 support must be explicitly enabled.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

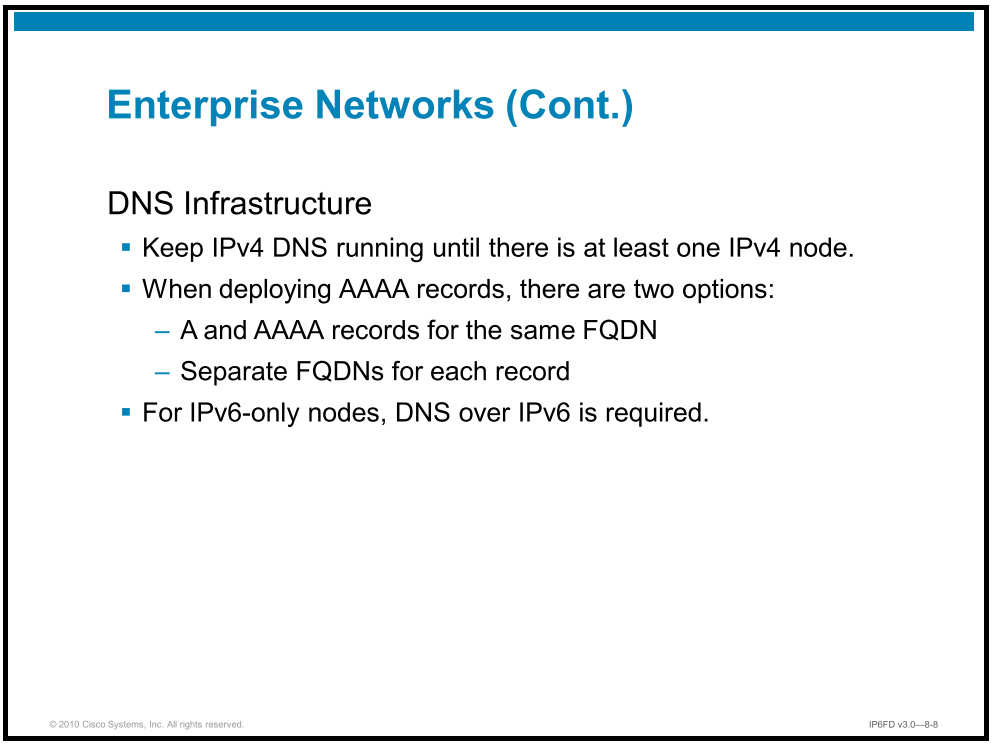

For DNS to be useful for resolving fully qualified domain names (FQDNs) to IPv6 addresses, your DNS software must support AAAA records. Moving entire DNS infrastructure transport solely to IPv6 is not advisable and not necessary.

For services available on both IPv6 and IPv4, there are two options for DNS records. The first option is to use the same FQDN for both protocols by setting two records, A and AAAA. In this case, for example, you have http://www.example.com, which has an A record pointing to 192.0.2.80 and an AAAA record pointing to 2001:db8::80. This works, as long as there is complete IPv6 and IPv4 connectivity between the server and any potential client. It may cause problems if both the server and the client have IPv6 enabled. The client prefers IPv6, but there is a connectivity issue on the path from the client to the server. In such a scenario, introducing IPv6 to the client may appear as a cause for a network slowdown. The client, preferring IPv6, will try to connect to the server by using IPv6, timeout and fail, and then fall back to IPv4, which would work. This scenario may be quite common and can cast a bad light on IPv6. Therefore, all applications on the server listed in the AAAA record must support IPv6.

To avoid such scenarios, some service providers may decide to use a different FQDN for the same service that is offered on IPv6. Possible choices are www6.example.com, w6.example.com, ipv6.example.com, and so forth. Using a different FQDN avoids the previous scenario. However, users wanting to connect over IPv6 must type a different hostname in the URL bar (for example, when using the web), which may be more cumbersome.

Leaving DNS service on IPv4 is not enough when using IPv6-only nodes. If that is the case, you must also enable DNS service on IPv6.

https://t.me/learningnets

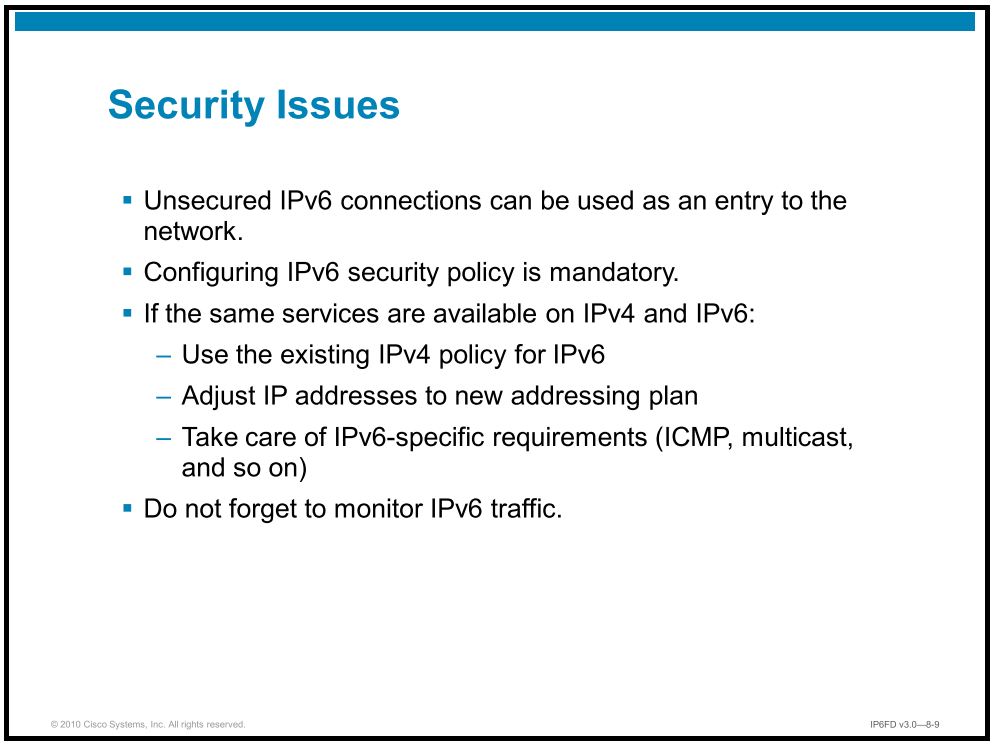

Introducing IPv6 in a secure network may create backdoors if IPv6 is not secured in the same way as IPv4. Existing IPv4 security policies will not be applied to IPv6 traffic. Instead, all IPv6 traffic will be passed unrestricted by default. Therefore, it is necessary to implement security measures before IPv6 is activated.

The recommended practice is to deploy the same policy for IPv6 as is implemented for IPv4. If some services are available for IPv4 only, this traffic type should be blocked in IPv6. If there is an IPv6-only service, then new rules must be written for IPv6 traffic.

IPv6 has some specific features not present in IPv4. When writing access rules for IPv6, take these features into account and permit specific traffic that is not present or mandatory in the IPv4 world. These features include: ICMP for neighbor discovery and router advertisements, and ICMP for Path Maximum Transmission Unit Discovery (PMTUD) and multicast traffic.

After IPv6 is deployed, configure your network monitoring systems to also monitor IPv6 traffic. (Some users may switch to IPv6 traffic, and their activity on the network may go unnoticed if monitoring systems report only IPv4 activity.)

https://t.me/learningnets

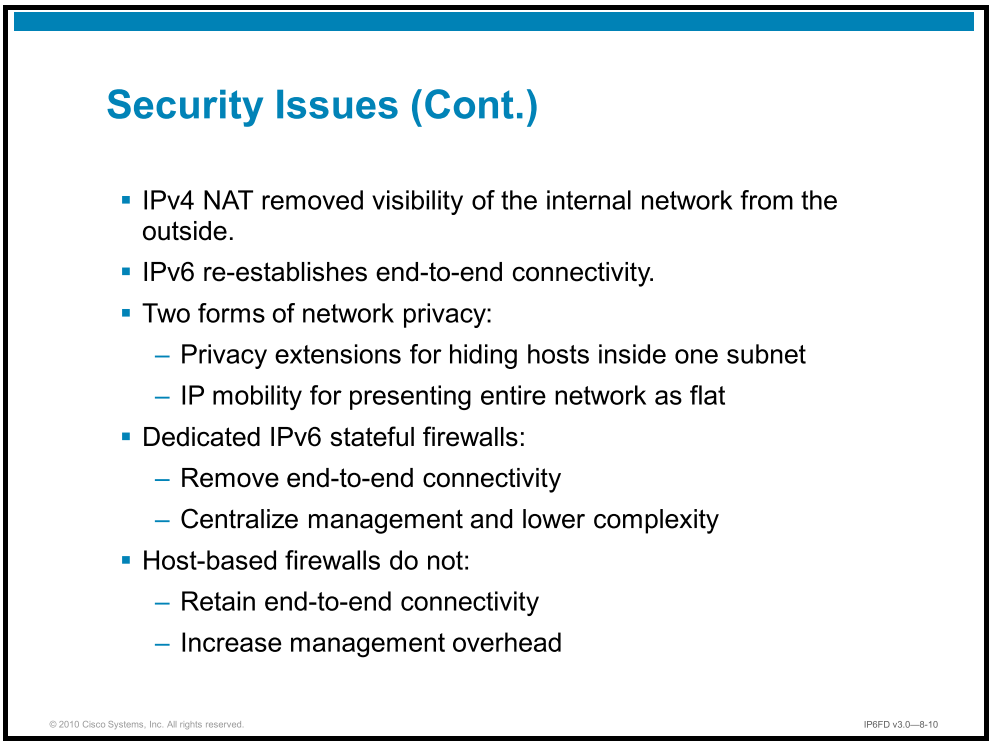

IPv4 Network Address Translation (NAT) hides the internal network from the outside, and Port Address Translation (PAT) must also keep track of sessions. As a result, IPv4 is often incorrectly used as a security measure. In an IPv6 world, there is no NAT, and such features as hiding and stateful connection checks are absent.

If end-to-end connectivity is required, then only host-based firewalls can be used. Otherwise, just like NAT, a stateful firewall will only allow return traffic to the protected nodes. Exotic protocols, like FTP, will fail to work unless specifically supported by the firewall device.

Network hiding can be performed, to some extent, by implementing IP mobility. This implementation may introduce additional overhead, which may, in turn, lower the overall security of the network. Careful analysis is required to weigh all options and make trade-offs.

https://t.me/learningnets

Impacts of Network Services

This topic describes the impact of network services on IPv6 enterprise networks.

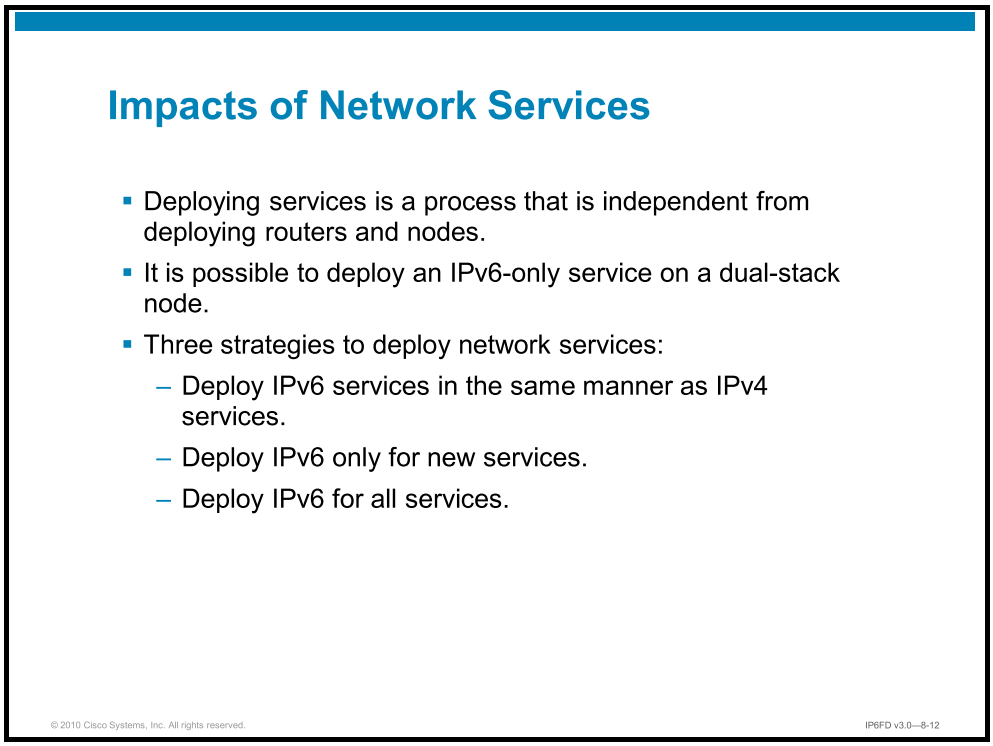

Applications must also be migrated to IPv6 or new applications must be developed. Complete this work in careful relationship to network and host (IP stack) upgrades. Coordinate the entire process (network, hosts, and applications) to yield the best results.

For existing IPv4-based applications, here are the three approaches to deploying IPv6:

- Upgrade current applications to be dual-stack, so that a single application provides services for both protocols. This upgrade is the cleanest solution, because there is only one application to maintain, and the functionality of the application is independent of the network transport.
- Introduce new applications running only on IPv6. Plan a quick transition from IPv4 nodes on the network and to retire legacy IPv4 applications in favor of improved IPv6 applications.
- Use only IPv6, which implies a conversion process for current applications to IPv6 and removing their existing IPv4 capability. Given the long transition that is expected for IPv4 and the number of important legacy applications most enterprises operate, this strategy will be practical only in very limited cases.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

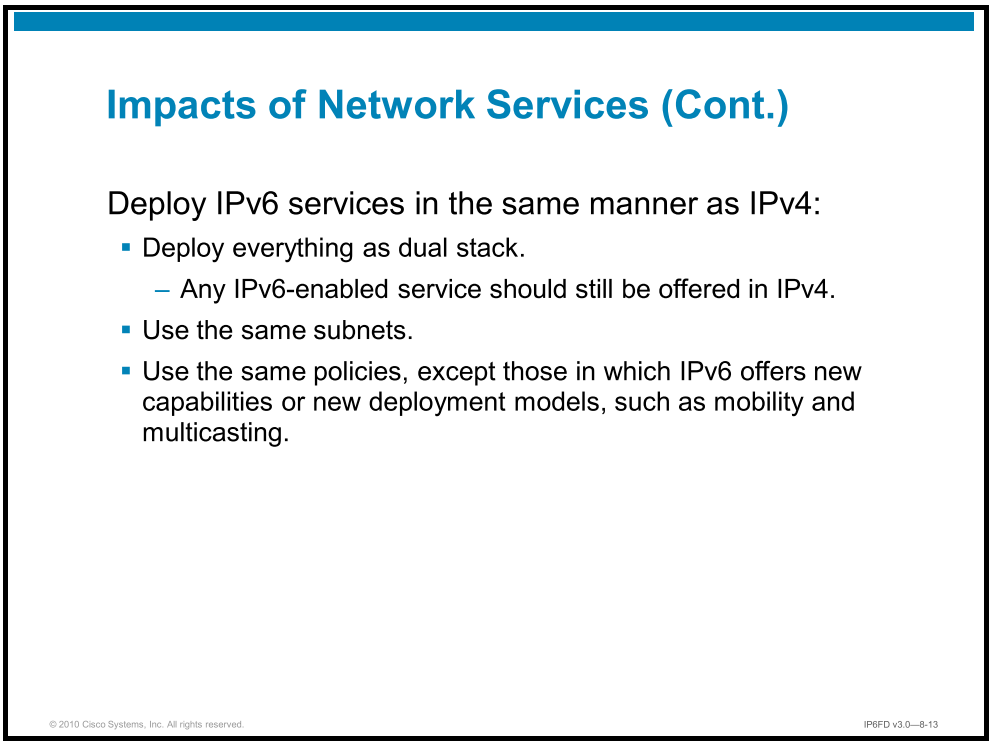

With this deployment strategy, both hosts and routers need to be upgraded to support IPv6 along with continued support for IPv4. This process is usually accomplished by enabling existing software features or by acquiring upgrades. The goal is to upgrade all operating systems, network systems, and services to provide IPv6 functionality in addition to an already existing IPv4 environment. This form of deployment may require dual routing tables with overlapping administration. Similar management and support are also required for both the IPv4 network and the IPv6 network (for example, security).

Dual stacking addresses most address-selection and DNS-resolution issues. A dual-stack node should observe DNS queries attempting to resolve an IPv6 address first. If the IPv6 address is not available, the DNS would use an IPv4 address. In some implementations, a dual-stacked node may request *any* IPv6 DNS record and process the decision of which address to use locally after receiving all the responses. Dual stacking allows hosts to continue to reach IPv4 resources while adding IPv6 functionality.

In this scenario, services throughout the network should map closely to the existing IPv4 environment. The same applications should be available to users, and the network infrastructure should be configured to map neatly to the existing network design. Use of the same high-level policies will also make coexistence more rational, but not at the expense of limiting new IPv6 capabilities.

https://t.me/learningnets

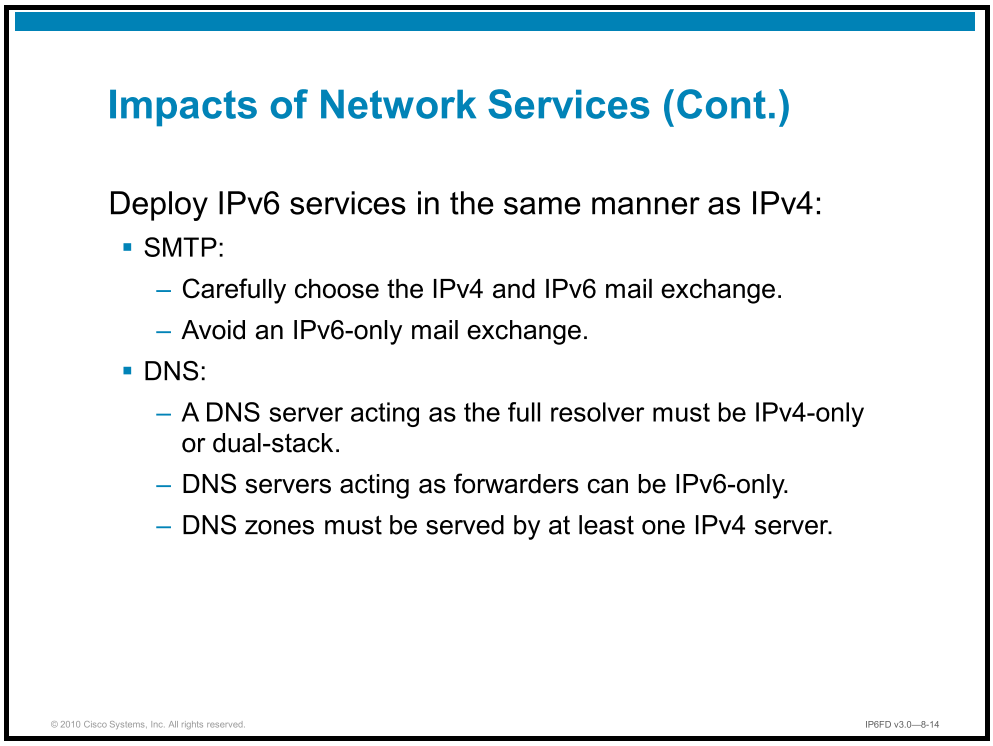

Email and DNS are critical applications for most enterprises. Implementations exist for both functions and must be coordinated with changes in network and host addressing. Simple Mail Transfer Protocol (SMTP) must be configured so that your email services work for the entire Internet-based SMTP community. Currently, the community is IPv4-based almost everywhere because this application inherently faces outward. Using an IPv6-only SMTP server will result in IPv4-only SMTP servers being unable to connect and deliver email. IPv6-capable nodes must not make too many assumptions about the IPv6 capability of other Internet-based systems.

For DNS systems, all interactions with upstream root DNS servers must be over IPv4. Local DNS operations, such as campus resolvers querying a local cache-only server, can use IPv6 transport.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

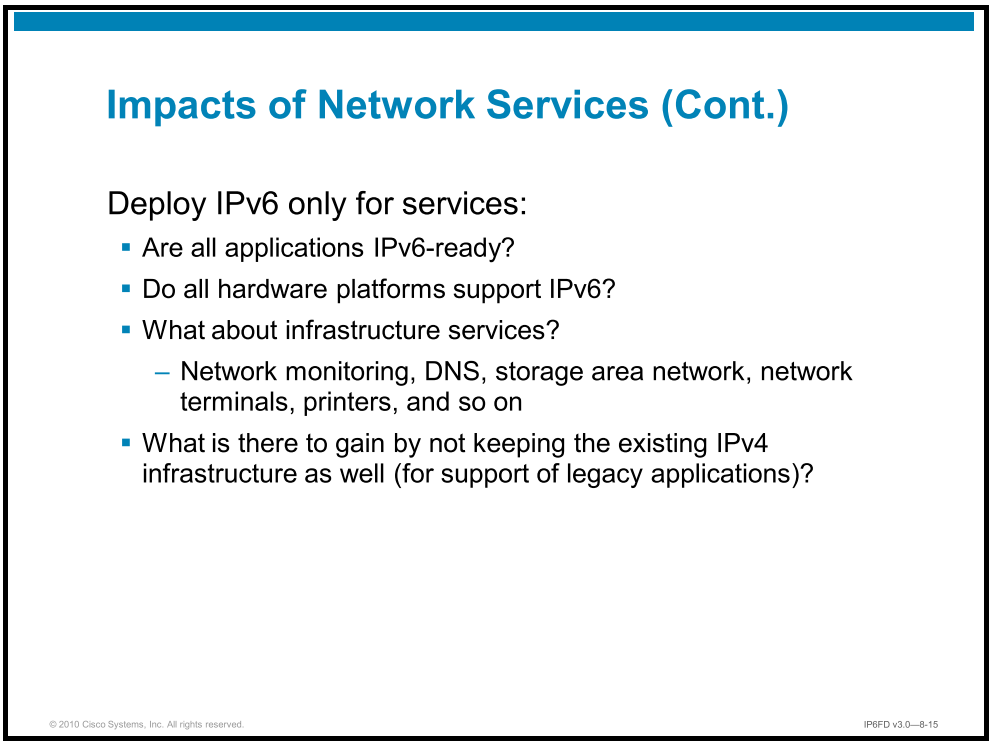

Business strategy will propel the technical planning. The technical community may be anxious to deploy an all-IPv6 network, but the timeline and strategy that remove the existing IPv4 capability should be needs-based. Before a complete changeover, all nodes, applications, network infrastructure, and network services must be all-IPv6 capable with plans for support and introduction in place. For most enterprises, a relatively quick build out of a dual-stack network infrastructure, followed by the gradual retirement of IPv4-only services and applications, will make the most sense.

https://t.me/learningnets

Dual Stack

Full Layer 3 infrastructure upgrade to implement dual stack:

- Software or hardware and software dependent on expected performance
- Requires IPv6 support on all components

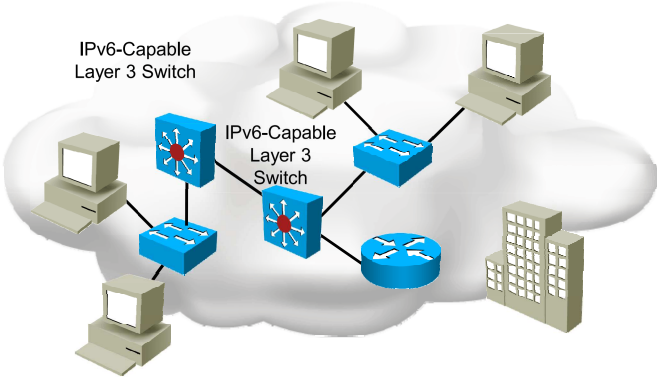

© 2010 Cisco Systems, Inc. All rights reserved.

There are two campus scenarios that are representative of common deployment strategies: IPv6-capable Layer 3 switches and IPv6-over-Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) tunnels. Each scenario has strengths and weaknesses. The two scenarios are described as follows:

- **IPv6-capable Layer 3 switches:** A native dual-stack or all-IPv6 deployment in which all Layer 3 devices in the campus are upgraded to run IPv6. This deployment can require hardware and Cisco IOS Software upgrades and configuration changes to enable IPv6. This scenario is the most direct but will probably require upgrading network equipment. This approach enables IPv6 unicast and, optionally, multicast services.

https://t.me/learningnets

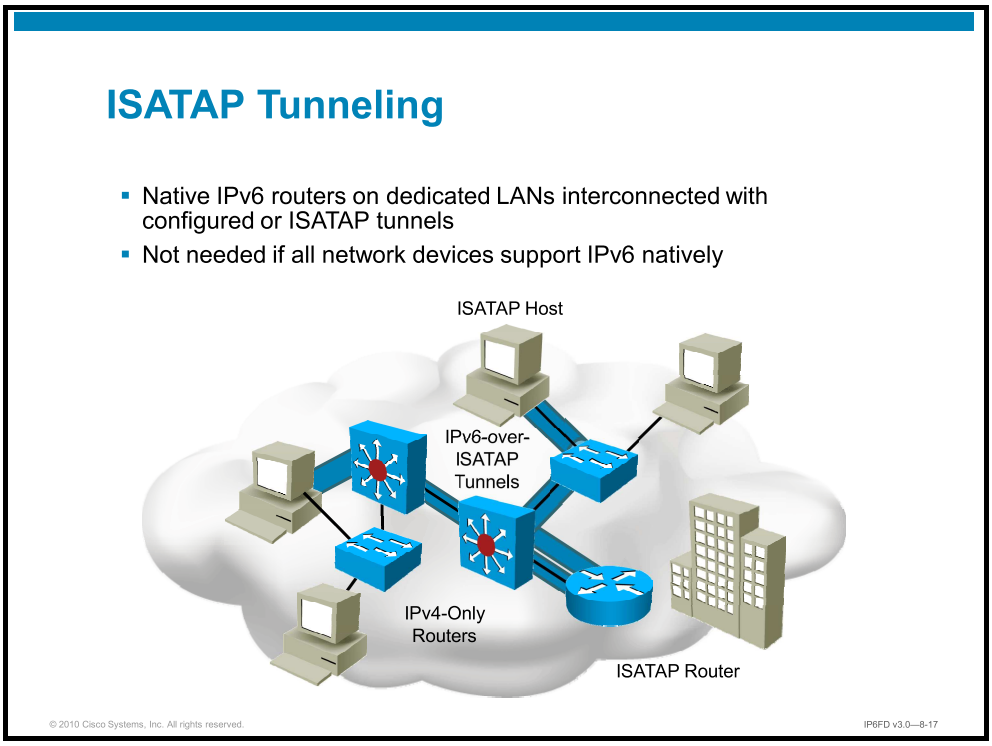

- **IPv6-over-ISATAP tunnels:** In the ISATAP case for intracampus IPv6 communications, the network infrastructure does not need upgrades. The IPv6 packets are encapsulated in IPv4 by ISATAP-capable hosts (such as Microsoft Windows XP) and delivered over the existing IPv4 network. The ISATAP router can be an upgraded existing router or a new dedicated router. Manual tunnels can also be built over the existing network. This approach enables only IPv6 unicast services; multicast is not supported over ISATAP.

https://t.me/learningnets

WAN Networks

This topic describes specific concerns about WAN networks that use IPv6.

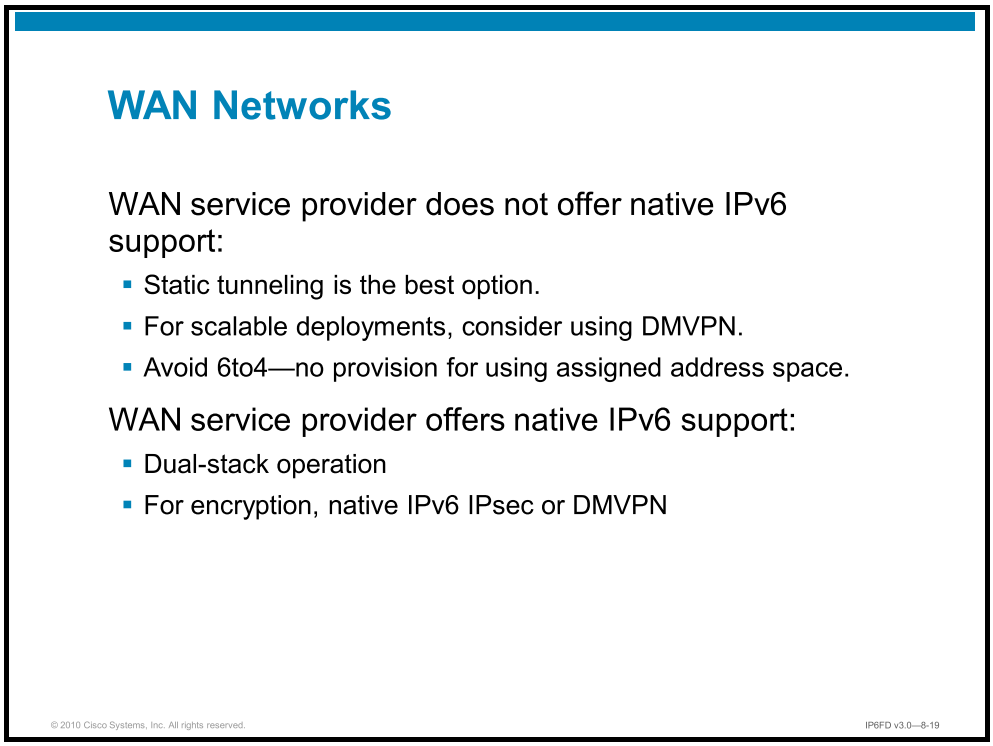

WAN networks are used to interconnect several geographically distant networks. The main characteristic that distinguishes WANs from LANs, in addition to the length of physical links, is the ownership of infrastructure. A service provider most often owns WAN networks, thus the service provider has the control over the network. Support for IPv6 rests with the provider. If the provider supports native IPv6, then this option is the best choice. Otherwise, the enterprise network administrator must use encapsulation to carry IPv6 over IPv4. There are three options:

- **Static tunnels:** These tunnels are most often used for point-to-point connections between large sites.
- **Dynamic Multipoint Virtual Private Network (DMVPN):** Used where scalability is required. DMVPN builds dynamic multipoint Generic Routing Encapsulation (GRE) tunnels, which are encrypted using IPsec. These tunnels are usually configured in a hub-and-spoke topology, but have the capability of direct spoke-to-spoke traffic, which does not burden the hub node.
- **6to4:** Can be considered as a last resort. It imposes the use of dedicated addresses for the entire site. You cannot use PA addresses for this type of tunneling.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Dual Stack: Advantages and Disadvantages

This topic describes the advantages and disadvantages of dual-stack integration in IPv6 enterprise networks.

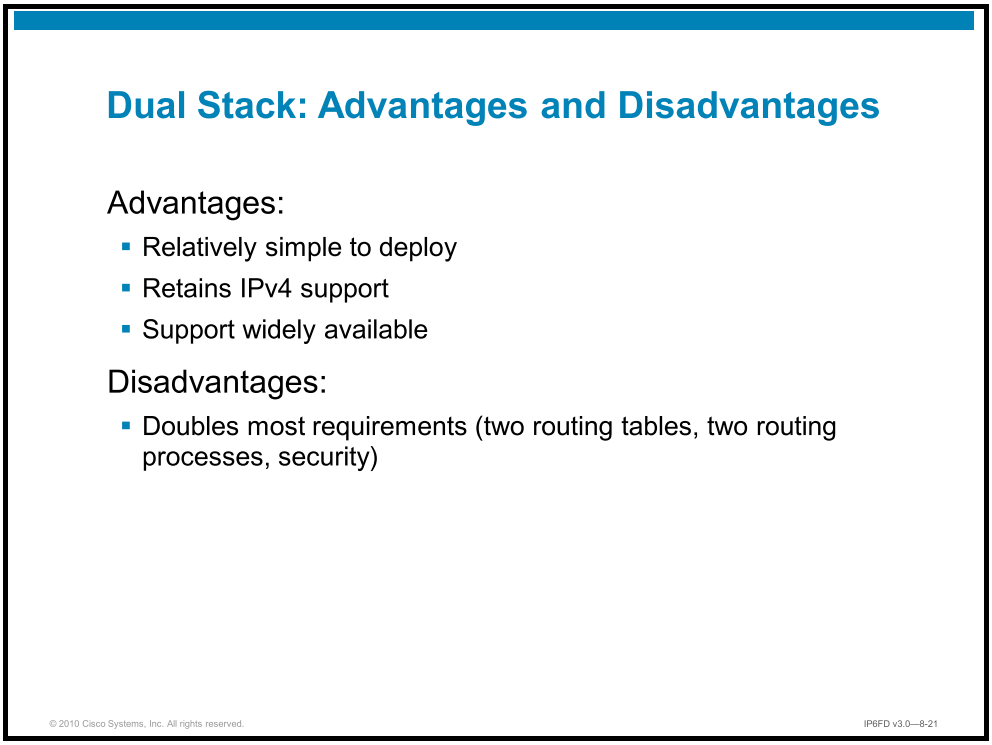

Dual stacking handles most address-selection and DNS-resolution issues. The default behavior that a dual-stack node should observe is to query DNS for all address record types and prefer IPv6 over IPv4 where both protocols are available. A well-built application will cycle through addresses in priority order if the initial connection attempt is not successful. In some implementations, a dual-stacked node may request “any” IPv6 DNS record and process the decision about which address to use locally after receiving all the responses.

Dual stacking allows hosts to continue to reach IPv4 resources while also adding IPv6 functionality.

Dual stacking provides the following benefits:

- It is simple to deploy.
- It allows backward compatibility (IPv4 support).
- It is available on most platforms.
- If there are problems with IPv6, the network can still operate on an IPv4 network.

A disadvantage of dual stacking is that a dual-stack network topology may require two routing tables and two routing processes, depending on the routing protocol. The Intermediate System-to-Intermediate System (IS-IS) protocol supports one Shortest Path First (SPF) tree for multiple protocols if the IPv4 and IPv6 networks are topologically identical.

IPv6 network security requirements are the same as for IPv4 networks. There can be an issue with flash memory size and upgrading to newer Cisco IOS Software images.

https://t.me/learningnets

Tunneling: Advantages and Disadvantages

This topic describes the advantages and disadvantages of using tunneling in an IPv6 enterprise.

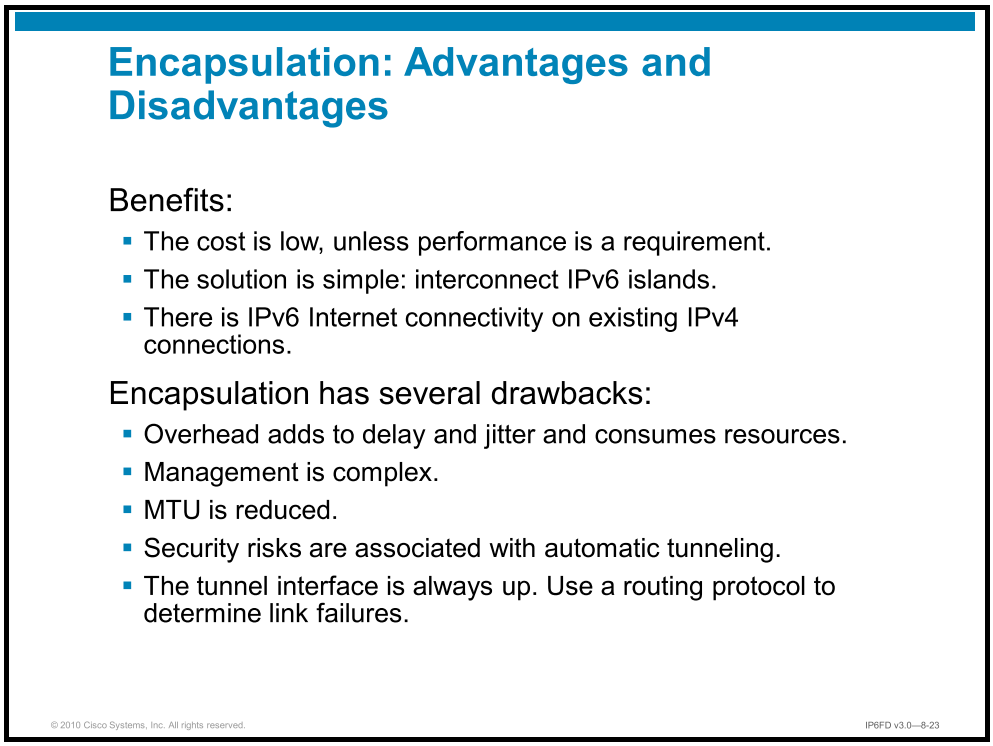

Tunneling provides a simple way to connect islands of IPv6 capability. This strategy allows an enterprise to leverage its existing IPv4 infrastructure, perhaps where equipment upgrades would otherwise be required. IPv6-in-IPv4 traffic is carried through these devices in a transparent manner.

Tunneling has the following drawbacks:

- By using tunnels, you can build IPv6 networks over IPv4 networks, but all communications will terminate at tunnel routers that cause delays because of encapsulation and de-encapsulation overhead.
- Tunnels are an interim step. Do not plan to use tunnels forever, because they will add complexity to the network as soon as IPv6 traffic grows. Tunnels also introduce their own class of problems, adding to the knowledge needed by network staff.
- Tunneling can cause packet fragmentation. Where the maximum transmission unit (MTU) of an IPv4 link may be, for example, 1500 bytes, traffic passing through a tunnel that is built over that link has a smaller effective MTU. The number of bytes used for tunnel overhead reduces the traffic. Some applications will not take this process into account, resulting in some packets being too large to pass without fragmentation, which also affects performance.
- When dynamic tunneling is used on the public Internet, it represents a security challenge because all traffic and connections are accepted. “Open relays” require tunnel endpoints to accept and de-encapsulate inbound packets from everywhere. Because the remote tunnel endpoint is anonymous, illegitimate nodes may pose as tunnel endpoints and overwhelm the campus endpoint with bogus packets. If used in a controlled environment, such as between well-known sites, tunnels can benefit from IPv4 IPsec and potentially dynamic key distribution, if they are configured for IPv4.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Translation: Advantages and Disadvantages

This topic describes the pros and cons of integrating address or protocol translation mechanisms in IPv6 enterprise networks.

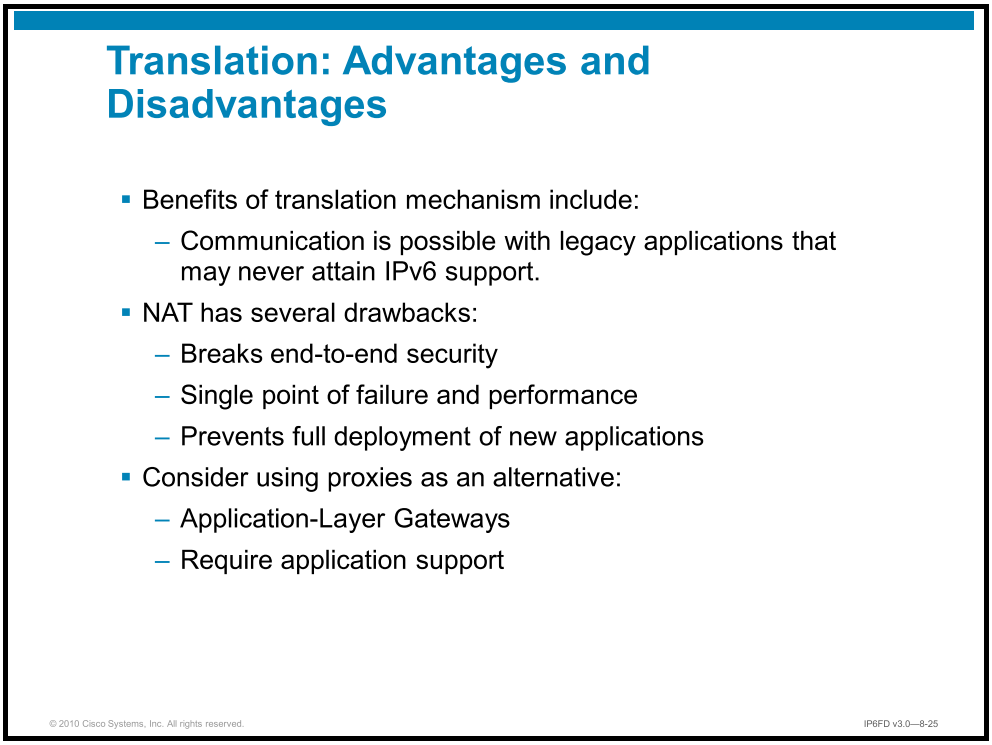

One of the translation mechanism was NAT-PT. It was an optional tool that you might have wanted to use in specific situations. NAT-PT, like NAT, requires an application layer gateway (ALG) for applications that embed an IP address in their data payload.

A benefit of network translation is that existing hosts do not require changes, because any configuration is performed at the translating router. To further illustrate the seamless transition, FTP can be used between IPv4-only and IPv6-only devices, just as within an IPv4 network, as long as the translating box supports an FTP ALG.

Packet fragmentation is enabled by default when IPv6 is configured, allowing IPv6 and IPv4 networks to resolve fragmentation problems between networks. Without the ability to resolve fragmentation, connectivity could become intermittent when fragmented packets are dropped or improperly interpreted.

Protocol translation should not be used when other native communication techniques exist. One major drawback of any IPv6 address translation mechanisms is that they have the same limitations as IPv4 NAT. Some of these generic NAT issues are as follows:

- NAT breaks IPsec-based encryption schemes because packets have been modified in transit.
- NAT causes problems for applications that embed IP addresses or Layer 4 service numbers in the payload (such as FTP, SIP, etc.).
- NAT limits the availability of services, because NAT is typically a single point of failure.

https://t.me/learningnets

Application Layer Gateways

By using proxies that can understand both IPv4 and IPv6, it is possible to proxy requests going from one network type to another. The ALG translates because everything but the application-level request is rewritten. The downside is that the application must support proxy settings, and the proxy must support the application protocol.

Note The NAT-PT feature is End-of-Life and is discouraged for implementation.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Summary

This topic summarizes the primary points that are discussed in this lesson.

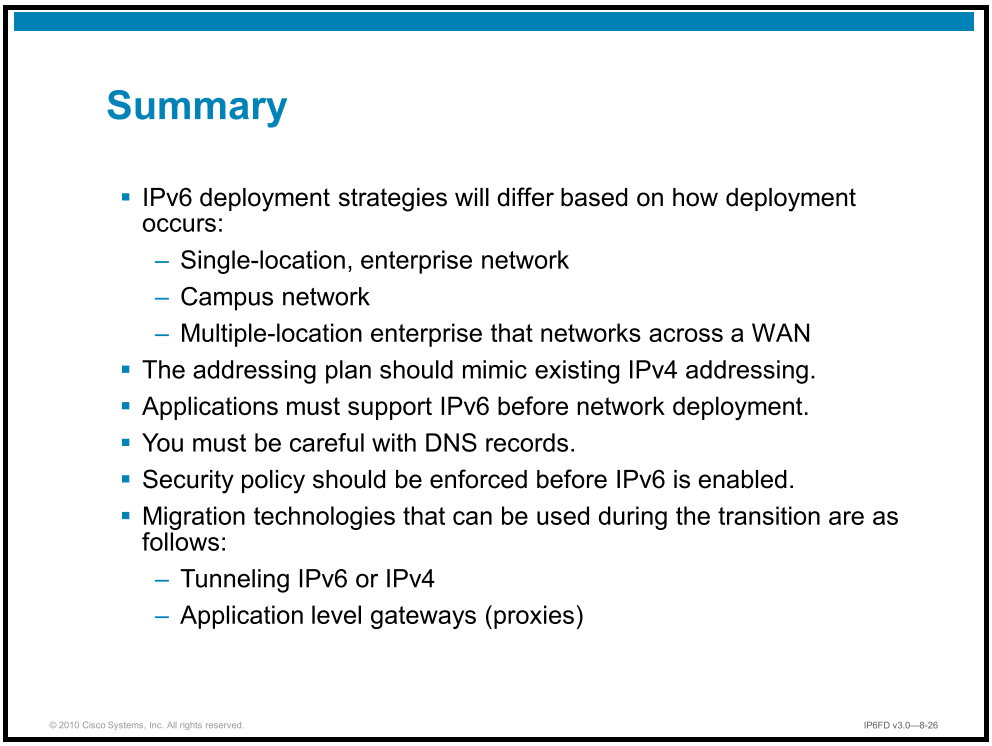

https://t.me/learningnets

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

This topic summarizes the key points that were discussed in this module.

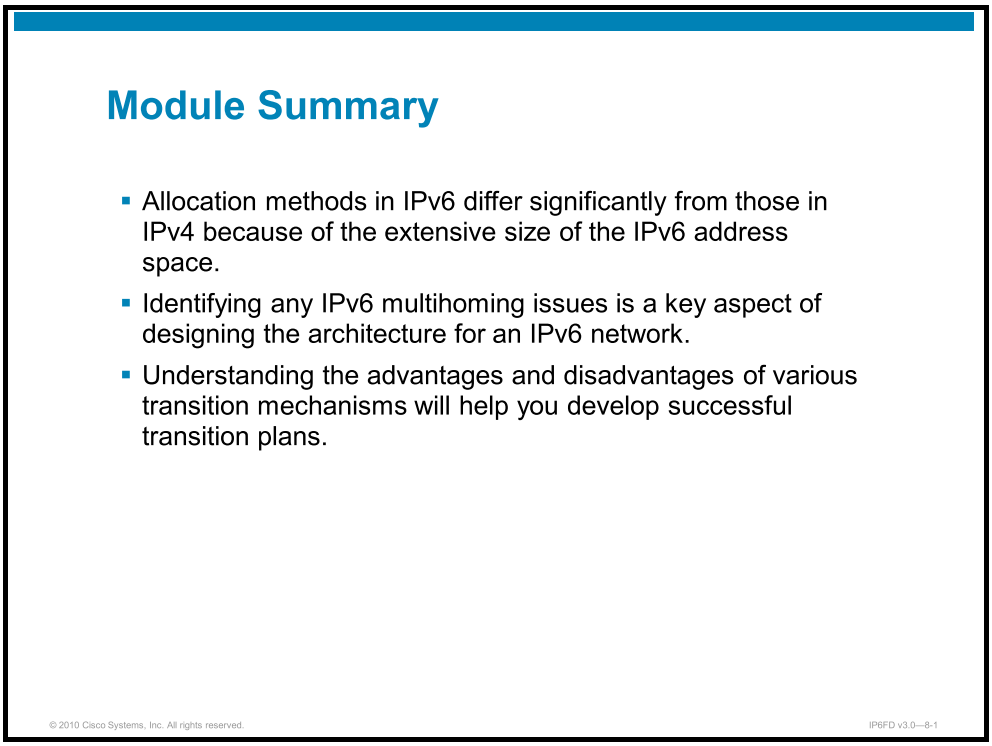

https://t.me/learningnets

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

Q1) A hierarchical chain of organizations distributes IPv6 addresses. Choose the correct hierarchy, from left to right, with the one on the far left being the top organization. (Source: Examining IPv6 Address Allocation)

A) IANA, LIR, RIR, NIR, EU
B) RIR, IANA, NIR, LIR, EU
C) IANA, EU, RIR, NIR, LIR
D) IANA, RIR, NIR, LIR, EU

Q2) What is the standard allocation for an IPv6 end site from an ISP, as specified in RFC 3177? (Source: Examining IPv6 Address Allocation)

A) /32
B) /48
C) /56
D) /64

Q3) IPv4 site-multihoming solutions are well understood by the networking community. IPv6 multihoming is somewhat dissimilar because of a key difference in the IPv4 and IPv6 address architecture. What is the difference? (Source: Understanding IPv6 Multihoming Issues)

A) In IPv4, multihomed enterprises have provider-independent address space, while IPv6 usually uses provider-allocated address space.
B) The IPv6 prefix space is less robustly aggregated than IPv4 address space, which does not lend itself naturally to multihoming solutions.
C) IPv4 multihoming does not work reliably on the Internet, while IPv6, with its large address space, should work perfectly without stressing the Default Free Zone (DFZ) on the IPv6 Internet.
D) IPv4 multihoming is, broadly speaking, implemented using IPv4 options in the IPv4 header, while IPv6 implements multihoming using the IPv6 routing extension header.

Q4) In the SHIM6 site multihoming strategy, the initial interface locator is which of these values? (Source: Understanding the IPv6 Multihoming Issues)

A) SHIM6 world-local identifier
B) MAST ID
C) interface IPv6 address
D) interface link-local address, which is built upon the globally unique 38-bit MAC address

https://t.me/learningnets

Q5) In designing the enterprise edge connection to the public network (Internet), what is one way that IPv4 and IPv6 will differ substantially? (Source: Identifying IPv6 Enterprise Deployment Strategies)

A) IPv4 traffic is likely to be run natively through the edge device, whereas IPv6 traffic will almost certainly be encapsulated in an IPv4/GRE tunnel.
B) IPv6 will outperform IPv4, and most traffic will quickly transition to IPv6 as a result.
C) IPv4 traffic will be NAT-translated at the network edge, breaking the peer-to-peer model, while IPv6 traffic will likely be routed, preserving the peer-to-peer model.
D) IPv4 and IPv6 traffic will be inspected by an edge device capable of deep inspection of their IPsec-encrypted traffic.

Q6) Tunnel techniques have several advantages over NAT-PT translation as a transition mechanism. What is a key advantage to tunneling? (Source: Identifying IPv6 Enterprise Deployment Strategies)

A) Tunneling preserves the end-to-end, peer-to-peer model of the network, whereas NAT-PT breaks end-to-end connectivity.
B) Tunneling makes it possible to fully use the Layer 2 MTU size, as opposed to NAT-PT, which suffers from the attachment of additional IPv6 extension headers.
C) NAT-PT supports IPsec via UDP encapsulation, whereas a tunneled IPv6 packet does not arrive intact at the destination.
D) IPv6-in-IPv4 tunneling delivers IPv6 packets perfectly intact at the ultimate destination, allowing IPsec to work, whereas NAT-PT breaks IPsec.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answer Key

Q1) D

Q2) B

Q3) A

Q4) B

Q5) C

Q6) A

https://t.me/learningnets

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Module 9

IPv6 and Service Providers

Overview

This module describes IP version 6 (IPv6) deployment practices for service provider networks.

Module Objectives

Upon completing this module, you will be able to describe the deployment strategies that service providers might consider when deploying IPv6. This ability includes being able to meet these objectives:

- Identify IPv6 service provider deployment strategies
- Deploy IPv6 over an MPLS network
- Explain the 6VPE service that service providers offer
- Understand IPv6 broadband access services and DSL-based access in particular

https://t.me/learningnets

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1

Identifying IPv6 Service Provider Deployment

Overview

There are many ways to deliver IP version 6 (IPv6) services to end users. Service provider environments have unique characteristics that sometimes are not present in enterprise environments. When working with service providers, you must understand the various deployment strategies, to address the specific needs of the service providers. This lesson describes IPv6 service provider deployment strategies.

Objectives

Upon completing this lesson, you will be able to identify IPv6 service provider deployment strategies. This ability includes being able to meet these objectives:

- Describe the three main methods that IPv4-only ISPs use to offer IPv6 services to their existing or new customers
- Identify the pros and cons of dual-stack deployment
- Identify the pros and cons of IPv6-only deployment
- Identify the pros and cons of encapsulation
- Identify the pros and cons of various IPv6 services
- Describe two significant service provider strategies for a phased deployment of IPv6
- Describe broadband address allocation in IPv6 networks
- Describe support for encapsulation in broadband IPv6 networks and guidelines for choosing the encapsulation type

https://t.me/learningnets

IPv6 Service Provider Deployment

Service provider networks can be broadly classified into MPLS-based networks and non-MPLS (IPv4 only) networks. This topic describes the three main methods that IPv4-only ISPs use to offer IPv6 services to their existing or new customers.

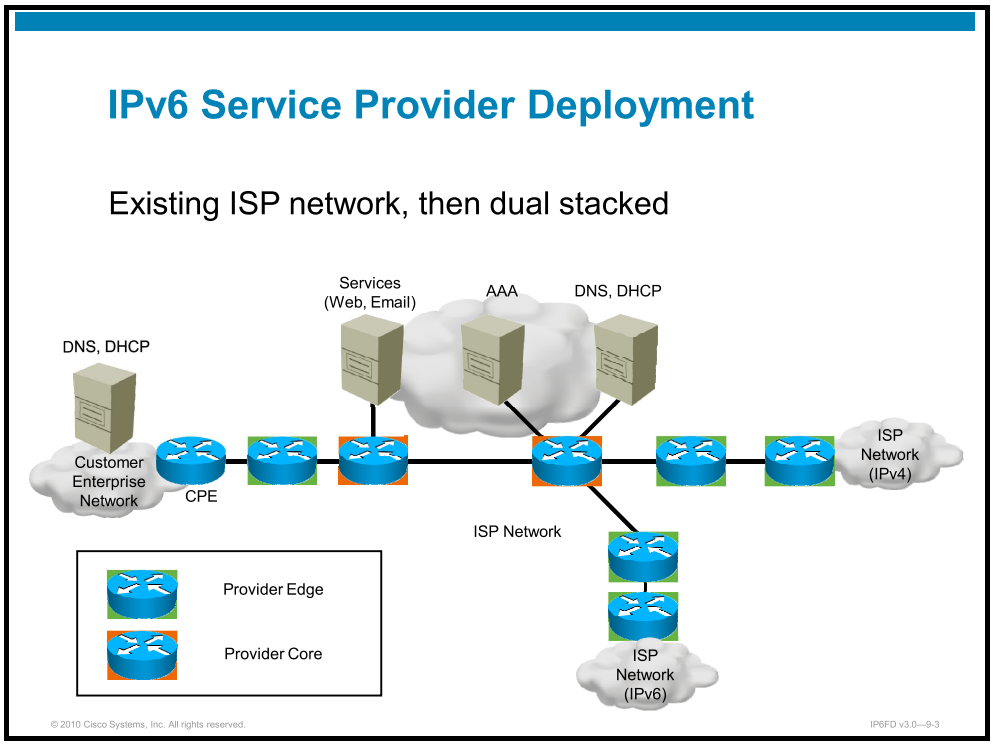

Three main methods are used by IPv4-only ISPs to offer IPv6 services to their existing or new customers. Each strategy has strengths and weaknesses that determine its ability to fit a particular ISP.

IPv4-only ISPs use three methods to offer IPv6 services to their customers:

- Upgrade the entire network to dual stack.
- Deploy a new parallel, IPv6-only network.
- Upgrade the network edge and use tunneling through an IPv4-only core.

Upgrade the Entire Network to Dual Stack

This solution is clean and straightforward, providing customers with the most functionality and performance. However, this solution can be a large undertaking for the ISP. The solution involves upgrading all ISP network infrastructure, from the core to the edge, to dual stack, in which each link and interface is configured for both IPv4 (existing) and IPv6 (added). This process requires capital upgrades to network infrastructure if nodes are not upgradable to IPv6 or have insufficient memory or other resources to run both protocols concurrently. This solution also requires a change in routing policy, new peering arrangements, and so on. However, with this approach, customers receive the benefits of native IPv6 connectivity while maintaining the same IPv4 infrastructure.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Deploy a New Parallel, IPv6-Only Network

This solution does not involve the existing IPv4 network. Instead, it involves building a new IPv6-only network and connecting IPv6 customers to that network. This option is probably the most expensive in terms of capital outlay and WAN costs, but it serves as a migration path for ISPs that are interested in retiring their existing network. This solution is probably the safest deployment in terms of ensuring that newly deployed features do not disrupt current performance, features, or availability of the existing, money-generating IPv4 infrastructure.

https://t.me/learningnets

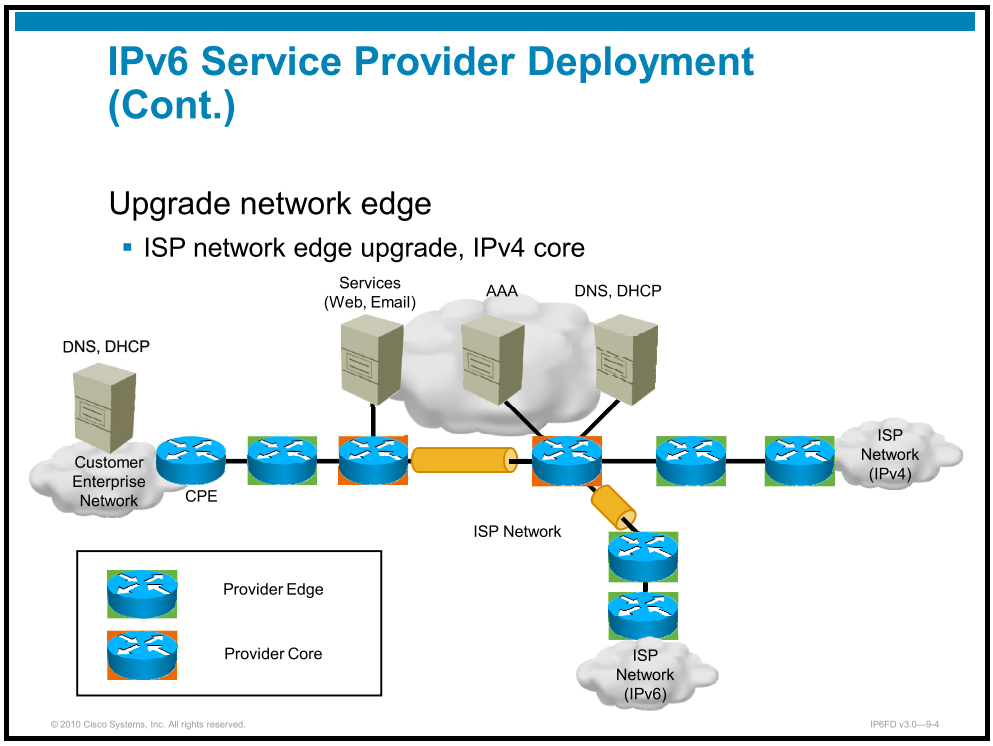

Upgrade the Network Edge

This solution involves upgrading or deploying new IPv6 capabilities at the edge of the network and tunneling IPv6 through the core network. This solution is the most balanced for many ISPs. Using this method, you can deploy IPv6 connectivity at the edges of the ISP network, near the customers who want that service, either by upgrading existing edge equipment to dual stack or by deploying new dual-stack equipment. At the same time, you can leave the core of the network alone and tunnel the IPv6 traffic to the IPv6 Internet exchange point (IXP) and other IPv6-capable customers.

Although there are numerous methods to accomplish this solution, those service providers that already support Multiprotocol Label Switching (MPLS) will find that using IPv6 integration techniques such as Cisco IPv6 Provider Edge over MPLS (6PE) and Cisco IPv6 VPN Provider Edge over MPLS (6VPE) might present cost-effective solutions that use equipment and skills that are already available. Cisco 6VPE is an emerging standard.

The figure shows the nodes that must be upgraded to support this strategy. Notice that the core routers are untouched and that only the devices near the edges, at IPv6 network setpoints, are upgraded.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Dual-Stack Deployment

This topic describes the advantages and disadvantages of dual-stack deployment in an IPv6 environment.

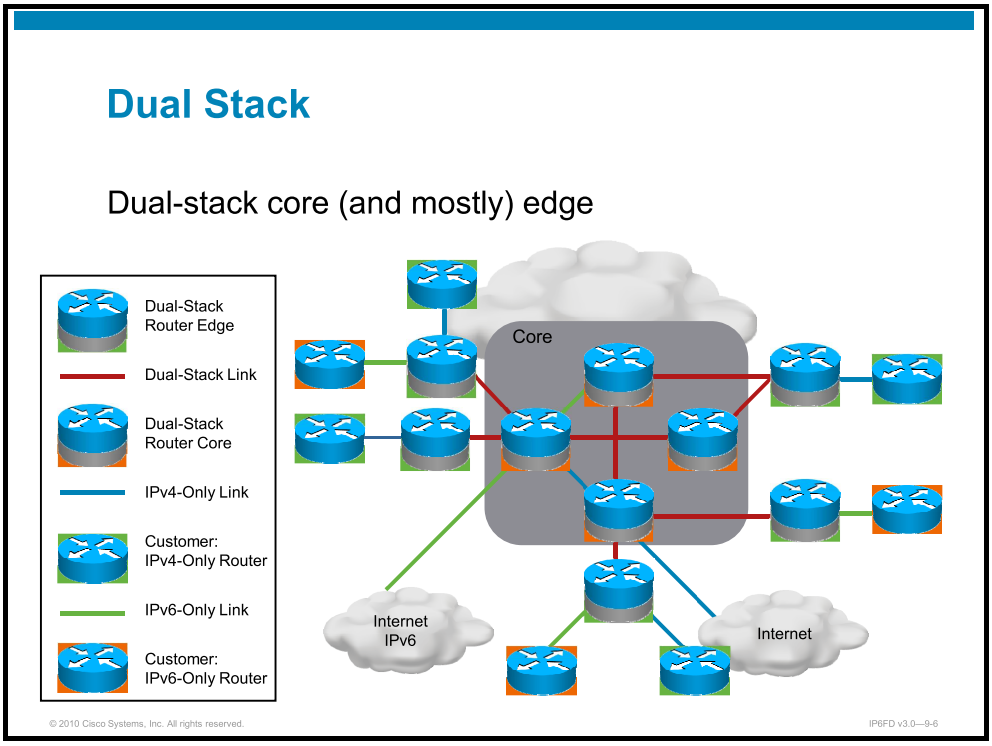

The method of deploying IPv6 by using dual-stack backbones allows IPv4 and IPv6 applications to coexist in a dual IP layer routing backbone. All routers in the network need to be upgraded to dual stack, with IPv4 communication that uses the IPv4 protocol stack and IPv6 communication that uses the IPv6 stack.

Dual-stack end systems allow applications to migrate individually from IPv4 to IPv6 transport. Applications that are not upgraded and that support only the IPv4 stack can coexist with upgraded applications on the same end system. Applications choose between IPv4 and IPv6, according to DNS name lookups. Both the IPv4 and IPv6 addresses may be returned from the Domain Name System (DNS), with the application (or the system, according to the rules that are defined in the Internet Engineering Task Force [IETF] document, *Default Address Selection for IPv6*) selecting the correct address, according to IP traffic type and particular communication requirements.

Rather than running different networks for different types of traffic (IPv6-only network) or tunneling over the existing network (requiring network upgrades and re-engineering), the dual-stack solution allows native operation of IPv4 and IPv6.

https://t.me/learningnets

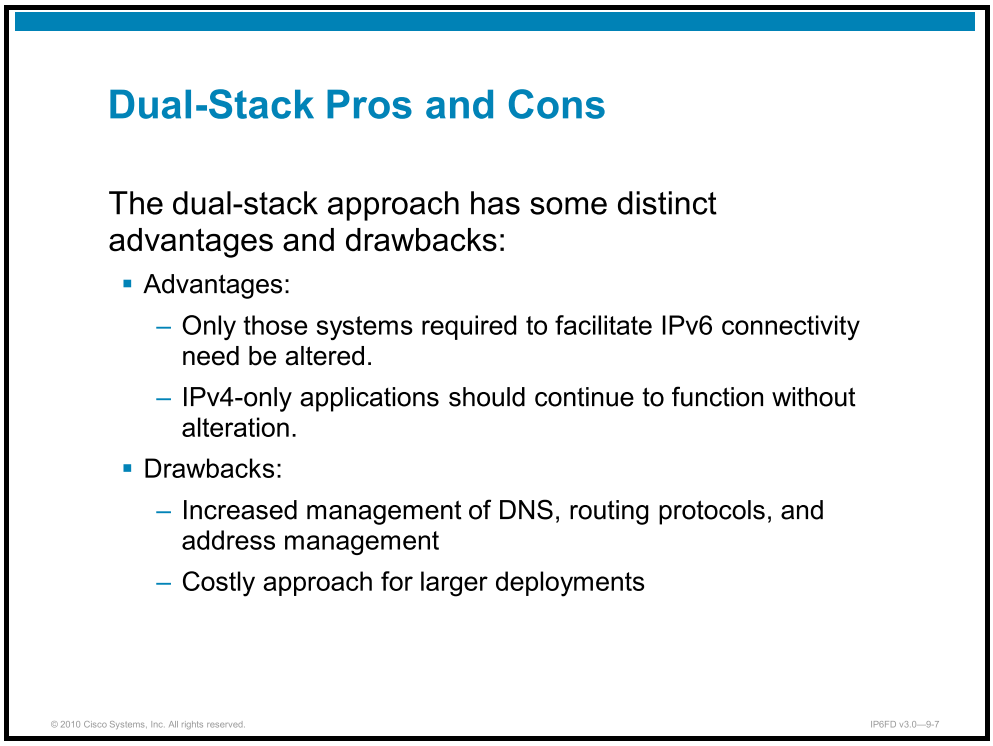

Today, dual-stack routing is a valid deployment strategy for specific network infrastructures that have a mixture of IPv4 and IPv6 applications (such as on a campus or an aggregation point of presence) and that require configuration of both protocols. Apart from the obvious need to "upgrade" all routers in the network by deploying a new protocol, the limitations of this approach are that the routers require definition of a dual-addressing scheme, dual management of IPv4 and IPv6 routing protocols, and configuration that has enough memory for both IPv4 and IPv6 routing tables.

Several other factors must be considered:

- Dual stack is easy to implement for small campus networks that have a mixture of IPv4 and IPv6 applications.
- Dual stack might require dual management of routing protocols and a major upgrade for large networks.
- All routers must be dual stack with IPv4 and IPv6 addresses.
- Dual stack requires access to IPv6 DNS (AAAA records).
- Sufficient memory is required for both IPv4 and IPv6 routing tables. In practice, this issue has not been significant. However, the size of the new Cisco IOS images that support advanced features, and the flash limitations of older equipment, can make this requirement an issue.
- The size and scope of dual stacking an entire enterprise network ensures that this solution will be costly in terms of people time and soft costs, if not in terms of capital investment. These costs are likely to greatly outweigh any equipment costs.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

IPv6-Only Deployment

This topic describes the advantages and disadvantages of an IPv6-only deployment.

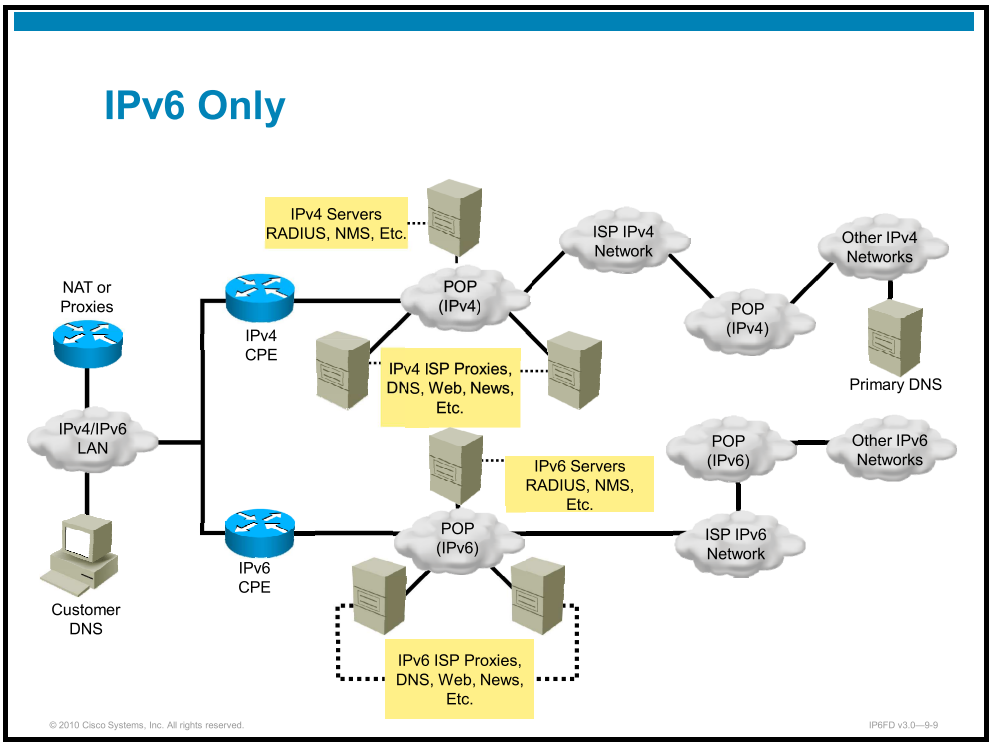

The technique of deploying IPv6 over dedicated IPv6-only data links enables isolated IPv6 domains to communicate with IPv6 by using separate Frame Relay or ATM permanent virtual connections (PVCs), separate transport circuits or optical links, or dense wavelength-division multiplexing (DWDM). The IPv4 and IPv6 networks can also share the network at Layer 2, if appropriate and if the Layer 2 networks have the capacity. Doing so creates a situation in which the IP networks are separate but share the already deployed Layer 2 infrastructure. In most cases, however, one infrastructure will support both IPv4 and IPv6, and few providers are likely to fund a separate infrastructure.

The figure shows an ISP that has deployed a completely separate network to service IPv6 customers. Notice that the customer might be using a dual-stack strategy.

An ISP that offers an IPv4 service will find that there are different ways to add IPv6 to this service. During the introduction of IPv6, the network goes through different stages of IPv6 maturity. In addition, there must be a transition between these stages, to make them feasible for implementation. In some cases, it might not be possible to introduce IPv6 as a native part of the network because of hardware limitations or similar issues. Instead, some coexistence mechanism needs to be used. This situation leaves two choices: a native IPv6 transport in a part of the network, or a mechanism such as dual stack or encapsulation to transport IPv6 over the existing IPv4 network.

https://t.me/learningnets

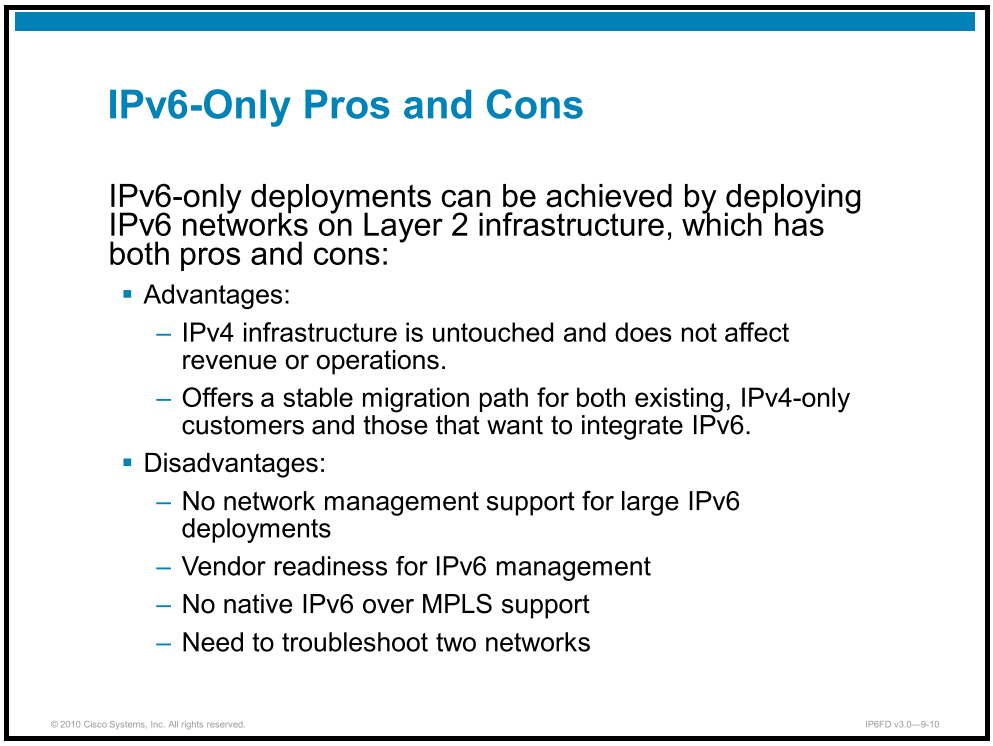

Service provider WANs or metropolitan-area networks (MANs) that use ATM, Frame Relay, or DWDM might find building IPv6-only network cores across the Layer 2 infrastructure advantageous. Such networks can provide end-to-end IPv6, with no impact on the IPv4 traffic and revenue. However, some currently deployed hardware lacks IPv6-specific hardware acceleration and support for IPv6 network management. Finally, the service provider must provide access to the WAN through a dual-stack router with IPv4 and IPv6 addresses, as well as access to IPv6 DNS (AAAA records).

Building IPv6-only networks is advantageous for some ISPs with these considerations:

- The ISP does not want to make changes to its existing production IPv4 network.
- The ISP wants to build a new network to provide IPv6 services to customers and to give customers a migration path in which current IPv4 customers will eventually end up on the new network. Networks that offer mostly IPv6, with some support for IPv4, are not very common.
- Network management is based on existing tools and packages that run only on IPv4, and will be for some time. Secure Shell (SSH), Telnet, and TFTP are important management tools that run over IPv6 transport. However, large networks require large network-management packages that will not run on IPv6-only networks any time soon.
- Providers with MPLS networks cannot run on an IPv6-only core. MPLS requires IPv4 and an IPv4 interior gateway protocol (IGP) to distribute routes in the core.
- Running two networks, one for IPv4 and another one for IPv6, increases troubleshooting efforts.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Encapsulation

This topic describes the advantages and disadvantages of encapsulation in an IPv6 network.

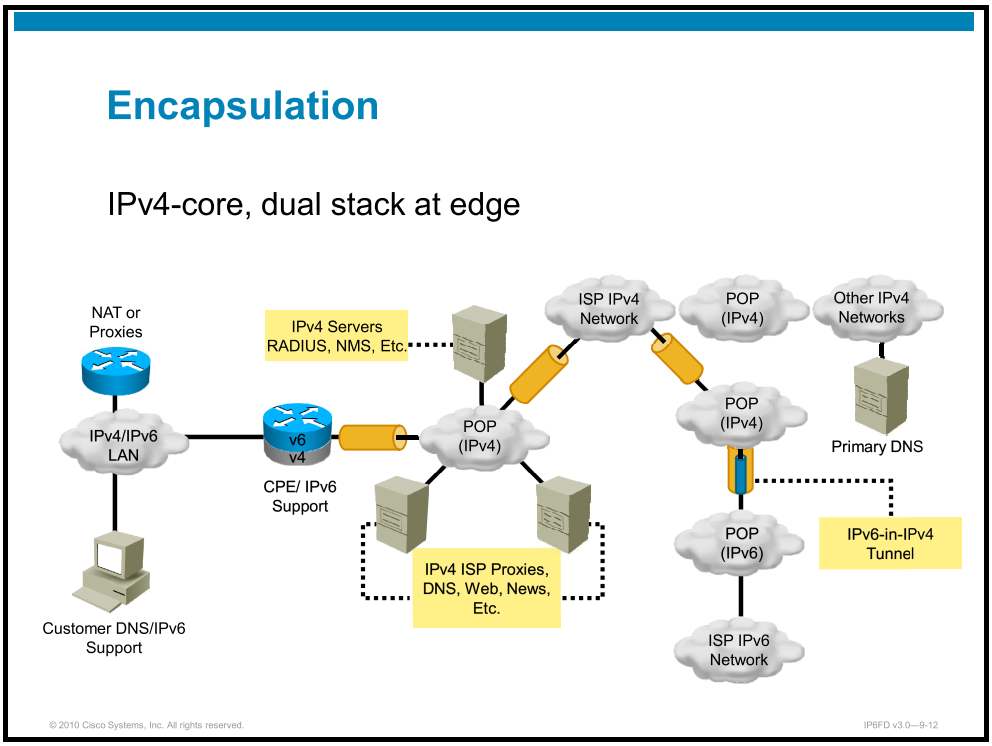

The tunneling mechanism is primarily for communication between isolated IPv6 sites or for connection to remote IPv6 networks over an IPv4 backbone. All tunneling mechanisms require tunnel endpoints to run both IPv4 and IPv6 (dual-stack mode). The dual-stack routers run both IPv4 and IPv6 protocols simultaneously and thus can interoperate directly with both IPv4 and IPv6 end systems and routers. The design is similar in concept to running IP and Internetwork Packet Exchange (IPX), DECnet, or AppleTalk on the same router, something Cisco IOS Software has done since its inception. In many senses, tunneling IPv6 in IPv4 is a hybrid strategy that still requires dual stack at tunnel endpoints.

As mentioned previously, service providers that run MPLS cores will find solutions such as Cisco 6PE to be a viable way to provide IPv6 transport over the network core. Cisco 6PE is covered in detail in another lesson.

The figure shows a dual-stack customer environment that is tunneled through an existing IPv4 ISP network to an intra-ISP setpoint with a dual-stack-capable router that then transmits customer traffic to the IPv6 Internet.

https://t.me/learningnets

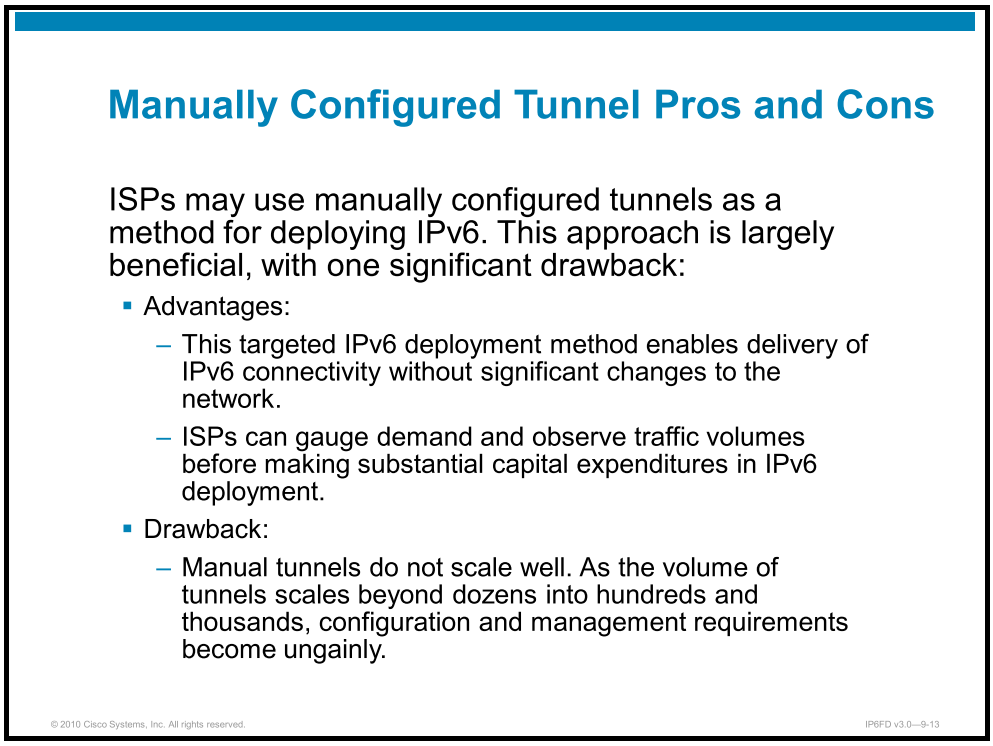

For an ISP, the advantages of using manually configured IPv6 tunnels to provide IPv6 services to customers include the following:

- Provides IPv6 connectivity to IPv6 customer sites, without disrupting the existing IPv4 service of the ISP and without a large financial or resource commitment from the ISP
- Creates new revenue streams for the ISP, by offering new IPv6 services and support for new IPv6 applications
- Enables the ISP to validate and gauge future demand for IPv6 while generating revenue from initial customer sites that request IPv6 services
- Provides stable, secure connections among the ISP, the customer site, and the IPv6 exchange point
- Enables the ISP to monitor individual traffic statistics for the customer site and to charge the customer site for IPv6 services
- Enables the ISP to change its IPv6 service offerings without affecting the /48 IPv6 prefix that is delegated to the customer site

The disadvantage of using manually configured IPv6 tunnels to provide IPv6 services to customers is that supporting additional IPv6 customers requires the ISP to manually configure additional IPv6 tunnels. Management overhead increases as the number of tunnels increases.

At some point, the number of manually configured IPv6 tunnels that each ISP border router can support might be exhausted, and the ISP will need to upgrade the border router hardware. The ISP will then need to re-engineer its IPv6 service offering.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

IPv6 Services

This topic describes the advantages and disadvantages of IPv6 services.

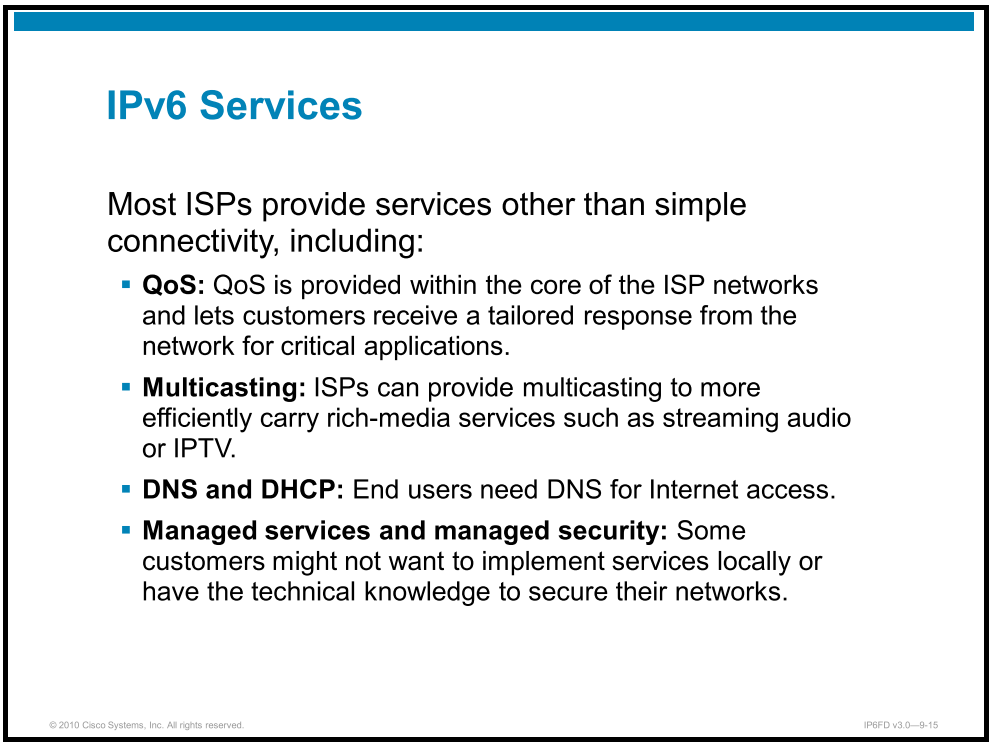

Most service providers offer services beyond basic connectivity, Internet access, and IP transport. For smaller ISPs that are closer to the end consumer, products such as web hosting and email are quite common. Larger ISPs might also provide dedicated circuits and features such as quality of service (QoS) support and IP multicasting. In many cases, ISP services are affected by the integration and support of IPv6. Some of the services that are affected by IPv6 integration are the following:

- **QoS support:** The introduction of IPv6 should not impair existing service level agreements (SLAs) or similar quality reassurances. Depending on the deployment of the IPv6 service, the service might be best effort, at least initially, even if the IPv4 service had an SLA. Both Integrated Services (IntServ) and differentiated services (DiffServ) are equally applicable and work similarly in IPv6 and IPv4. Of these two, typically only DiffServ is implemented.

  IPv6 and IPv4 currently have the same QoS capabilities. IPv6 has area service differentiation as an add-on, derived from an IPv6 flow label field. In the future, this field might add pricing information and service differentiation. QoS is also increased indirectly. Always-on connections, increased network performance, and similar features also increase QoS. As a network improves, customer expectations rise. Better-quality networks alone probably will not increase customer satisfaction. A service provider that offers these connectivity services should consider that competitors are improving their networks as well.

  In the future, charging for access quality will likely become more common as many networks implement QoS. This charging model enables a guaranteed, predetermined level of network capacity between specific locations—crucial for certain clients. This model also enables dynamic provision of different service levels to a customer, based on what the customer is willing to pay for access at any given time.

https://t.me/learningnets

QoS features that are supported for IPv6 environments include packet classification, queuing, traffic shaping, weighted random early detection (WRED), class-based packet marking, and policing of IPv6 packets. These features are available at both the process switching and Cisco Express Forwarding switching paths of IPv6.

- **Multicasting services:** IPv6 multicasting is well developed on both the network and host sides; implementations exist for Windows XP, Linux, Solaris, and so on. Many ISP implementations of multicast on IPv4 use tunnels, because few ISPs have fully multicast-enabled networks. This scenario can result in an implementation in which multicast traffic is tunneled IPv6-in-IPv6 for IPv6 networks, or in which multicast traffic is tunneled IPv6-in-IPv4 between multicast-enabled routers. Both tunneling and multicast have become well-understood concepts.

  Multicast can operate simply by using the unicast routing table as a base to build the multicast routing information, but most ISPs run a multicast-enabled routing protocol to manage the networks separately. Multiprotocol Border Gateway Protocol (MP-BGP) is used as the interdomain protocol.

- **DNS and DHCP:** DNS is required for normal operation of Internet clients and must be offered by a service provider. If clients are to use stateless autoconfiguration, their DNS settings must either be static, or acquired over DHCP. In latter case, ISP must also provide DHCP service.

- **Managed services:** Whenever the client lacks the resources or technical knowledge to set up and maintain a set of services, the service provider may offer to do so for the client.

- **Managed security:** As with managed services, an ISP that provides managed security offloads security checking, policy maintenance, and equipment from the customer. The customers do not need to set up and maintain security systems; the ISP does it for them. This arrangement might be unacceptable for customers that want to have complete control over all communications. Offloading some services to the ISP might put the ISP in a position in which critical systems are not under customer control.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Key Service Provider Strategies

This topic describes two significant service provider strategies for a phased deployment of IPv6.

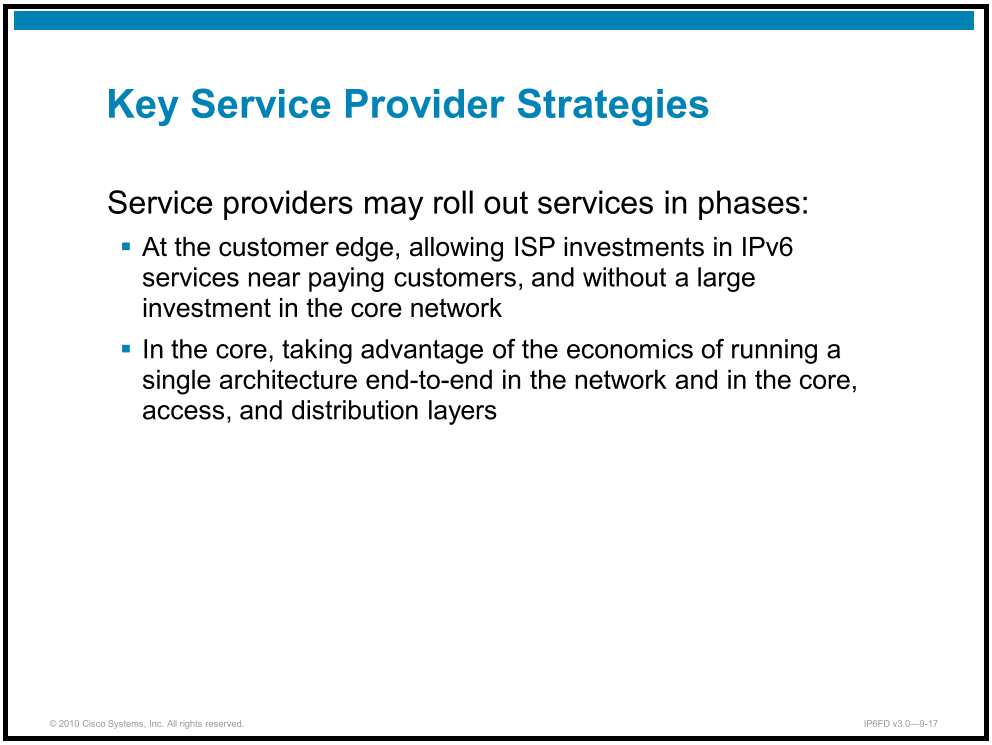

There are two main strategies for phased service provider deployment of IPv6:

- **Provide an IPv6 service at the customer access level:** Starting the deployment of IPv6 at the customer access level permits an IPv6 service to be offered immediately, without a major upgrade to the core infrastructure and without affecting current IPv4 services. This approach allows an evaluation of IPv6 products and services before complete network implementation and allows an assessment of the future demand for IPv6, without substantial investment early in the process.
- **Run IPv6 within the core infrastructure:** At the end of the initial evaluation and assessment stage, as support for IPv6 within the routers (particularly IPv6 high-speed forwarding) improves and as network management systems fully embrace IPv6, the network infrastructure can be upgraded to support IPv6. This upgrade path can involve dual-stack routers or, as IPv6 traffic gradually becomes predominant, IPv6-only routers.

ISPs that are ready to begin the IPv6 deployment process should consider these issues:

- Delegation of IPv6 global addresses
- Migration of the DNS infrastructure
- Network migration
- Customer access control and accounting

https://t.me/learningnets

Service Layer Address Allocation

This topic describes broadband address allocation in IPv6 networks.

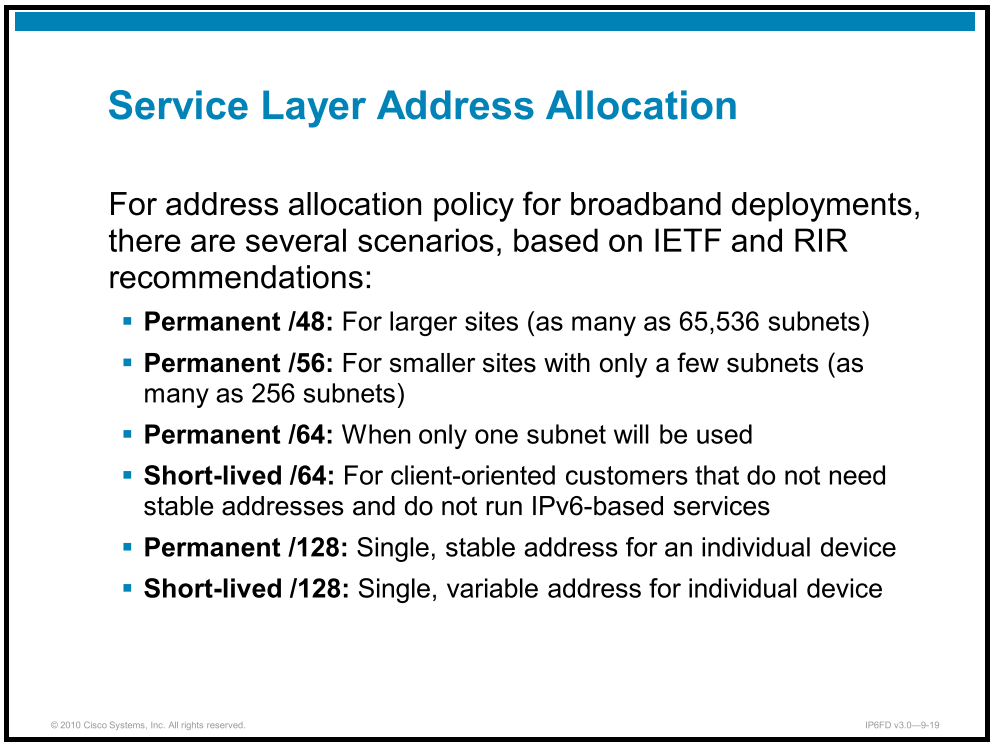

The principle for address allocation "recommends the assignment of /48 in the general case, /64 when it is known that one and only one subnet is needed," according to RFC 3177, *IAB/IESG Recommendations on IPv6 Address Allocations to Sites*.

Broadband address allocation for IPv6 presents several choices for the ISP. Several addressing schemes are available, depending on the customer environment, the need for permanent or temporary addresses, and the need to support single hosts, small networks, or large networks.

These are the most common implementations:

- **Permanent /48:** This allocation provides a customer the same /48 prefix each time it connects. The /48 prefix allocation means that the customer can have multiple networks, all capable of autoconfiguration (requires /64 prefix).
- **Permanent /56:** This allocation is like a permanent /48 prefix but for fewer possible subnets at the customer site.
- **Permanent /64:** This allocation supports a customer site that has either one device or a router with one stub network, which desires stable addresses and autoconfiguration, behind the router.
- **Short-lived /64:** This allocation supports the same environment as the permanent /64 prefix but does not provide the same /64 network to the customer each time it connects.
- **Permanent /128:** This static allocation is used to connect single devices only, because only a single IPv6 address is allocated to the customer.
- **Short-lived /128:** This variable allocation is used to connect single devices only, because only a single IPv6 address is allocated to the customer.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Address allocation is independent of the method that the customer uses to connect to the ISP network access point (NAP). These alternatives are determined for each customer (per-user), and configuration information is set up on the network access server (NAS) and on a RADIUS server that is used both for customer authentication and to establish connection parameters.

https://t.me/learningnets

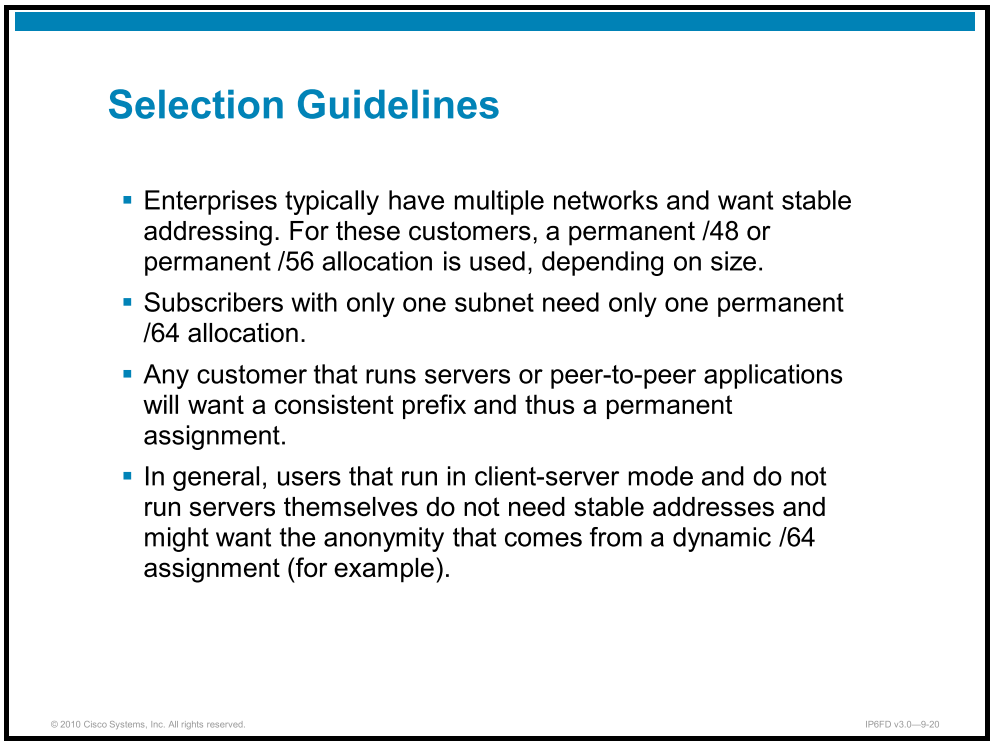

ISPs use various methods to allocate addresses to customers, and most NAS implementations have several methods in use at all times, all driven by the information that is stored in a RADIUS server or alternative authentication, authorization, and accounting (AAA) solution.

These guidelines can be used to determine which solution is best for a given situation:

- Understand the needs of the customers, in terms of how many devices and networks they need to support.
  - If a customer has multiple networks, the recommended allocation is /48, which gives the customer enough subnet bits for 64-K networks, in which each network supports autoconfiguration. Smaller allocations, such as /60 (which allows for 16 customer networks) are also technically possible. These allocations are always permanent; to make them temporary would require ongoing renumbering of the enterprise network.
  - If the customer has only one link or device (host), then the recommended allocation is /64. This allocation allows for the customer premises equipment (CPE), when it is a router, to use the bridging mode to build a multilink network. In the network, hosts behind the router can autoconfigure addresses from the same /64 that is used to connect the CPE to the provider edge (PE).
- For customers that need a /64 or a /128 allocation, determine whether the allocation should be permanent:
  - If the customer will be running peer-to-peer or server IPv6 applications or acting as a server, then the address assignment should be permanent so that stable addresses can be put into DNS. The customer will always obtain the same prefix for a /64 allocation, in which the customer can use static IPv6 address assignment.
  - For customers that use client-only applications or that want a measure of privacy from a changing prefix or address, transient allocation is acceptable and might be preferred.

https://t.me/learningnets

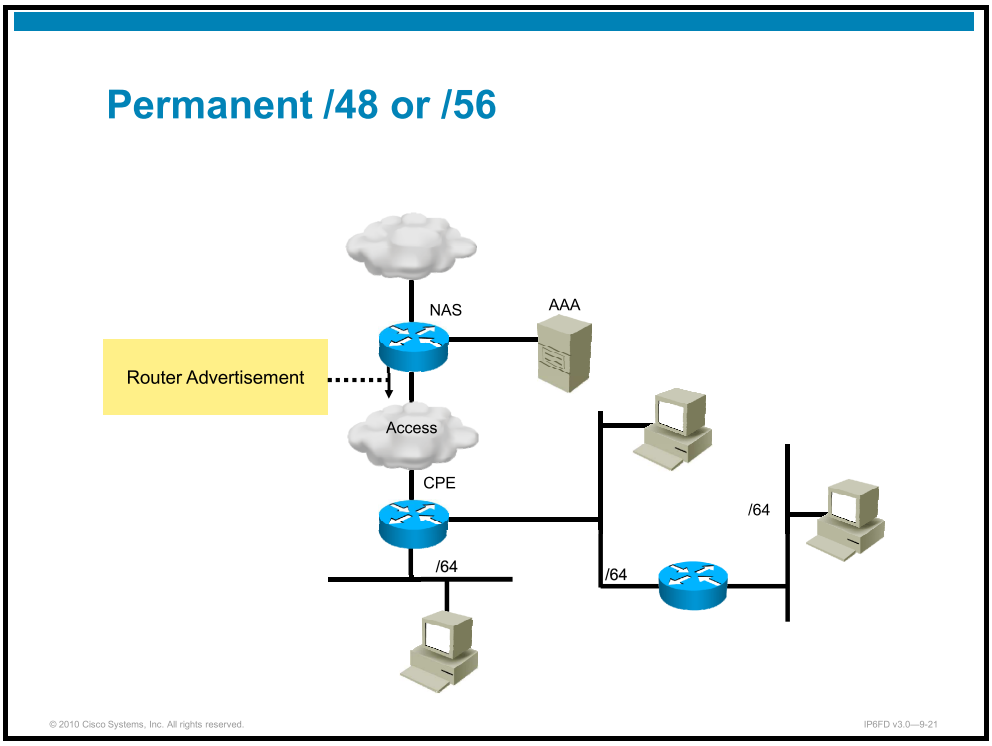

Assigning a /48 address breaks the network into the CPE-PE link (a /64) and the /48 networks behind the router. The CPE acts as a router, not a bridge. Dynamic assignment of prefixes to a downstream router (such as the CPE in this example) is not done via standard access methods but by using DHCP version 6 (DHCPv6) prefix delegation. Assignment may also be done offline, by arrangement.

The primary RADIUS attribute for this implementation is the "route" attribute:

```
cisco-avpair = "ipv6:route=2001:db8:400::/48" (Cisco Vendor
Specific Attribute [VSA])
(or)
Framed-IPv6-Route = "2001:db8:400::/48 2001:db8:100:40::2"
(syntax generalized)
where "2001:db8:400::/48" is the customer prefix and
"2001:db8:100:40::2" is the next hop.
```

This attribute specifies the route to the customer /48 network, which is installed in the NAS when the customer connects and is authenticated. Note that the Cisco vendor-specific attribute (VSA) does not include the next-hop gateway. Because the CPE-PE link is configured cooperatively between the devices, the NAS should know the next hop to this customer. Therefore, the next hop does not need to be specified.

If DHCPv6 prefix delegation is not used, then the /48 networks would be allocated manually within the customer network.

Note that this approach provides the customer with a /48 prefix for use inside its enterprise but does not provide any addressing for the PE-CPE link. This addressing would come from either a permanent /64 or permanent /128 allocation and would not be used inside the customer network.

https://t.me/learningnets

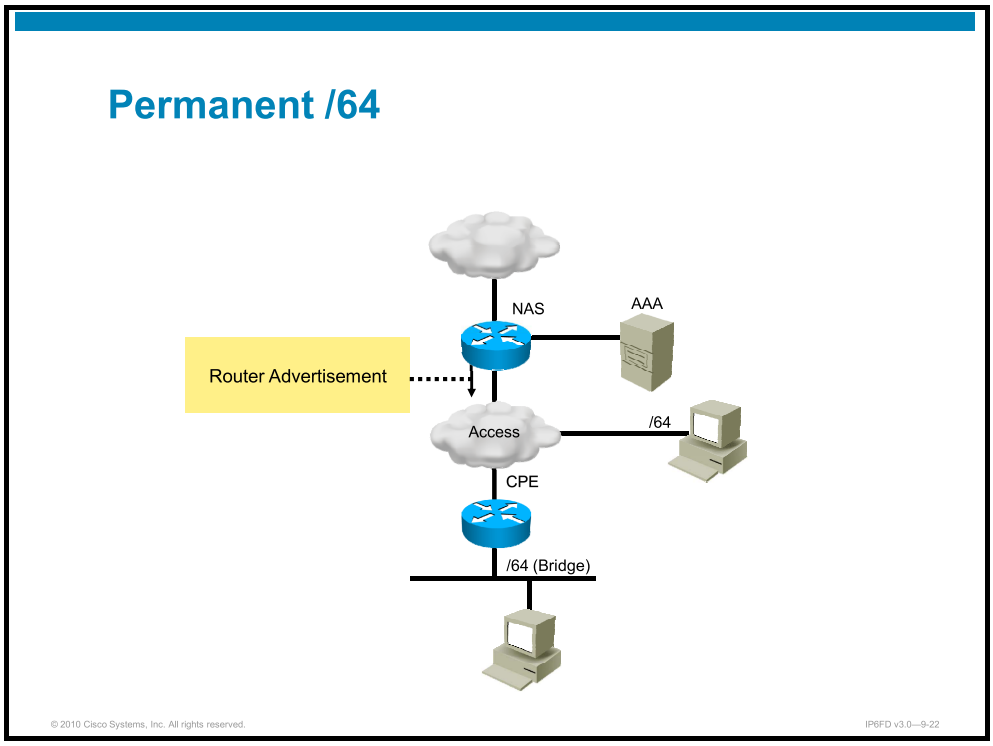

For typical home users and very small businesses in which the customer CPE is either a single device (home PC) or a router (such as a Cisco 837 Router) with one network segment behind it, allocation of a /64 network is an excellent solution. If the CPE is a router, then the router bridges the CPE-PE link to the internal link, creating a multilink network that uses a single /64 network. The NAS sends route advertisements along the link, which the CPE bridges to the internal network for host or node autoconfiguration.

To assign a permanent /64 allocation, the “prefix” parameter is stored in RADIUS (per-user) and uses either of these configurations:

```
cisco-avpair = "ipv6:prefix#1=2001:db8:200:abcd::/64"
(or)
Framed-IPv6-Prefix = "2001:db8:200:abcd::/64" (generalized)
```

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

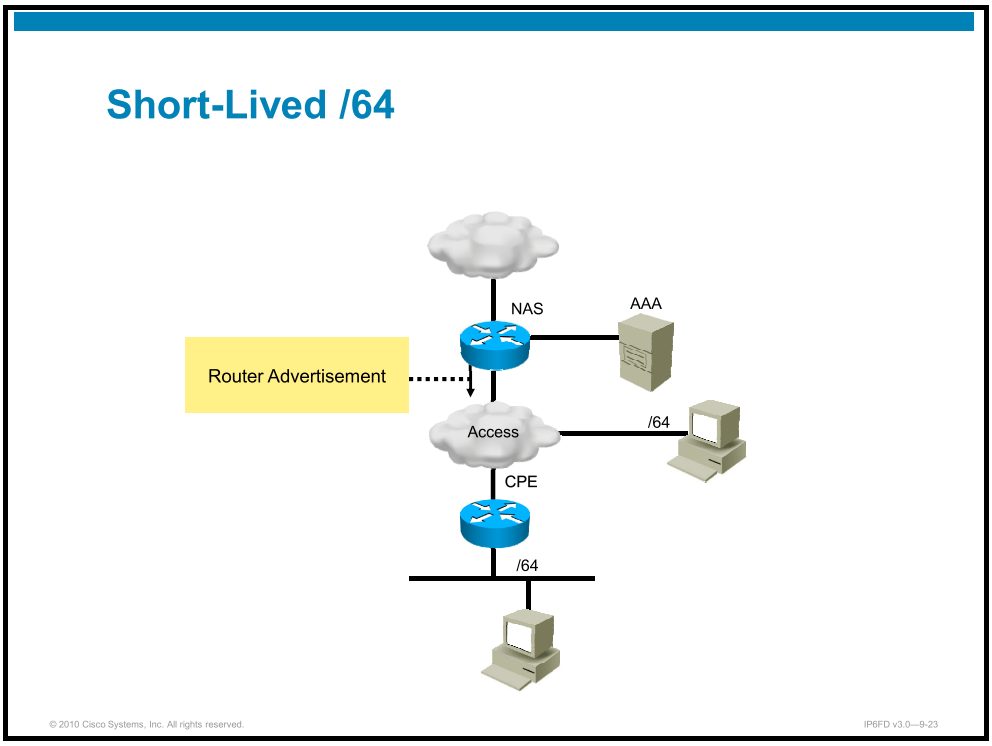

This allocation is acceptable for home users and small businesses in which the customer CPE is either a single device (PC) or a router with one network segment behind it. Allocation of a short-lived /64 network is an excellent solution.

There are two main variations on this implementation:

- For the short-lived /64 allocation, the CPE prefix is assigned from a pool of prefixes that are defined on the NAS. The RADIUS configuration does not specify which prefix to allocate to the customer (as it does for permanent prefixes) but specifies which pool to use when allocating a dynamic prefix to the customer:

```
In RADIUS, specify the existing pool to allocate the prefix
from, either:
cisco-avpair = "ipv6:addr-pool=foo"
(or)
Framed-IPv6-Pool = "foo"
(where "foo" is the name of the pool)
```

- Alternatively, the RADIUS configuration can perform only authentication, and the NAS can be configured with a local prefix pool. The customer is allocated a prefix from this pool:

```
"ipv6 local pool foo 2001:db8:b::/56 64"
(where "56" is the length of the pool and "64" is the prefix
length from the pool assigned to users)
(has 256 /64 prefixes available to allocate)
```

When the pool is defined in the RADIUS server, the customer will always get a /64 prefix from the same pool, even if the CPE connects to a different NAS at different times.

When the pool is a local pool, the customer will have a /64 allocated from different pools on different connections to the NAS, especially if the CPE does not always connect to the same NAS.

https://t.me/learningnets

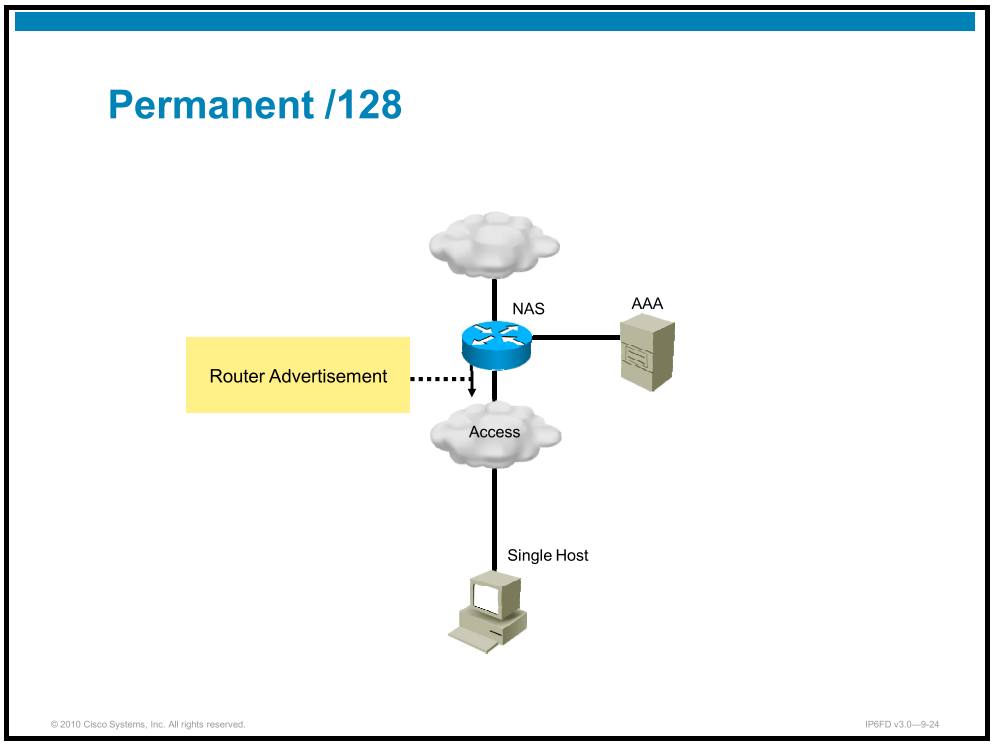

For customers that have only one device to connect, a /128 allocation from a shared pool of addresses is an acceptable solution. This allocation is most like the IPv4 solution, although the address is still autoconfigured rather than negotiated by IP Control Protocol (IPCP). This solution provides no facility for other devices; multiple devices cannot share a single address. Use of a specified interface identifier makes the address assignment complete (all 128 bits are specified). The prefix that the CPE uses comes from a shared pool.

In a permanent /128 allocation, the customer is provided with one permanent address. The parameter is stored in RADIUS (per-user) and uses either of these configurations:

```
cisco-avpair = "ipv6:pool=foo-128"
(or)
Framed-IPv6-Pool = "foo-128" (generalized)
```

The permanent /128 scenario uses a RADIUS option, which allows the NAS to choose the interface identifier that is negotiated for the CPE interface that faces the NAS, so that the entire address is known. This option is accomplished via the following:

```
Framed-Interface-ID = "0:0:0:2"
(where the address shown is the CPE interface)
```

The pool must be defined on each NAS to which the customer may connect on more than one successive connection, as in the permanent /64 case:

```
ipv6 local pool foo-128 2001:db8:f::/64 128
(where "64" is the length of the pool, and "128" is the
allocation to the user)
(has 2^64 interface identifiers [and thus addresses] to
assign)
```

Permanent /128 allocation is an excellent solution when the customer has one device and wants a permanent address for name-to-address resolution.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

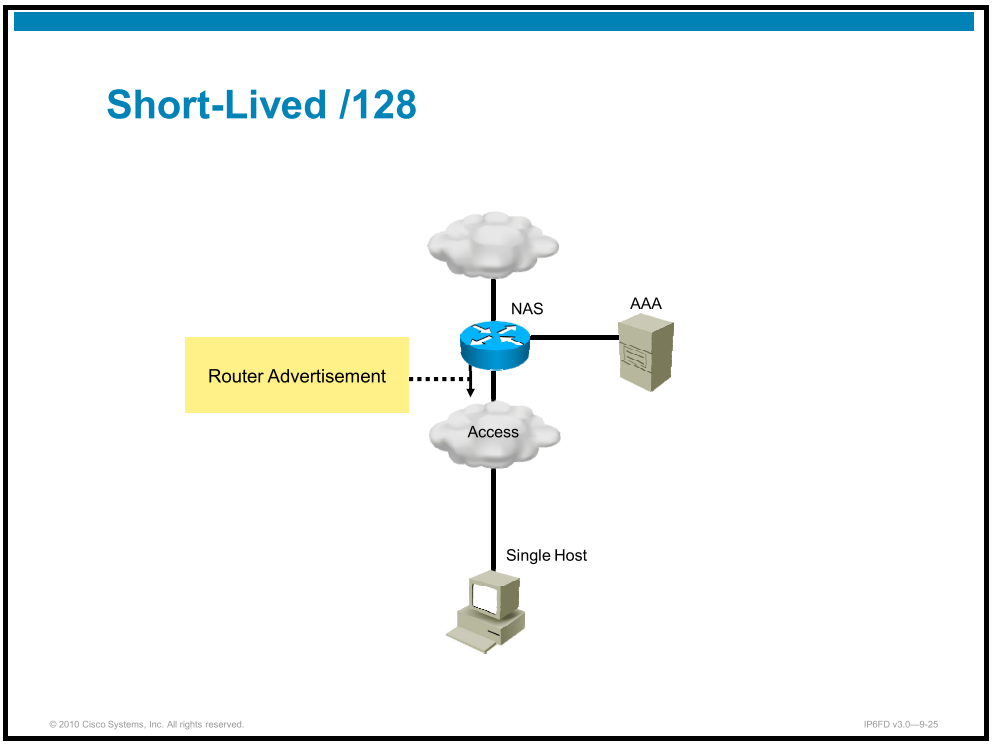

For customers that have only one device to connect, a /128 allocation from a shared pool of addresses is an acceptable solution. This allocation is most like the IPv4 solution, although the address is still autoconfigured rather than negotiated by IPCP. This solution provides no facility for other devices; multiple devices cannot share a single address. For this scenario, an interface identifier is not specified, and the CPE is allowed to configure whichever address is desired from the shared pool.

In a short-lived /128 allocation, the customer is provided with an address from a named shared pool. The parameter is stored in RADIUS (per-user) and uses either of these configurations:

```
cisco-avpair = "ipv6:pool=foo-128"
(or)
Framed-IPv6-Pool = "foo-128" (generalized)
```

The short-lived /128 scenario does not use the RADIUS option to specify the interface identifier for the CPE.

The pool must be defined on each NAS to which the customer may connect on more than one successive connection, as in the permanent /128 case:

```
"ipv6 local pool foo-128 2001:db8:f::/64 128
(where "64" is the length of the pool, and "128" is the
allocation to the user)
```

The customer typically autoconfigures the same address when it has the same interface identifier, because it will always be sent a router advertisement from the NAS, as specified in the foo-128 pool.

https://t.me/learningnets

Encapsulation Support

This topic describes support for encapsulation in broadband IPv6 networks and guidelines for choosing the encapsulation type.

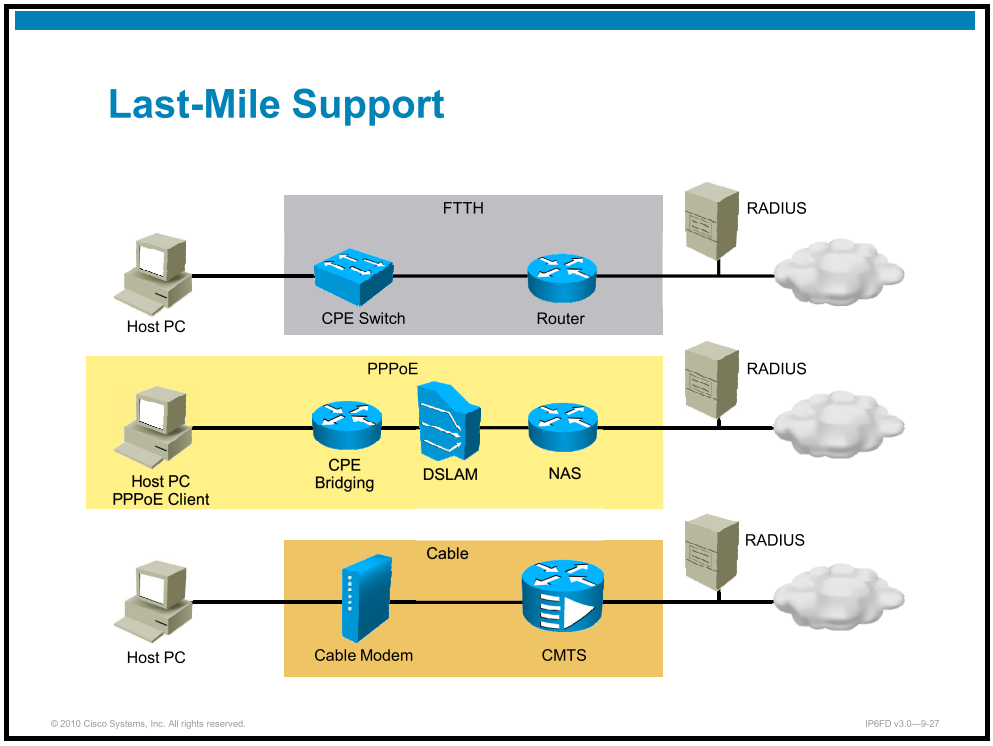

Broadband architectures for Internet service offer these last-mile implementations, among others:

- **Fiber to the home (FTTH):** This option has gained popularity recently, with the advent of fiber lines all the way to the residential buildings. The CPE device is typically a switch that has optical uplink and copper downlinks, where end users directly connect their own equipment or end nodes.
- **PPP over Ethernet (PPPoE):** PPPoE architecture typically is a CPE router with Ethernet on the customer side and ATM on the ISP-facing side. The CPE router can initiate the PPP session to the ISP PE or NAS, or the PPP session can be originated from each PC within the Ethernet network. The backend infrastructure is the same as for PPP over ATM (PPPoA).
- **Cable network (CATV):** Cable networks reach all homes that have cable TV. In such networks, the CPE device is a cable modem. Modern modems might have router capabilities; otherwise, they bridge Ethernet segments over CATV network.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Selection Guidelines

Several factors must be considered when choosing an encapsulation type for broadband access. The major issues relate to currently deployed equipment, staff skills in helping customers maintain encapsulation clients on their hosts, and the need for authentication and accounting. These are some selection guidelines for choosing an encapsulation type:

- Routed Bridged Encapsulation (RBE) requires no changes to the integrated routing and bridging (IRB)-based CPE equipment that many ISPs have deployed throughout their customer locations. To upgrade the capability to RBE, all the changes are made on the NAS equipment, and those changes are straightforward. No software is required on the customer-networked hosts. This solution is a good choice for small office, home office (SOHO) deployments that require a single /64 prefix allocation, because the route advertisements from the NAS are bridged onto the customer LAN.
- PPPoA requires no special software on the client host machines, and it supports IPv6 allocations of /48, /64, and /128. CPE configuration is different in these cases; for example, /64 would use bridging. In all cases, the CPE runs PPP to the NAS, where authentication can take place with a RADIUS server. This solution requires more CPE configuration than RBE does, because the CPE must be provisioned with username and password for PPP or AAA authentication. One other restriction is that, because the CPE is the endpoint for the PPP session, the ISP cannot view how many individual network users are using the CPE-PE-Internet connection.
- PPPoE requires client software on the network hosts, because the individual hosts terminate the PPP sessions on the customer end. As in PPPoA, PPPoE provides strong authentication. PPPoE is also ideal when the CPE equipment cannot be upgraded to use a PPP client itself. In this case, the CPE acts in bridge mode and all sessions simply pass through the CPE. The PPPoE session can also be terminated at the CPE, removing the need for host client software. This practice results in a configuration like PPPoA, in which only one session exists between the CPE and the NAS.

https://t.me/learningnets

Office Broadband Deployment

PPPoE can be terminated on:

- PCs
- CPE router

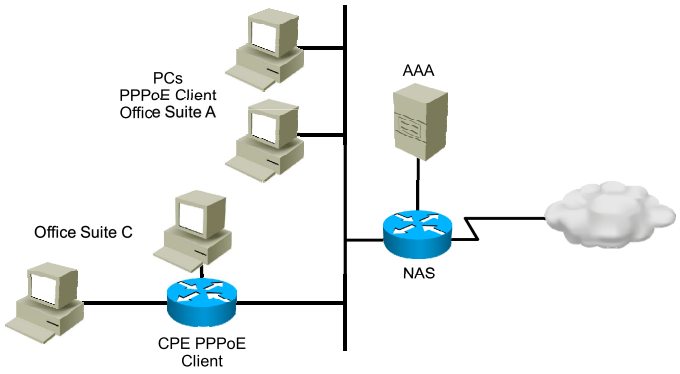

© 2010 Cisco Systems, Inc. All rights reserved.

A professional center with many small offices, each of which is an independent business, provides an opportunity to deploy PPPoE effectively. In this case, the building is wired with Ethernet, much like a campus or enterprise-dedicated building. Hosts from each business connect to the shared building Ethernet, which provides each host access to the ISP CPE equipment at the wiring closet.

The CPE uses Ethernet on the building side and ATM facing the telephone company (telco) network. The CPE is set up in bridge mode, passing packets transparently through to the building LAN.

Customer hosts are configured with PPPoE clients. Each host is an independent ISP customer; their PPP sessions are terminated on their local hosts and on the ISP NAS. These clients can be serviced and billed separately, and each can obtain different IPv6 services; for example, some customers can receive permanent /128 addresses, whereas others receive temporary /128 addresses.

Customers can also install their own in-office network and connect to the ISP from a private CPE, where the PPPoE client is on the in-suite CPE. This approach allows the customer to have multiple hosts but only one session with the ISP.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

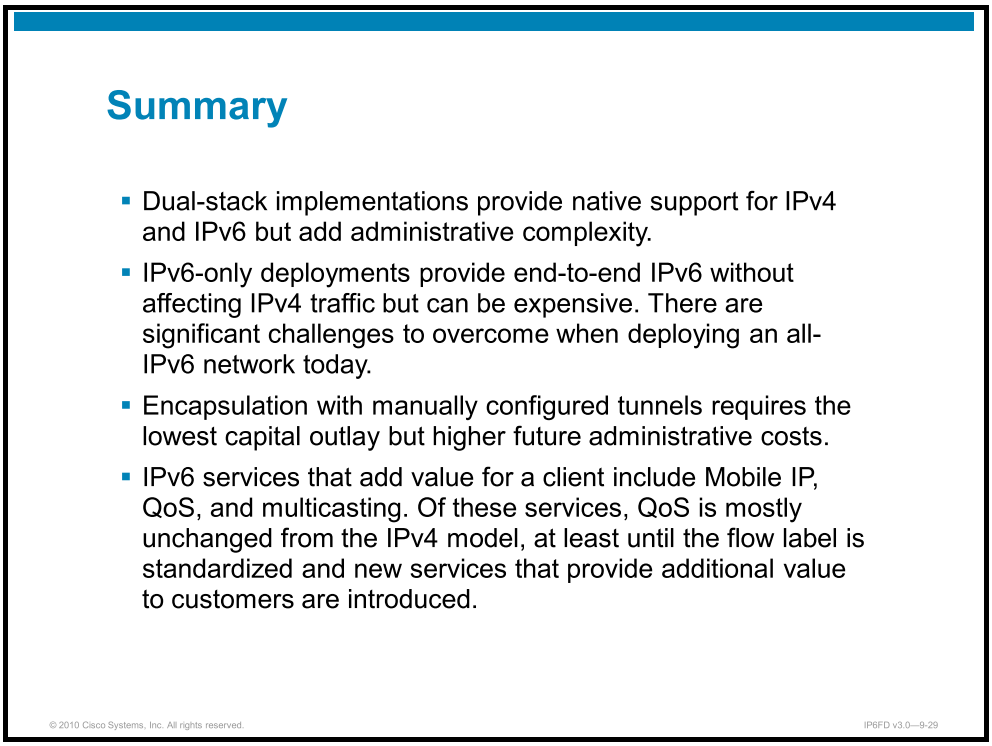

https://t.me/learningnets

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Understanding Support for IPv6 in MPLS

Overview

When deploying IP version 6 (IPv6), you need to understand how to deploy IPv6 over a Multiprotocol Label Switching (MPLS) network. MPLS is an important enabling technology for IPv6 WANs.

Objectives

Upon completing this lesson, you will be able to deploy IPv6 over an MPLS network. This includes being able to meet these objectives:

- Explain the operations of MPLS
- Describe the various methods that are used to deploy IPv6 over MPLS
- Explain IPv6 tunnels that are configured on CE routers
- Explain IPv6 over AToM
- Describe Cisco 6PE
- Describe how Cisco 6PE is deployed in an MPLS network

https://t.me/learningnets

MPLS Operations

This topic describes the different MPLS operations that are involved in a transitional IP version 4 (IPv4)-to-IPv6 network environment.

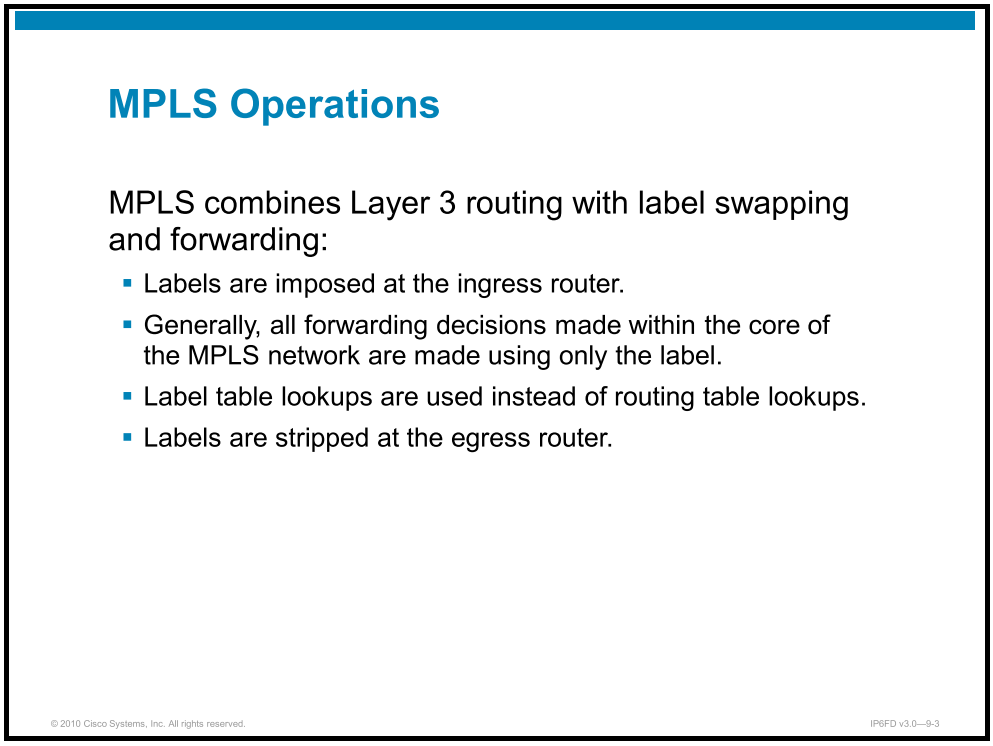

MPLS is a packet-forwarding technology that uses labels to make data-forwarding decisions. With traditional routing, each router holds an entire routing table and forwards to the next hop (destination-based routing). With MPLS, the Layer 3 header analysis is done just once—at the edge, when the packet enters the MPLS domain. Label inspection influences subsequent packet forwarding.

MPLS provides these beneficial applications:

- Virtual private network (VPN)
- Traffic engineering
- Quality of service (QoS)
- Any Transport over MPLS (AToM)
- The core routers do not need to know or carry customer prefixes, thus there is less memory and CPU overhead

MPLS also decreases the forwarding overhead on the core routers. MPLS technologies are applicable to any network layer protocol.

MPLS over a native IPv6 network is being considered but is not yet available.

Tip MPLS networks use IPv4 in the core, but they use it strictly on the control plane. The data plane uses MPLS label switching to switch packets across the core. Providers are heavily invested in their MPLS core infrastructures, and are generally not interested in upgrading all their MPLS devices (core and edge) to use IPv6.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

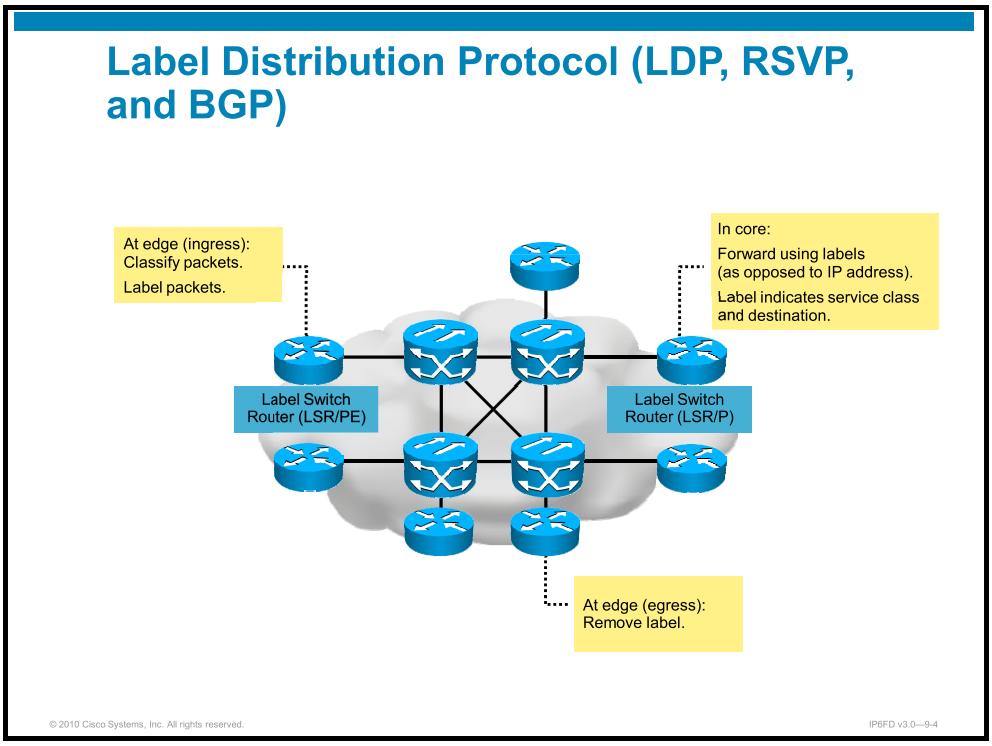

The figure shows how MPLS works, including terminology. At the edge of the MPLS network, inbound packets are processed, classified, and given an MPLS label that determines how they will travel through the core to the egress router.

The core devices look at the label, make a switching decision, swap the label for a new label, and send the packet out the appropriate interface to the next hop.

https://t.me/learningnets

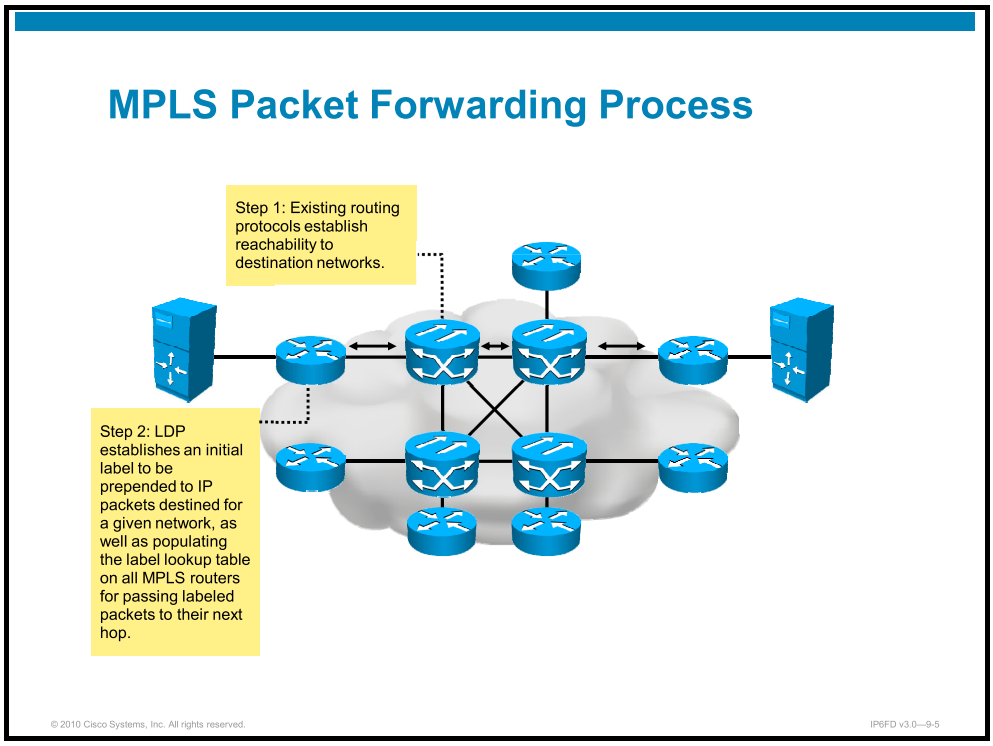

This example describes how a label-switched path (LSP) is established and how packets are forwarded through an MPLS network.

1. Existing routing protocols establish reachability to destination networks.

An interior gateway protocol (IGP) routing protocol such as Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), or Intermediate System-to-Intermediate System (IS-IS) determines the Layer 3 topology. A Cisco router or IP over ATM switch can have a routing function inside that takes care of this process. All devices in the network build the Layer 3 topology.

2. A label distribution protocol (LDP) establishes label-to-destination network mappings.

The LDP establishes label values for each device, according to the routing topology, to preconfigure paths to destination networks. Unlike ATM permanent virtual circuits (PVCs), in which the virtual path identifiers (VPIs) and virtual channel identifiers (VCIs) are manually assigned, labels are assigned automatically by the LDP.

https://t.me/learningnets

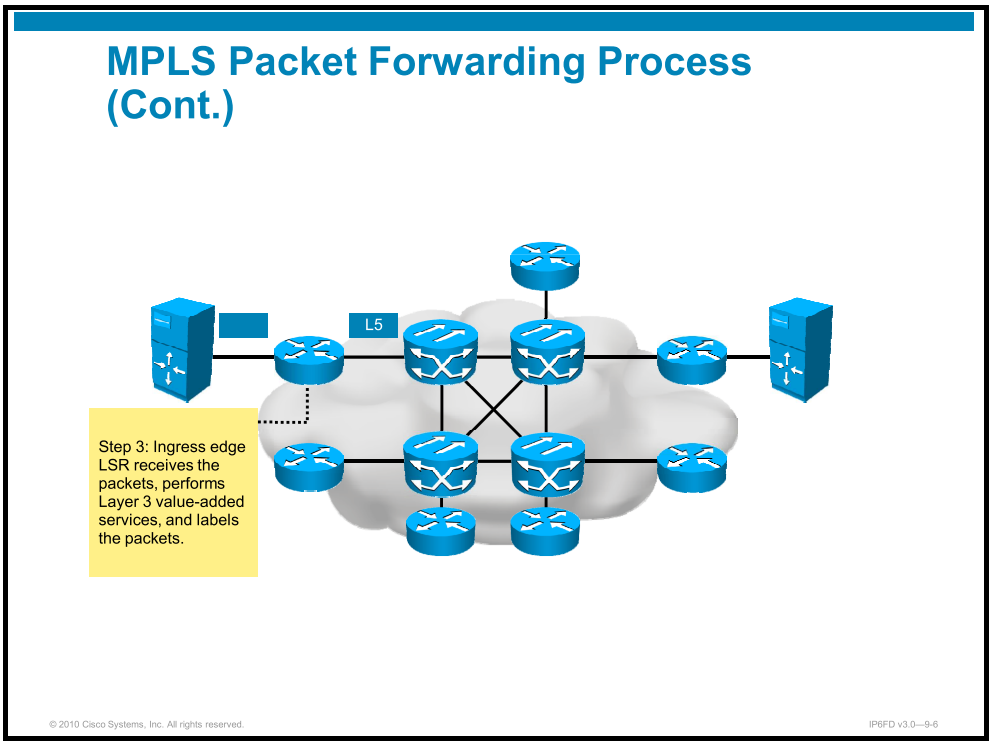

3. The ingress edge label switch router (LSR) receives packets, performs Layer 3 value-added services, and labels packets.

An ingress packet enters the edge LSR. The LSR consults the IP routing table on the edge router, and prepends an MPLS label to the packet. A single label may be used for a destination network, or multiple labels may exist to carry the packet via different paths or to provide some specific value-added service (such as priority handling across a more preferred path).

https://t.me/learningnets

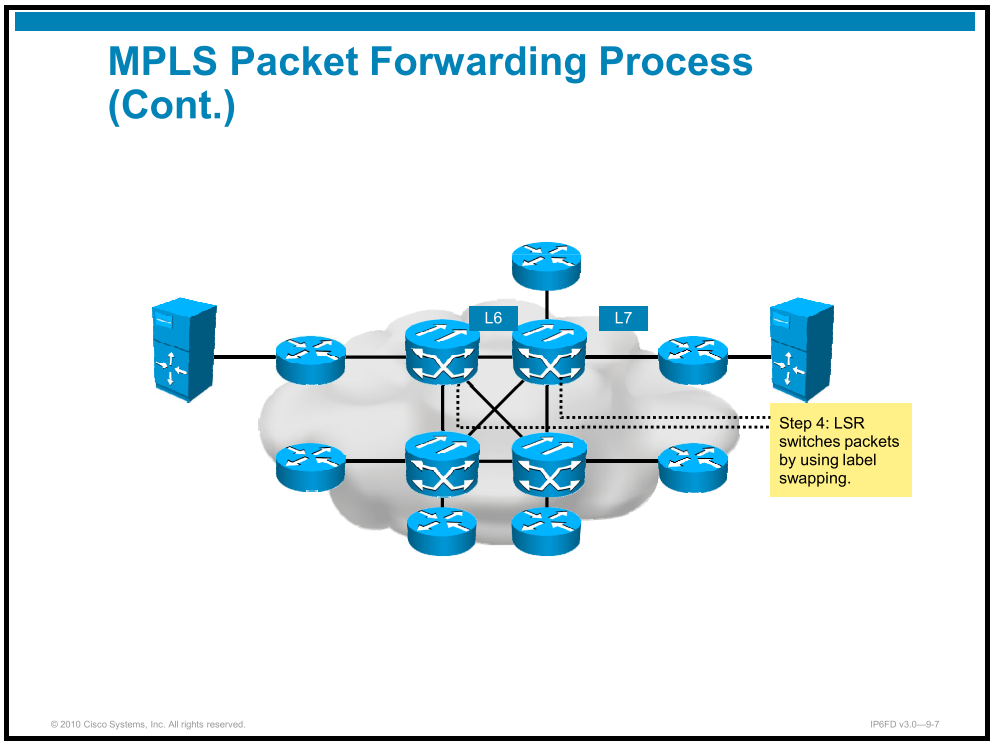

4. The LSR switches packets by using label swapping.

Core LSRs read the labels on each packet on their ingress interface and, based on the labels and the local label table lookup, forward the packets, with a new label, out the appropriate egress interface. This process is called label swapping.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

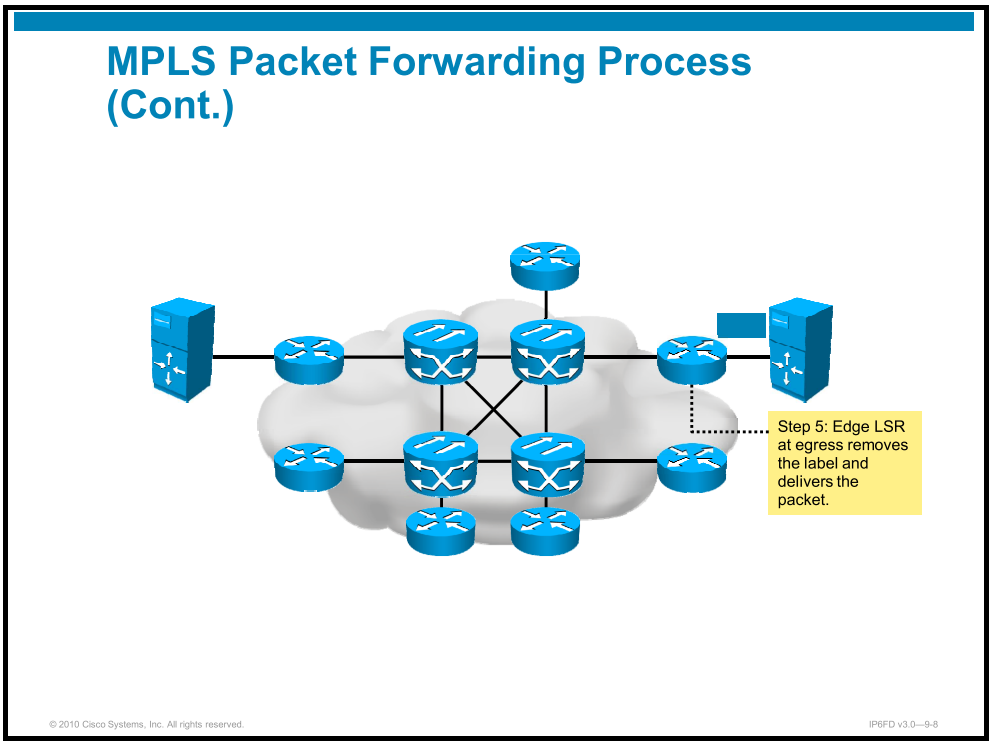

5. The edge LSR at egress removes the label and delivers the packet.

At the far edge of the MPLS network, the egress edge LSR strips the label and sends the packet towards its destination network.

https://t.me/learningnets

IPv6 over MPLS Deployment Scenarios

This topic describes the various methods that are used to deploy IPv6 over MPLS.

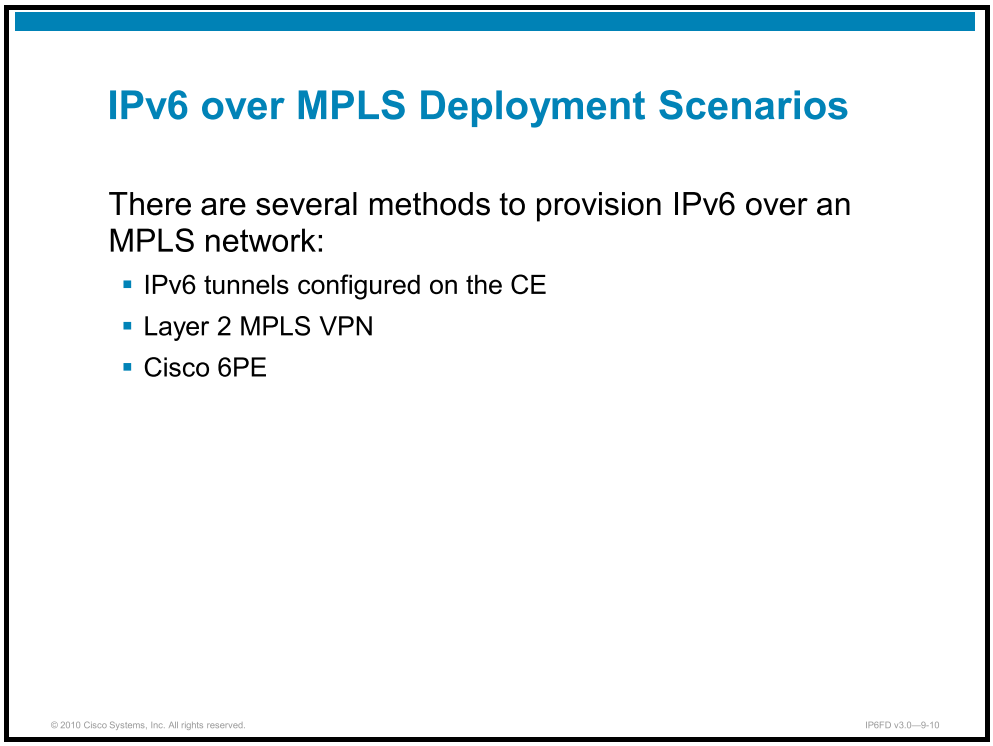

IPv6 over MPLS backbones enable isolated IPv6 domains to communicate with each other over an MPLS IPv4 network. Depending on the deployment scheme that is chosen and on the current MPLS implementation, enabling IPv6 services might require fewer backbone infrastructure upgrades and less reconfiguration of core routers than enabling IPv4 services, because forwarding is based on labels rather than on the IP header, thus providing a very cost-effective strategy for the deployment of IPv6.

Additionally, the inherent VPN and traffic engineering services that are available within an MPLS environment allow IPv6 networks to be combined into VPNs or extranets over an infrastructure that supports IPv4 VPNs and MPLS terminal equipment. A variety of deployment strategies are available or under development, as follows:

- Deploying IPv6 by using tunnels on the customer edge (CE) routers. This strategy has no impact on and requires no changes to the MPLS provider or provider edge (PE) routers. The strategy uses IPv4 tunnels that are terminated on the CE routers to encapsulate the IPv6 traffic, which then appears as IPv4 traffic within the MPLS provider network.
- Deploying IPv6 over a Layer 2 MPLS VPN. This strategy, applicable only on specific Cisco routers (such as the Cisco 12000 and 7600 Series Routers), also requires no changes to the core routing mechanisms, assuming that the provider is already capable of supporting Layer 2 MPLS VPNs.
- Deploying IPv6 on the PE routers via Cisco IPv6 Provider Edge Router over MPLS (Cisco 6PE). This strategy requires changes to the PE routers to support a dual-stack implementation, but all the core functions remain IPv4-only, and so the provider (P) routers require no changes.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

IPv6 Tunnels Configured on CE Routers

This topic describes an IPv6 tunnel configuration scenario on CE routers.

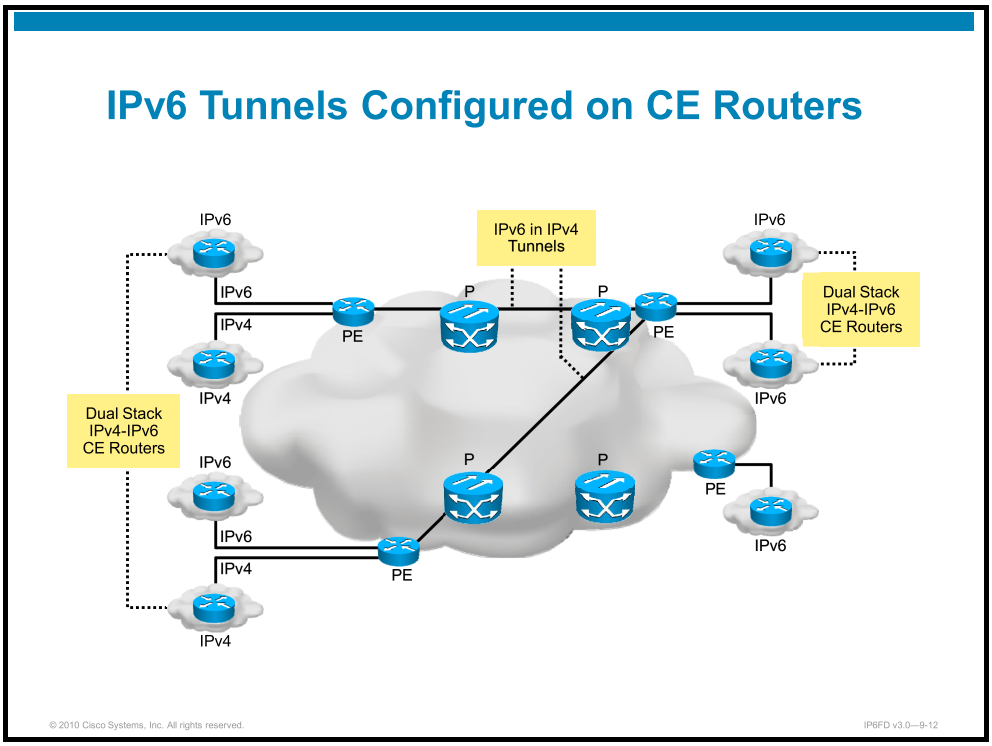

Using tunnels on the CE routers is the simplest way to deploy IPv6 over MPLS networks. This method has no impact on the operation or infrastructure of MPLS and requires no changes to the P routers in the core or to the PE routers that connect to customers.

IPv6 tunnels that are configured on CE have these characteristics:

- CE routers are IPv6-aware (dual-stack).
- Mesh of IPv6 over IPv4 tunnels connects CE to CE.
- Overhead includes IPv4 header and MPLS header.
- MPLS and VPN support native IPv4 and IPv6 tunnels.

Communication between the remote IPv6 domains uses standard tunneling mechanisms, running IPv6 over IPv4 tunnels like the way that MPLS VPNs support native IPv4 tunnels. The CE routers need to be upgraded to dual stack and configured by using manually configured or 6to4 tunnels, but communication with the PE routers is IPv4-only, and the traffic appears to be IPv4 to the MPLS domain. The dual-stack routers use the 6to4 tunnel addresses or an IPv6 prefix that is assigned by a virtual IPv6 ISP (possibly other than the MPLS provider). The figure shows an example for the deployment of IPv6 by using tunnels on the CE routers.

This alternative is attractive because it has no impact on the MPLS provider or PE routers, because it uses IPv4 tunnels to encapsulate IPv6 traffic. After being encapsulated, the traffic appears as IPv4 traffic within the network. However, this option does not permit the delegation of a /48 prefix address from the MPLS service provider address space to customers. Nonetheless, this alternative is probably useful for specific applications, such as adding IPv6 services to an MPLS VPN by configuring the tunnel on the CE. All Cisco IOS Software tunneling mechanisms support this capability.

https://t.me/learningnets

Pros

From the provider standpoint, using tunnels on CE routers is the simplest way to deploy IPv6 over MPLS networks. This approach has no impact on the operation or infrastructure of MPLS and requires no changes to either the P routers in the core or the PE routers that connect to customers.

Cons

Using tunnels on CE routers presents the same problems as any configured tunnel scheme. As the number of tunnel partners grows, the number of tunnels in an any-to-any mesh increases quickly, making management of the tunnels error-prone and expensive. This problem is not MPLS-specific but rather is a general issue with manually configured tunnels. Automatic tunnels (such as 6to4) can be used to reduce the amount of effort that is required to manage tunnels on the CE, but automatic tunnels have other issues that must be considered.

Delegating a global IPv6 prefix for an ISP is also difficult because the ISP does not furnish the IPv6 connectivity; it merely provides transport. Even when the ISP does allocate an address, early traffic levels will be modest, and later traffic might shift from IPv4 to IPv6. Therefore, there is not necessarily a net increase in carried packets. Because the service provider provides transport only over IPv4 over MPLS, this scenario offers little additional revenue-generating possibilities for the provider.

Tip From the provider standpoint, using tunnels on CE routers is not a good mechanism to carry IPv6 traffic across an IPv4-based MPLS network. The provider basically puts responsibility for IPv6 transit on the end customers, because the MPLS provider recognizes only IPv4 traffic.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

IPv6 over Layer 2 MPLS VPN

This topic describes the pros and cons of using Layer 2 MPLS VPNs to deploy IPv6 over MPLS networks.

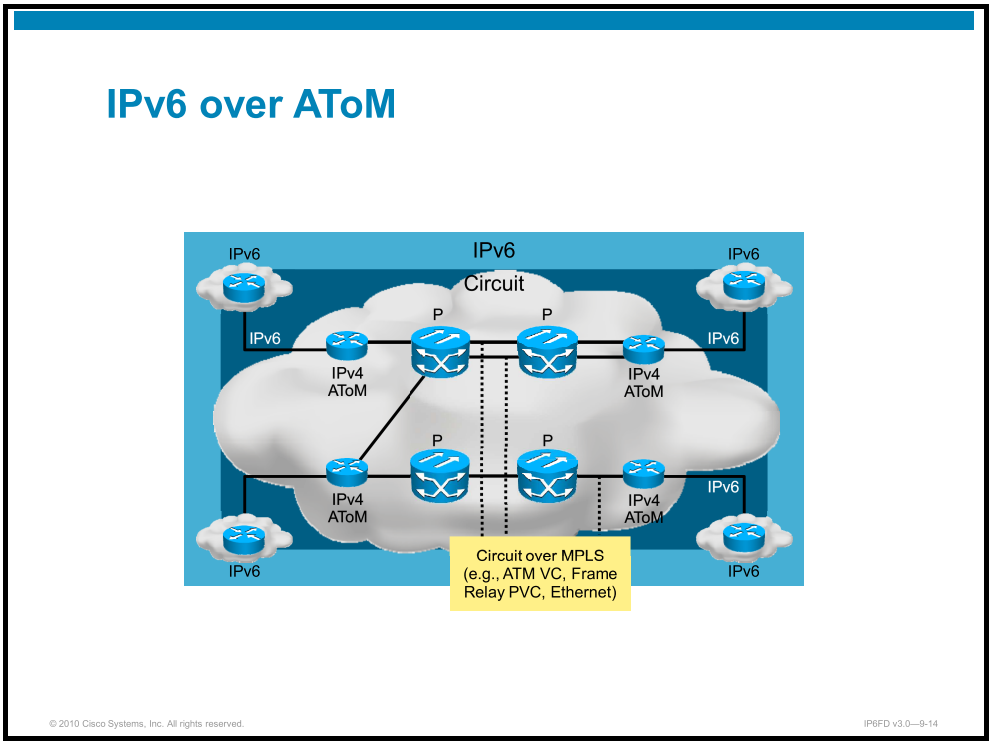

Using any circuit transport to deploy IPv6 over MPLS networks has no impact on the operation or infrastructure of MPLS. Circuit transport requires no changes to either the P routers in the core or the PE routers that connect to customers, assuming that the provider has already implemented circuit-over-MPLS technology. This is a Layer 2 solution.

IPv6 using a Layer 2 MPLS VPN has these characteristics:

- IPv6 run over pseudowires (see the Internet Engineering Task Force [IETF] document draft-ietf-pwe3-wildcard-pw-type-02.txt).
- Edge MPLS routers need to support a Layer 2 MPLS VPN.
- Mesh of virtual circuits connects PE to PE.
- PE routers are regular IPv4 routers that are used to aggregate customer IPv6 routers.

Communication between remote IPv6 domains runs native IPv6 protocols over a dedicated link, in which the underlying mechanisms are fully transparent to IPv6. The IPv6 traffic is tunneled by using a Layer 2 MPLS VPN or Ethernet over MPLS (EoMPLS), with the IPv6 routers connecting through an ATM or Ethernet interface, respectively. This means that the PE will require an upgrade to a Layer 2 MPLS VPN if it does not already contain that support. The figure shows an example of IPv6 deployment over any circuit transport over MPLS.

The IPv6 over Layer 2 MPLS VPN alternative is likely to be used by service providers that have ATM or Ethernet links to CE routers. This alternative is fully transparent to users. However, like the previous option, this option does not permit the delegation of a /48 prefix address from the service provider address space to customers. In addition, IPv6 over Layer 2 MPLS VPN does not scale as well as 6PE because service providers need to create circuits through their MPLS networks for their customers.

https://t.me/learningnets

Pros

These are some advantages of using IPv6 over a Layer 2 MPLS VPN:

- Simple solution for service providers with ATM or Ethernet links to CE routers, and Circuit over MPLS capability
- Fully transparent IPv6 communication, from the provider perspective
- No required reconfiguration of the MPLS core

Cons

These are some disadvantages of using IPv6 over a Layer 2 MPLS VPN:

- Layer 2 VPN complexity equivalent
- Upgrading PE routers that connect to customers to a Layer 2 MPLS VPN
- Difficult to delegate a global IPv6 prefix for an ISP

Description of Pros and Cons

These are some additional points regarding IPv6 over Layer 2 MPLS VPN:

- The use of a Layer 2 MPLS VPN to deploy IPv6 over MPLS networks does not affect the operation or infrastructure of MPLS. Neither the P router software version nor the configuration is changed. However, PE routers that connect to customers must be upgraded to Layer 2 MPLS VPN. Delegating a global IPv6 prefix for an ISP is also difficult.
- Communication between remote IPv6 domains runs native IPv6 protocols over a dedicated link in which the underlying mechanisms are fully transparent to IPv6. The IPv6 traffic is tunneled via a Layer 2 MPLS VPN, with the IPv6 routers connected through an ATM or Ethernet interface, respectively. This potential heavy tunnel meshing is a drawback of the solution.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Cisco 6PE

This topic describes Cisco 6PE.

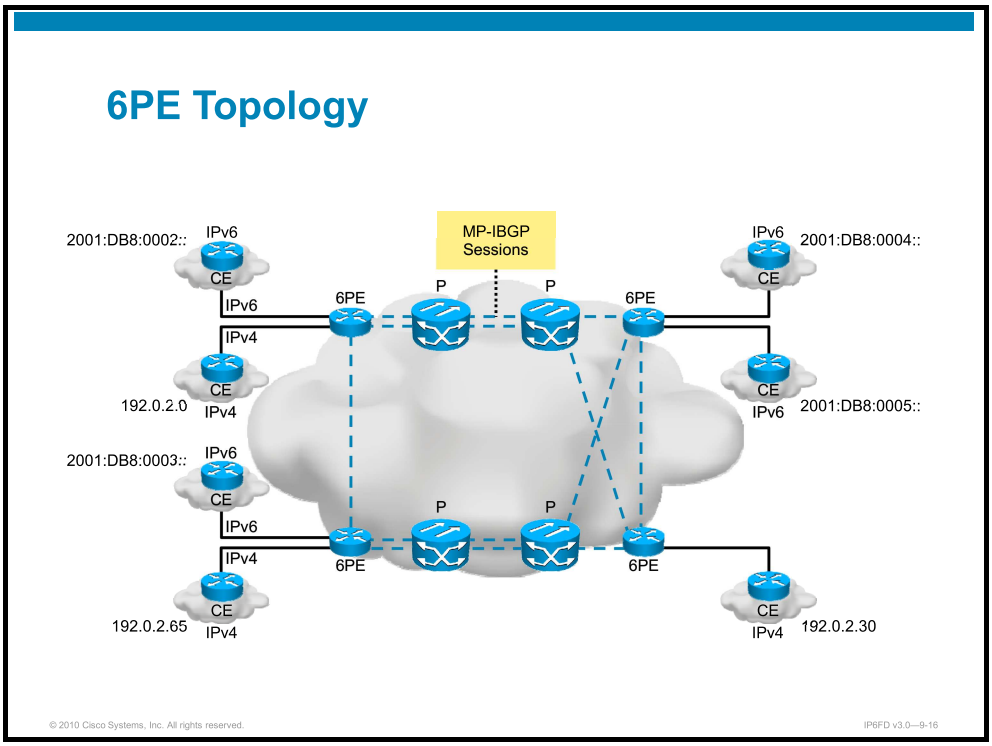

Cisco 6PE enables IPv6 islands to communicate with each other over an MPLS IPv4 core network by using MPLS LSPs. The method relies on Border Gateway Protocol (BGP) extensions in the IPv4 network PE routers (Cisco 6PE routers) to exchange IPv6 reachability information and an MPLS label for each IPv6 address prefix that is announced. Cisco 6PE routers are dual-stack (IPv4 and IPv6). This is a Layer 3 solution.

For the IPv6 transport to be transparent to all but Cisco 6PE routers, you must impose a hierarchy of labels at the Cisco 6PE ingress router. The top label provides connectivity inside the IPv4 MPLS core network. LDP, Tag Distribution Protocol (TDP), or Resource Reservation Protocol (RSVP) distributes this label. The inner label is used for IPv6 forwarding at each Cisco 6PE egress router. This label is distributed by Multiprotocol Border Gateway Protocol (MP-BGP) in the VPN-IPv6 address family.

https://t.me/learningnets

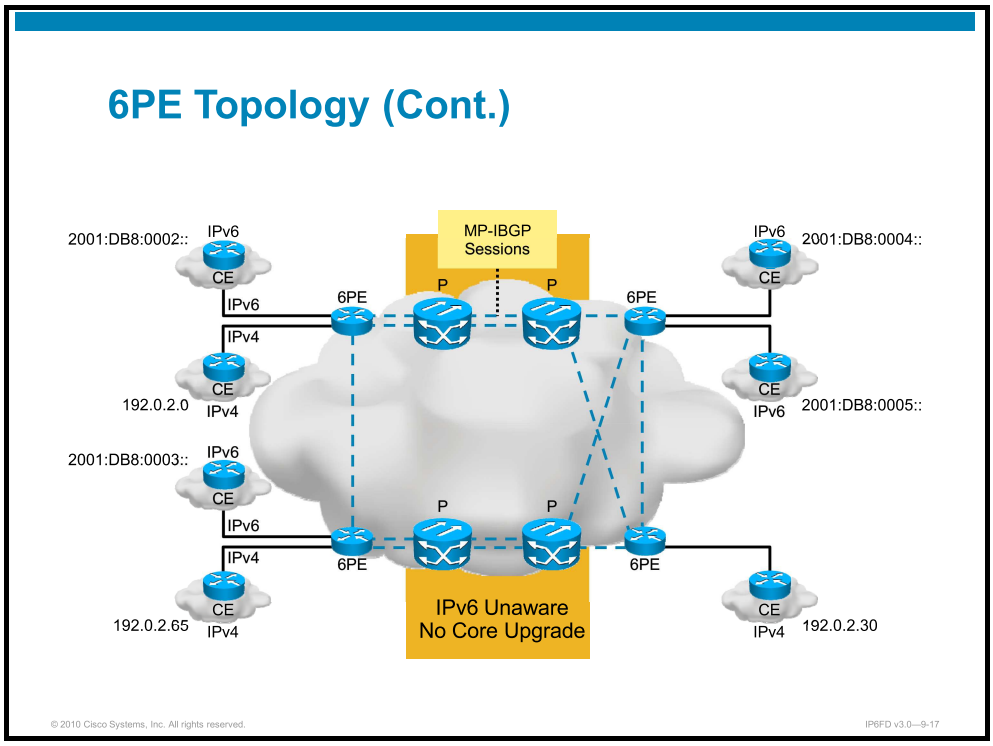

Cisco 6PE uses MPLS to minimize changes in the ISP core network.

With Cisco 6PE, IPv6 forwarding in the core of the MPLS network is done by label switching, eliminating the need for either IPv6-over-IPv4 tunnels or for an additional Layer 2 encapsulation and allowing the appearance of a native IPv6 service to be offered across the network.

Each PE router that must support IPv6 connectivity needs to be upgraded to dual stack (becoming a Cisco 6PE router) and configured to run MPLS on the interfaces that connect to the core. Depending on the site requirements, each router can be configured to forward IPv6 or IPv6 and IPv4 traffic on the interfaces to the CE routers, thus providing the ability to offer only native IPv6 or both IPv6 and native IPv4 services. The Cisco 6PE router exchanges either IPv4 or IPv6 routing information with CE router through any of the supported routing protocols (depending on the connection) and switches IPv4 and IPv6 traffic over the native IPv4 and IPv6 interfaces that do not run MPLS.

The Cisco 6PE router uses MP-BGP to exchange reachability information with the other 6PE routers in the MPLS domain, and shares a common IPv4 routing protocol (such as OSPF or IS-IS) with the other P and PE devices in the domain. IGP is used for deriving next hop information.

The Cisco 6PE routers encapsulate IPv6 traffic by using two levels of MPLS labels. The top label is distributed by an LDP or TDP that devices in the core use to carry the packet to the egress Cisco 6PE by using IPv4 routing information. The second or bottom label is associated with the IPv6 prefix of the destination through MP-BGP version 4 (MP-BGP4).

Multiprotocol Extensions for BGP4

Multiprotocol extensions for BGP4 are characterized as follows:

- BGP4 is the primary mechanism to exchange IPv6 reachability information for Cisco 6PE and Cisco IPv6 VPN Provider Edge Router over MPLS (6VPE).
- Label binding information is piggybacked with the prefix advertisement in a Network Layer Reachability Information (NLRI) attribute.

https://t.me/learningnets

- BGP protocol extensions, which carry routing information about other protocols, include multicast, MPLS, IPv6, VPN-IPv4, labeled IPv6 unicast (Cisco 6PE), and Cisco 6VPE.
- The exchange of multiprotocol NLRI must be negotiated at the session setup. The exchange uses BGP capabilities advertisement negotiation procedures.

MP-BGP is the supported exterior gateway protocol (EGP) for IPv6. MP-BGP extensions for IPv6 support the same features and functionality as IPv4 BGP. IPv6 enhancements to MP-BGP include support for an IPv6 address family, NLRI, and next-hop attributes that use IPv6 addresses.

The Cisco 6PE multipath feature uses Multiprotocol Internal BGP (MP-IBGP) to distribute IPv6 routes over the MPLS IPv4 core network and to attach an MPLS label to each route.

Note See RFC 2858, *Multiprotocol Extensions for BGP4* for more information. (RFC 2283 is obsolete.)

https://t.me/learningnets

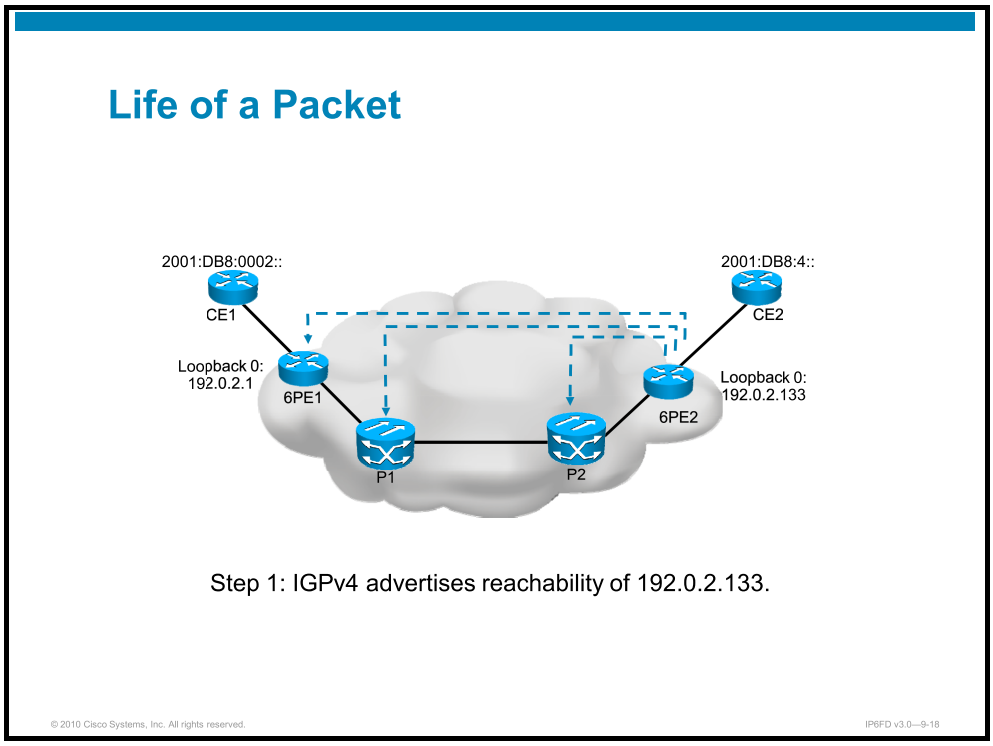

Life of a Packet (Control Plane)

This example describes how IPv6 transport is achieved across an IPv4-based MPLS provider network that is running an IPv4 IGP such as OSPF version 2 (OSPFv2), and how reachability is established between Cisco 6PEs.

Step 1 IGPv4 advertises the reachability of PE routers.

The Cisco 6PE routers exchange IPv6 routing information through MP-BGP. The Cisco 6PE routers peer together through MP-BGP4, to exchange IPv6 reachability with the MPLS cloud and to perform IPv6 forwarding. IGP advertises the reachability of the BGP next-hop address; that is, 192.0.2.133.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Life of a Packet (Cont.)

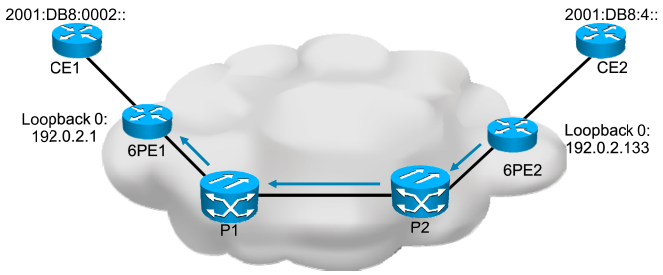

Step 2: LDPv4 binds label to 192.0.2.133.

© 2010 Cisco Systems, Inc. All rights reserved.

Step 2 LDPv4 binds labels to create LSP.

Life of a Packet (Cont.)

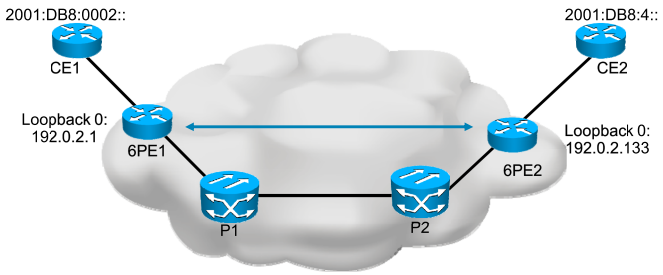

Step 3: MP-BGP advertises 2001:DB8:4:: and binds a (second-level) label. IPv6 next hop is an IPv4-mapped IPv6 address built from 192.0.2.133.

© 2010 Cisco Systems, Inc. All rights reserved.

Step 3 The 6PE2 router sends MP-BGP advertisements to the ingress PE router, 6PE1.

https://t.me/learningnets

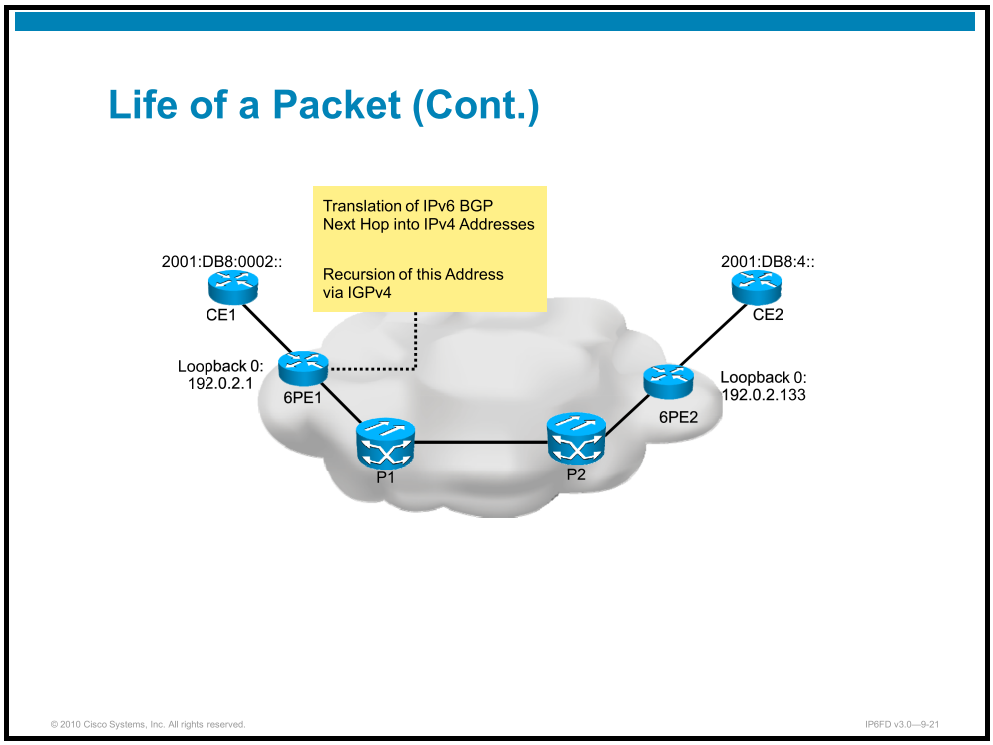

Step 4 The BGP next hop for the destination IPv6 prefix is set to an IPv4-compatible address, which is built from the 6PE2 router.

The MPLS label binds to the IPv4- or IPv6-mapped address, which is formed from the IPv4 address of the 6PE2 (192.0.2.133). Thus, other Cisco 6PE routers (which are part of the mesh edges around the MPLS core) see 6PE2 ::FFFF:192.0.2.133 as the next-hop IPv6 address, when a packet needs to be forwarded to the 2001:DB8:4:: network.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

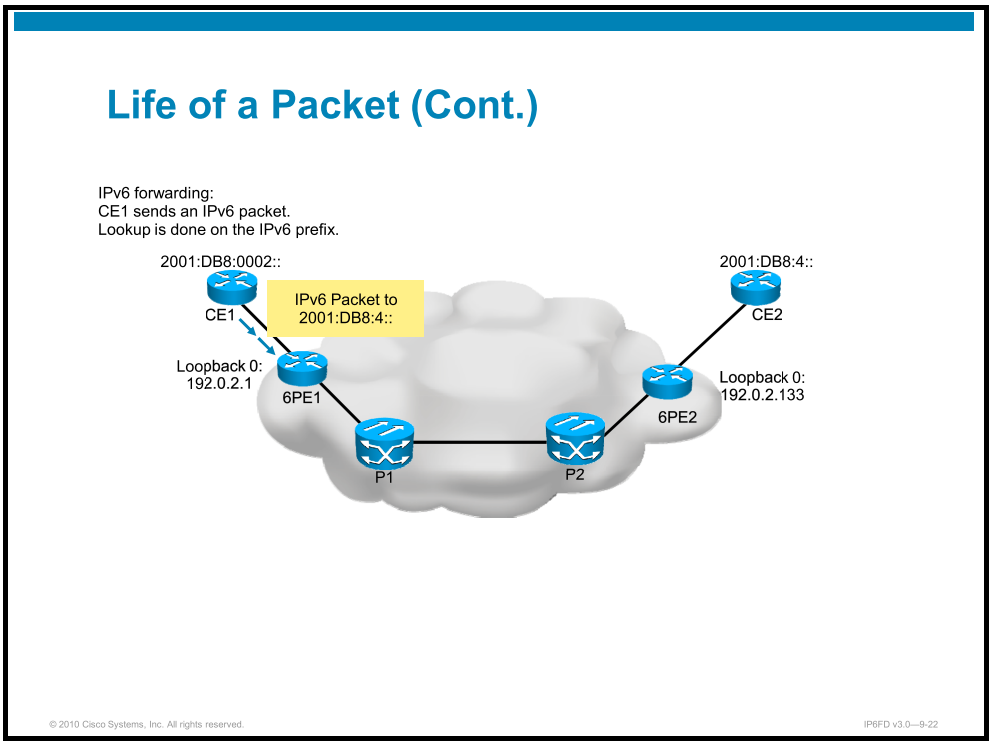

Life of a Packet (Data Plane)

Step 1 An IPv6 packet arrives at the ingress Cisco 6PE router.

An IPv6 packet that is destined to a network with prefix 2001:DB8:4:: enters through 6PE1. The MPLS core needs to be used to carry the packet to its destination, via 6PE2. The routing configuration via IBGP has already advertised the next-hop IPv6 address (::FFFF:192.0.2.133) to 6PE1.

https://t.me/learningnets

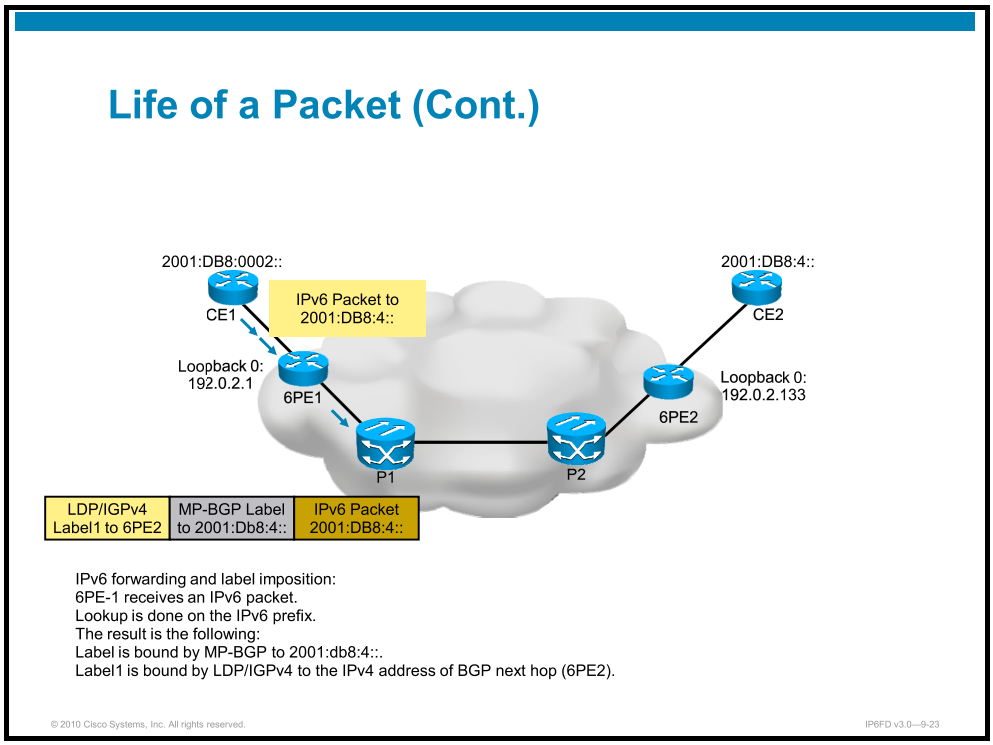

Step 2 The ingress Cisco 6PE router binds MP-BGP (inner) and LDP/IGPv4 (top, or outer) labels to the packet.

6PE1 does a lookup on the IPv6 prefix. Two labels are then created and bound to the packet. MP-BGP binds 2001:DB8:4:: and LDP/IGPv4 binds its label to the IPv4 address of 6PE2 (192.0.2.133). In this way, MP-BGP can carry either IPv4 packets or IPv6 packets from one Cisco 6PE router to another.

https://t.me/learningnets

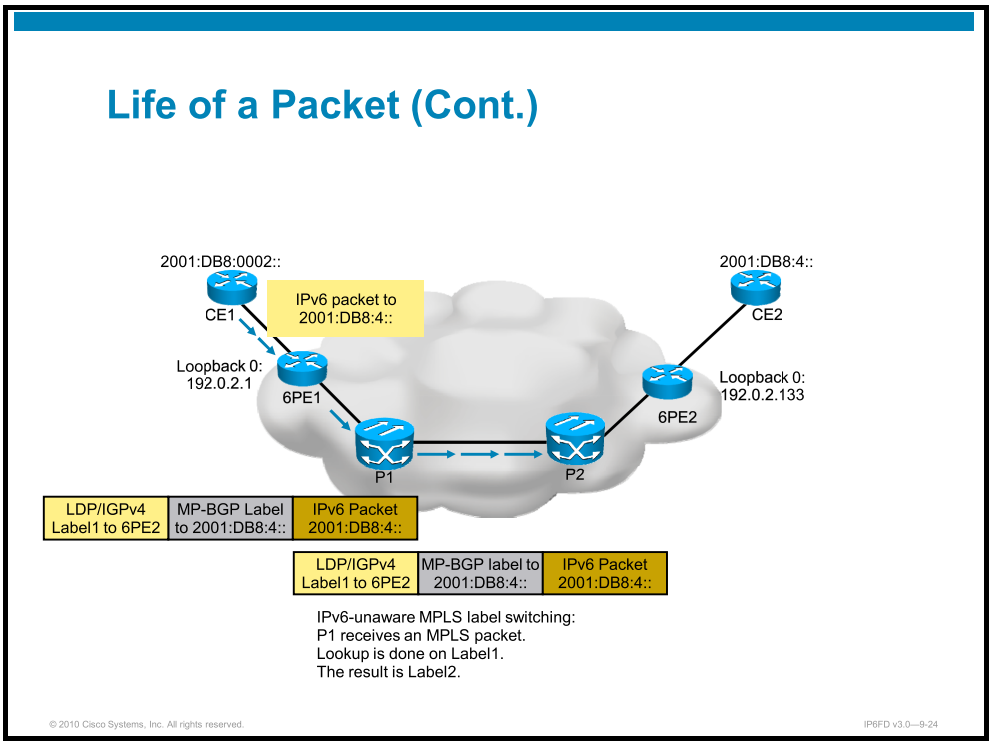

Step 3 The IPv6 packet is forwarded to the first P router.

MPLS label switching is IPv6-unaware. When P1 receives the MPLS packet, it looks only at the LDP-distributed MPLS label, and a lookup is performed. The result is the imposition of Label2, and the MPLS packet is forwarded to P2. The IPv6 traffic is transparent to the core of the MPLS network.

https://t.me/learningnets

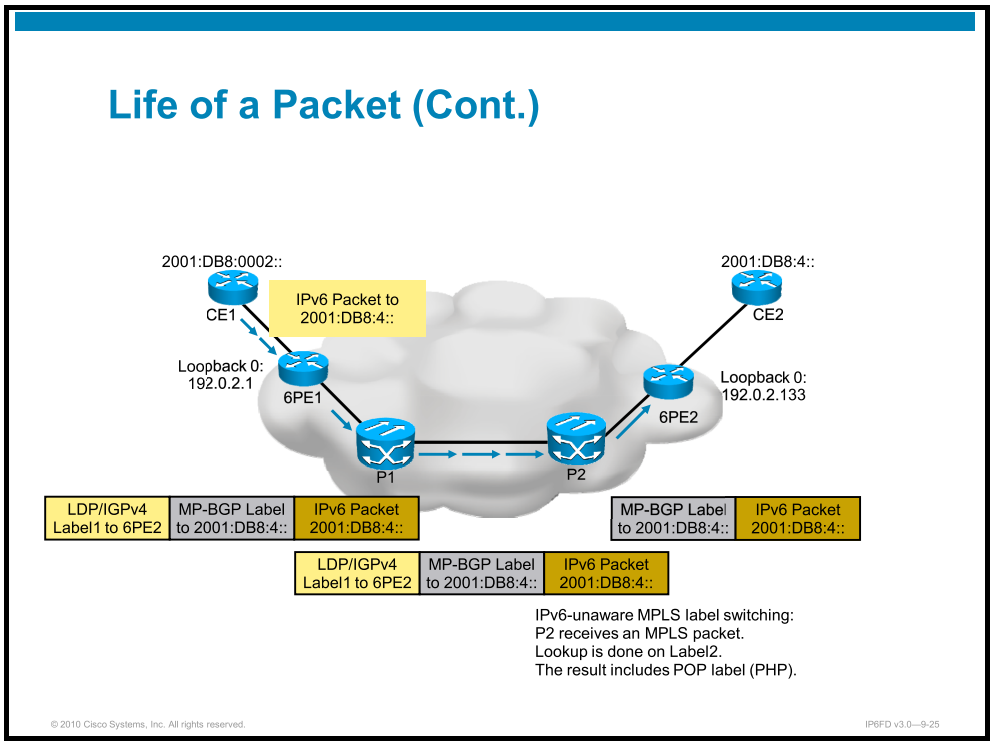

Step 4 The IPv6 packet is forwarded to the next P router.

P2 receives an MPLS packet, which performs an MPLS lookup on Label2 and results in the point of presence (POP) label. This lookup results in the forwarding of the packet to the IPv4 address to which it is bound. The packet still has the inner MPLS label, which is required because a raw IPv6 packet is not allowed within the MPLS core.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

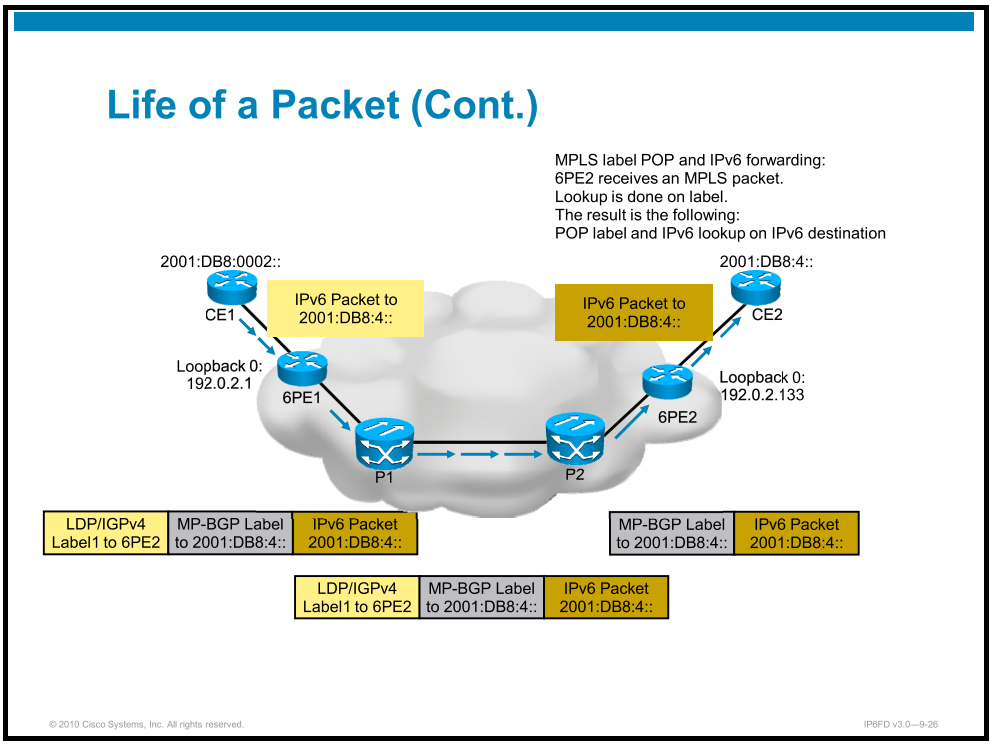

Step 5 Egress PE routers remove MPLS labels and forward the IPv6 packet.

The 6PE2 router examines the MPLS packet, and the result is that the remaining MPLS label is removed and an IPv6 lookup on the IPv6 destination address, using normal IP destination-based routing, is performed. The Cisco 6PE router forwards the IPv6 packet on to the CE device, based on the results of the IPv6 lookup on the IPv6 destination address.

Tip This step reveals the reason for the MPLS label stack, in the absence of MPLS VPNs. If the MPLS network uses penultimate hop popping (PHP), in which the next-to-last router in the transit path removes the MPLS label, the un-encapsulated packet must be an IPv4 packet, because the core of the MPLS network understands only IPv4 and MPLS. Because the inner packet here is an IPv6 packet, it must remain inside an MPLS packet (actually, behind an MPLS label) until it reaches the Cisco 6PE router.

https://t.me/learningnets

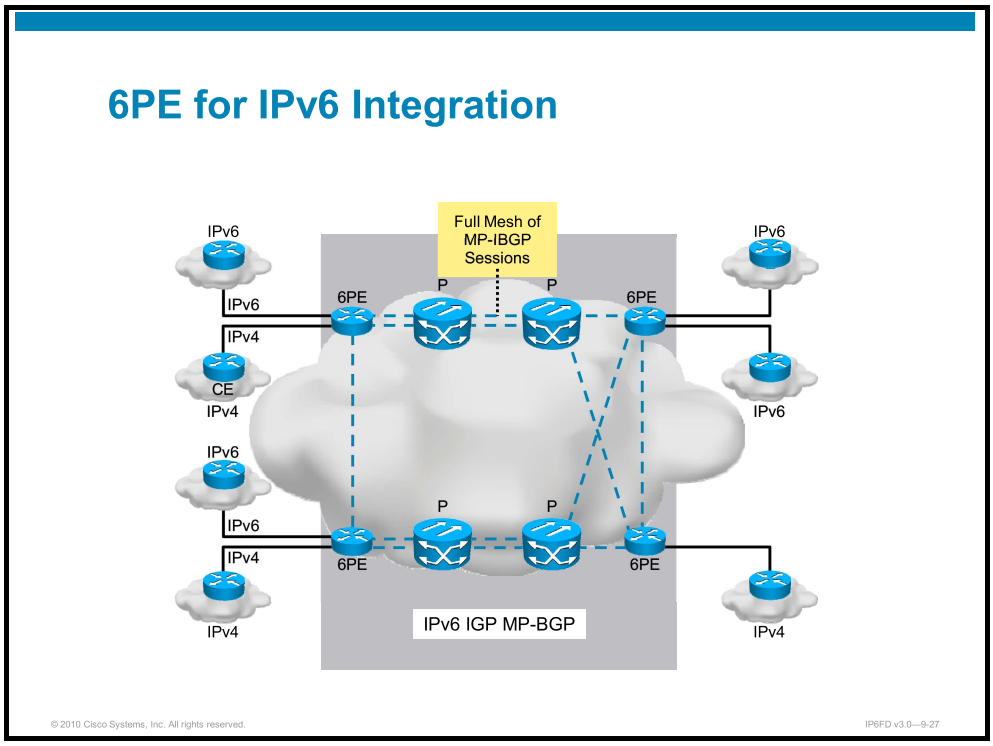

Benefits

Using Cisco 6PE for IPv6 integration has these benefits:

- IPv6 CE has only one PE routing peer, regardless of how many remote IPv6 CEs it communicates with.
- No changes are required on an IPv6 CE when remote CEs are added or removed. Reachability is automatically learned via MP-BGP.
- No tunnel or circuit configuration is required.
- Cisco 6PE offers a scalable and flexible solution. Benefits are analogous to the Layer 3 VPN solution for IPv4, in RFC 2547bis.

Service providers that already deploy MPLS, or plan to do so, can obtain these benefits from Cisco 6PE:

- There is minimal operational cost and risk, and no impact on existing IPv4 and MPLS service, because only PE routers are upgraded. A Cisco 6PE router can be an existing PE router or a new router that is dedicated to IPv6 traffic.
- There is no impact on IPv6 CE routers. The ISP can connect to any customer CE that runs static, IGP, or EGP ready-for-production services. An ISP can delegate IPv6 prefixes to the customers.
- Nondisruptive IPv6 introduction into an existing MPLS service Cisco 6PE router can be added at any time.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Drawbacks

These are some of the drawbacks to using Cisco 6PE for integration:

- Makes sense only when the network already runs MPLS
- Requires knowledge of MPLS and BGP technologies
- Requires dual-stack and software upgrade on existing PE, or deployment of dedicated Cisco 6PE routers

https://t.me/learningnets

How to Deploy Cisco 6PE on MPLS Networks

This topic describes how Cisco 6PE is deployed in an MPLS network.

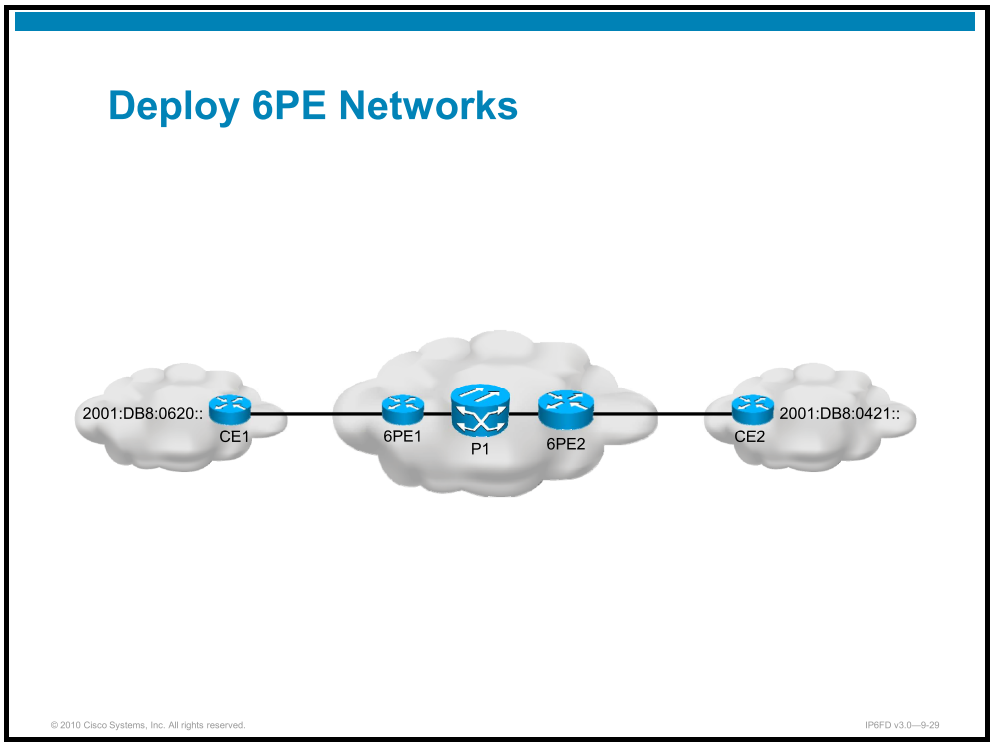

Follow these steps to deploy Cisco 6PE:

Step 1 Configure Cisco 6PEs as dual-stack routers.

Step 2 Configure Cisco 6PEs to be part of the IPv4 network and configure MPLS.

Step 3 Specify the interface from which locally generated packets take their source IPv6 addresses.

Step 4 Bind and advertise aggregate labels.

Example

Using the network that is shown in the figure, two configuration tasks are required at the 6PE1 router to enable the Cisco 6PE feature.

The CE router, CE1, is configured to forward its IPv6 traffic to the 6PE1 router. The P1 router in the core of the network is assumed to be running MPLS, a label distribution protocol, an IPv4 IGP, and Cisco Express Forwarding or distributed Cisco Express Forwarding. P1 does not require any new configuration to enable the Cisco 6PE feature. New configuration tasks are also not required for the CE1 router.

The Cisco 6PE routers, 6PE1 and 6PE2, must be members of the core IPv4 network. The Cisco 6PE router interfaces that are attached to the core network must run MPLS, the same LDP, and the same IPv4 IGP as in the core network.

The 6PE routers must also be configured to be dual stack to run both IPv4 and IPv6.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Restrictions

These restrictions apply when implementing the Cisco 6PE feature:

- Core MPLS routers support MPLS and IPv4 only, so they cannot forward or create any IPv6 Internet Control Message Protocol (ICMP) messages.
- Cisco 6PE does not provide load balancing between an MPLS path and an IPv6 path. If both are available, the MPLS path is always preferred. Load balancing between two MPLS paths is possible.

When MP-IBGP multipath is enabled on the Cisco 6PE router, all labeled paths are installed in the forwarding table with MPLS information (label stack), when MPLS available. This functionality enables Cisco 6PE to perform load balancing.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in the lesson.

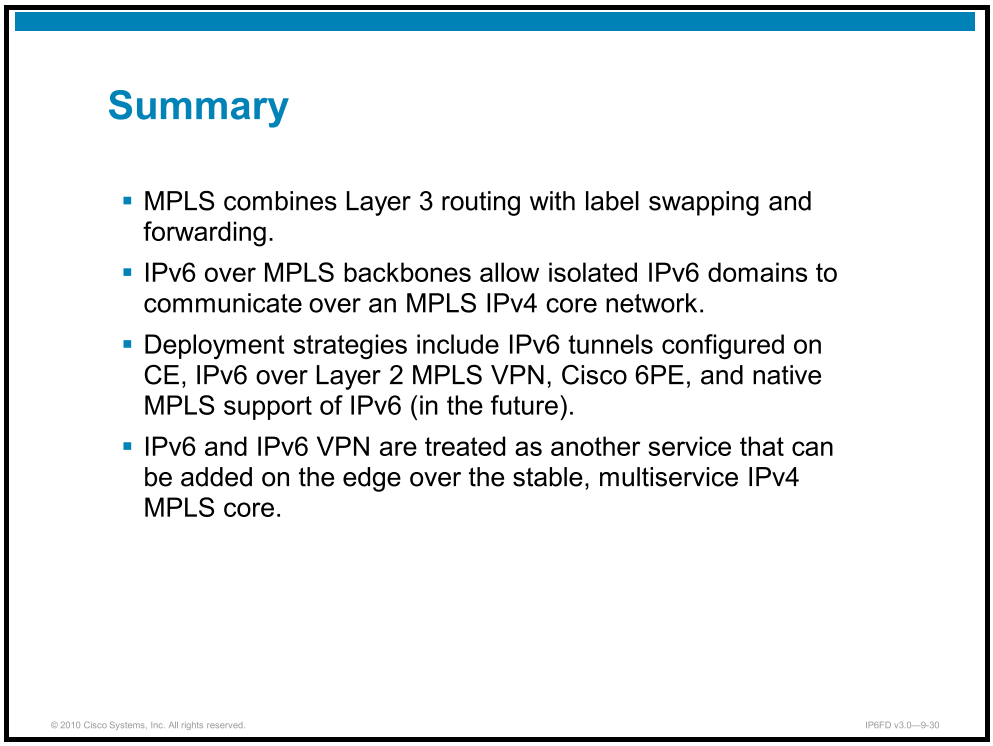

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3

Understanding 6VPE

Overview

IP version 6 (IPv6) Virtual Provider Edge Router (6VPE) over Multiprotocol Label Switching (MPLS) network facilitates the RFC 2547bis-like virtual private network (VPN) model for IPv6 networks. 6VPE is more like a regular IP version 4 (IPv4) MPLS VPN provider edge, with the addition of IPv6 support within Virtual Routing and Forwarding (VRF). It provides logically separate routing table entries for VPN member devices.

Objectives

Upon completing this lesson, you will be able to explain the 6VPE service that service providers offer. This includes being able to meet these objectives:

- Describe how Cisco 6VPE is deployed in an MPLS network
- Implement Cisco 6VPE

https://t.me/learningnets

Cisco 6VPE

This topic describes how Cisco 6VPE is deployed in an MPLS network.

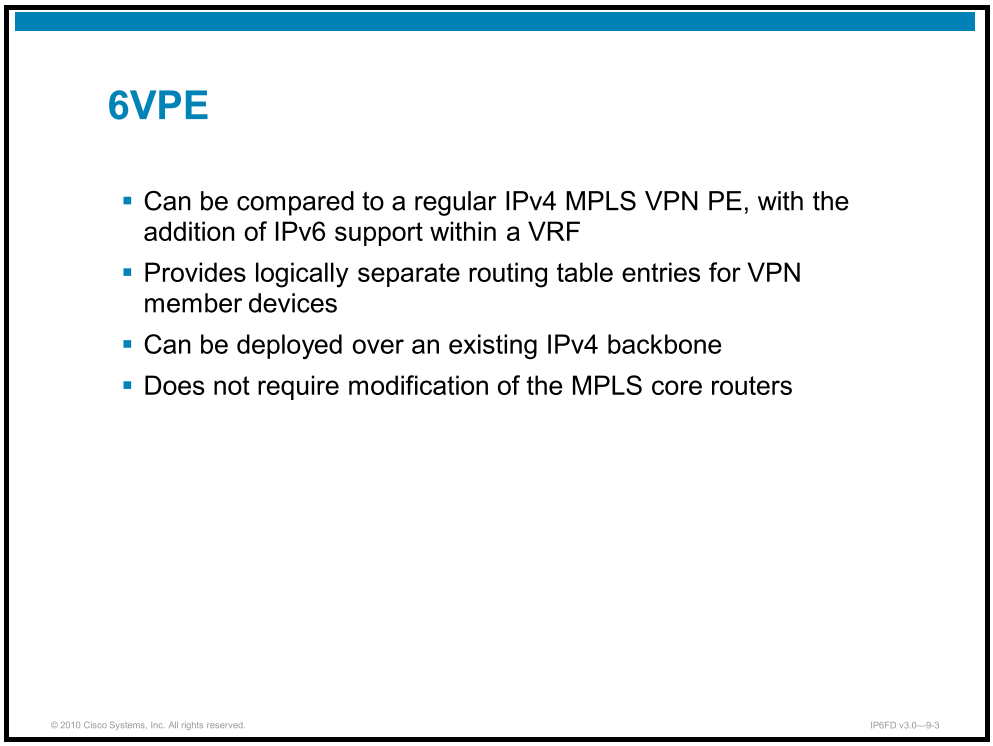

IPv6 MPLS VPNs are like IP version 4 (IPv4) MPLS VPNs. Customer sites are connected over the shared MPLS backbone so that each customer has reachability to their own VPN partner sites only. As with IPv4 MPLS VPNs, this outsourced intranet remains private and is seen by the customer as a closed network that runs on a dedicated network.

The inherent VPN and terminal equipment services that are available within an MPLS environment allow IPv6 networks to be combined into VPNs or extranets over an infrastructure that supports IPv4 VPNs and MPLS terminal equipment.

For ISPs that offer MPLS VPNs for IPv4, use Cisco 6VPE to add IPv6 services. Using this approach has these benefits:

- There is no modification on the MPLS core.
- Cisco 6VPE supports both IPv4 and IPv6 VPNs concurrently on the same interfaces.
- Configuration and operation of IPv6 VPNs are the same as for IPv4 VPNs.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

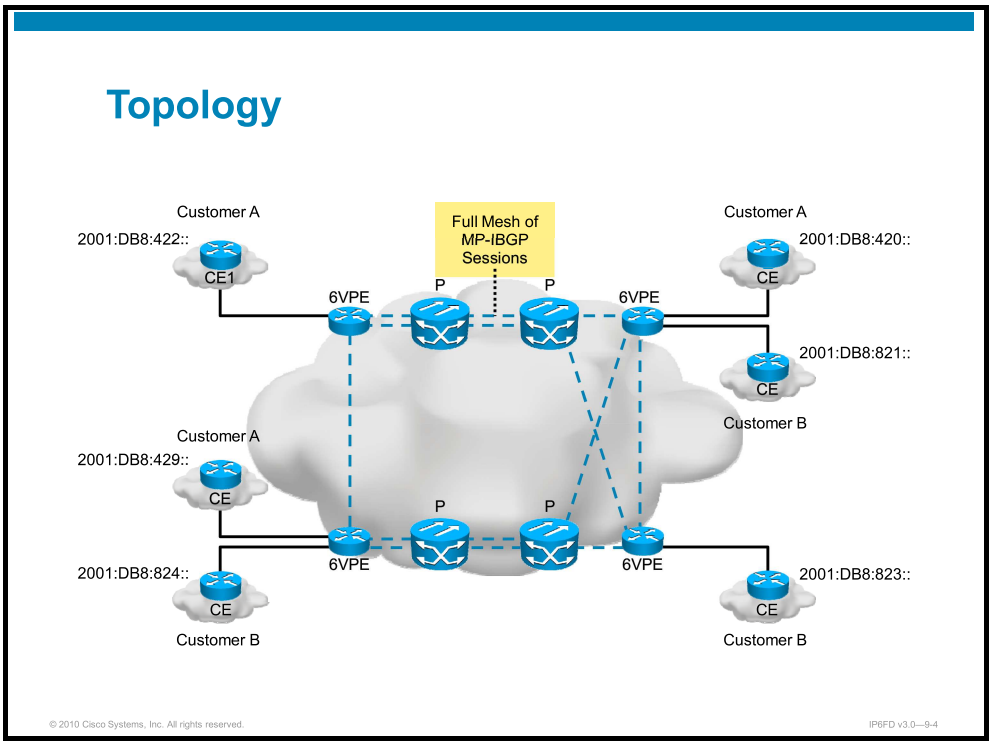

Service providers who offer MPLS IPv4 VPN services to their customers might look forward to adding IPv6 VPN services to their portfolios. A VPN is said to be an IPv6 VPN when a customer edge (CE) router turns on native IPv6 over an interface or subinterface to the provider edge (PE) router. Cisco 6VPE adds IPv6 VPN capability to Cisco IPv6 Provider Edge Router over MPLS (6PE) and enables an ISP to deliver similar services to IPv4.

As in IPv4 VPN route distribution, Border Gateway Protocol (BGP) and its extensions are used to distribute routes from an IPv6 VPN site to all other Cisco 6VPE routers that connect to a site of the same IPv6 VPN. Cisco 6PE routers use virtual routing and forwarding (VRF) tables to separately maintain the reachability and forwarding information of each IPv6 VPN (as with IPv4 VPN).

Standardization

The Internet Engineering Task Force (IETF) standardization defines a specification that lets a service provider use an MPLS-enabled IPv4 backbone to provide VPNs for its IPv6 customers. The standard method is based on VPNs as described in RFC 4364, *BGP/MPLS IP VPNs*. In a BGP/MPLS VPN, MPLS is used to forward packets over the backbone, and Multiprotocol Border Gateway Protocol (MP-BGP) is used to distribute VPN routes over the service provider backbone.

MP-BGP allows generic operations over both IPv4 and IPv6. For example, MP-BGP defines an IPv6 VPN address family and describes the route distribution and MPLS tunnel selection.

https://t.me/learningnets

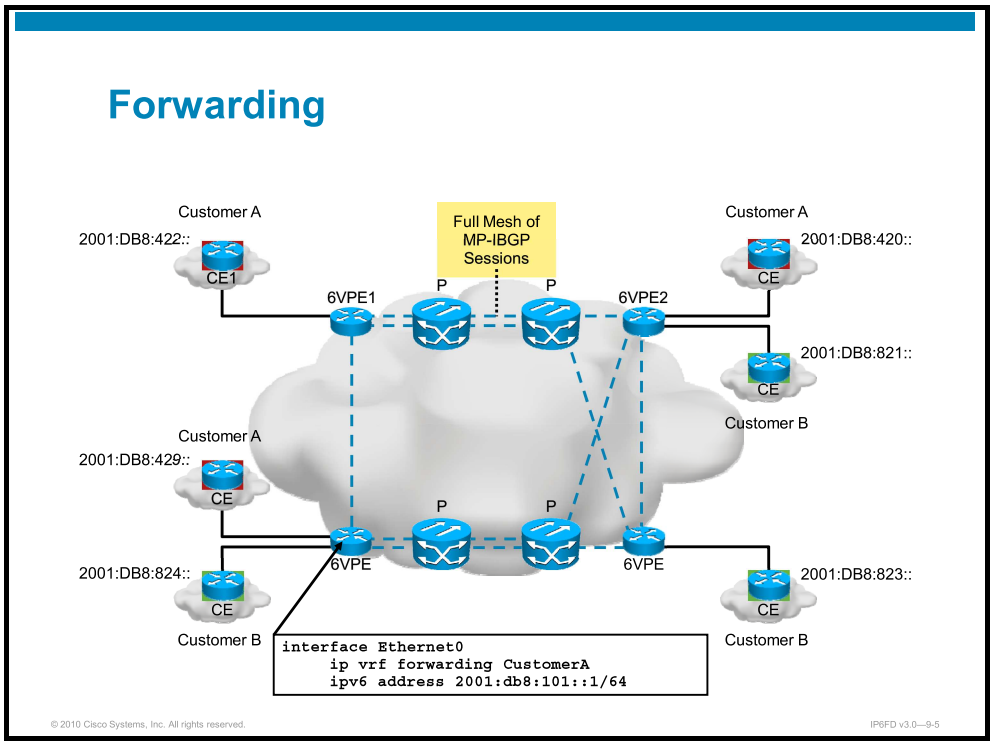

Forwarding in Cisco 6VPE follows these steps:

1. The ingress PE router receives an IPv6 packet and attaches an MPLS label.

When a Cisco 6VPE router receives an IPv6 packet from CE1, the router looks up the IPv6 packet destination address in the VRF table. This address enables 6VPE1 to find a VPN IPv6 route. The VPN IPv6 router has an associated MPLS label and an associated BGP next hop. This MPLS label is imposed on the IPv6 packet.

2. The ingress PE router attaches a second label to the IPv6 packet.

6VPE1 pushes another label, which is bound by Label Distribution Protocol (LDP) interior gateway protocol (IGP), to the IPv4 address of the BGP next hop, to reach 6VPE2 through the MPLS cloud on the label stack of the labeled IPv6 VPN packet. This top-most imposed label corresponds to the LSP that starts on 6VPE1 and ends on 6VPE2. The bottom label is bound to the IPv6 VPN prefix via BGP. All the provider (P) routers in the backbone network switch the VPN packet, based only on the top label in the stack, which points toward the 6VPE2 router. Because of the normal MPLS forwarding rules, the P routers never look beyond the first label and are thus completely unaware of the second label or the IPv6 VPN packet that is carried across the backbone network.

3. The egress router receives the labeled IPv6 packet, removes the labels, and forwards the packet to the CE router.

The egress PE router, 6VPE2, receives the labeled IPv6 VPN packet, drops the first label, and performs a lookup on the second label, which uniquely identifies the target VRF table and sometimes even the outgoing interface on 6VPE2. A lookup is performed in the target VRF table, and the IPv6 packet is sent toward the proper CE router in the IPv6 domain or site.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Configuring 6VPE

This subtopic describes configuration steps for implementing 6VPE.

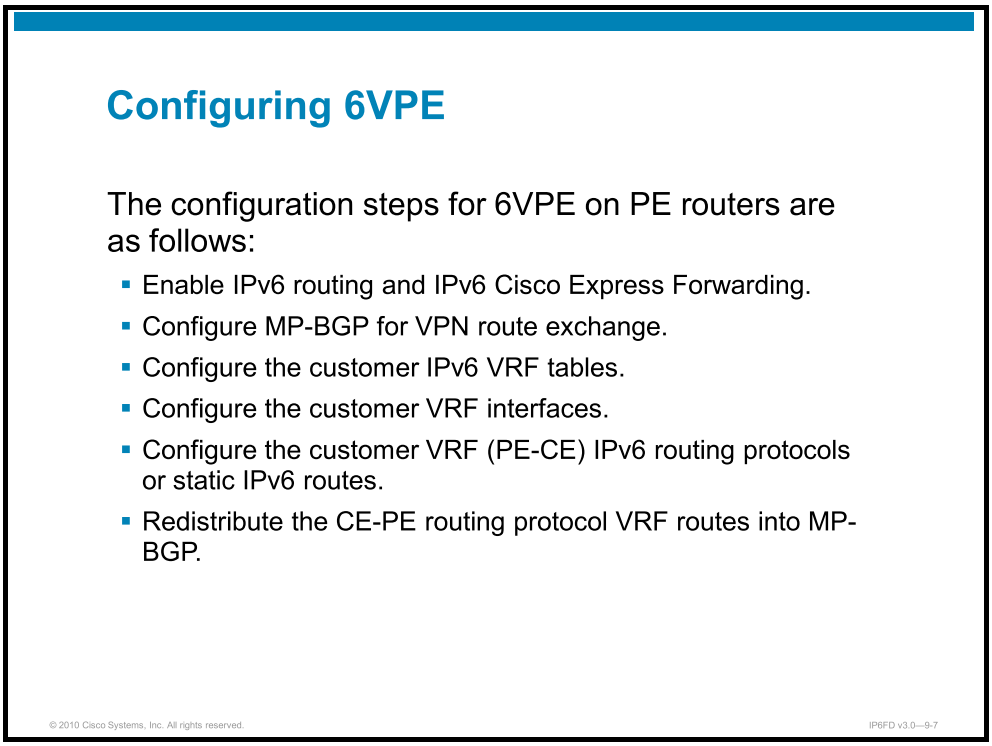

To successfully implement 6VPE, you must apply all the steps that the figure shows.

This lesson assumes that these steps have already been completed:

- Loopback interface is configured.
- LDP is configured.
- MPLS is enabled on interfaces.
- Backbone IGP (Open Shortest Path First [OSPF] or Intermediate System-to-Intermediate System [IS-IS]) is configured.

https://t.me/learningnets

Configuring 6VPE Example

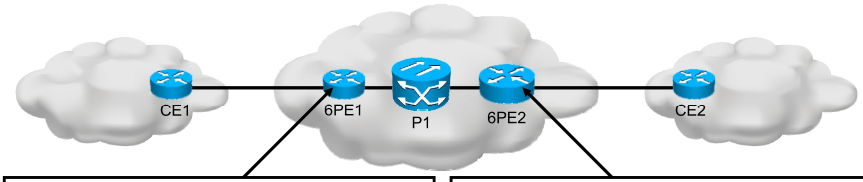


```
hostname 6PE1
!
vrf definition VPN A
 rd 100:100
 route-target export 100:100
 route-target import 100:100
!
router bgp 64512
 neighbor 10.1.1.4 remote-as 64512
 neighbor 10.1.1.4 update-source Loopback0
 !
 address-family vpnv6
 neighbor 10.1.1.4 activate
 neighbor 10.1.1.4 send-community both
!
 address-family ipv6 vrf VPN A
 redistribute connected
!
ipv6 route vrf VPN A 2003:430:210:1::/64
Serial2/0
```

```
hostname 6PE2
!
vrf definition VPN A
 rd 100:100
 route-target export 100:100
 route-target import 100:100
!
router bgp 64512
 neighbor 10.1.1.1 remote-as 64512
 neighbor 10.1.1.1 update-source Loopback0
 !
 address-family vpnv6
 neighbor 10.1.1.1 activate
 neighbor 10.1.1.1 send-community both
!
 address-family ipv6 vrf VPN A
 redistribute connected
!
ipv6 route vrf VPN A 2003:440:210:1::/64
Serial3/3
```

© 2010 Cisco Systems, Inc. All rights reserved.

The figure shows a sample configuration of PE routers. Both PE routers have an MP-BGP session for route exchange and a configured VPN (VPN A).

Enable IPv6 Routing and CEF

```
router(config)#
ipv6 unicast-routing
```

- Enables IPv6 routing

```
router(config)#
ipv6 cef
```

- Enables IPv6 Cisco Express Forwarding

© 2010 Cisco Systems, Inc. All rights reserved.

The figure displays the global configuration commands to enable the router for IPv6 routing and Cisco Express Forwarding, respectively.

https://t.me/learningnets

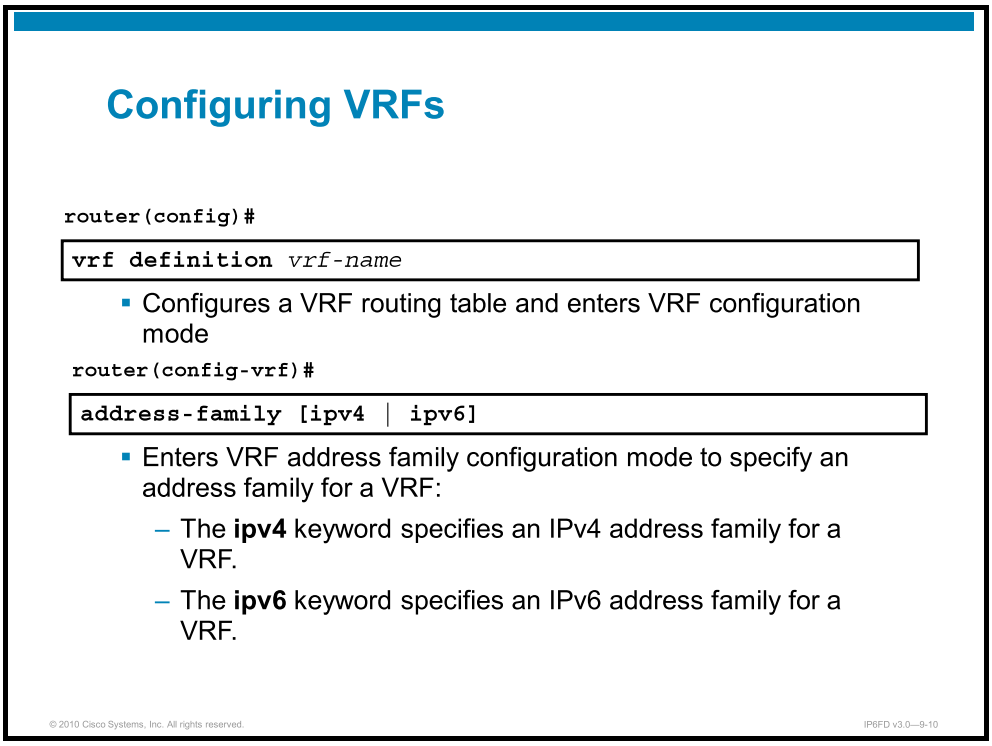

The figure shows common commands that are used to configure the IPv6 VRF VPN definition and corresponding IPv4 and IPv6 address families.

https://t.me/learningnets

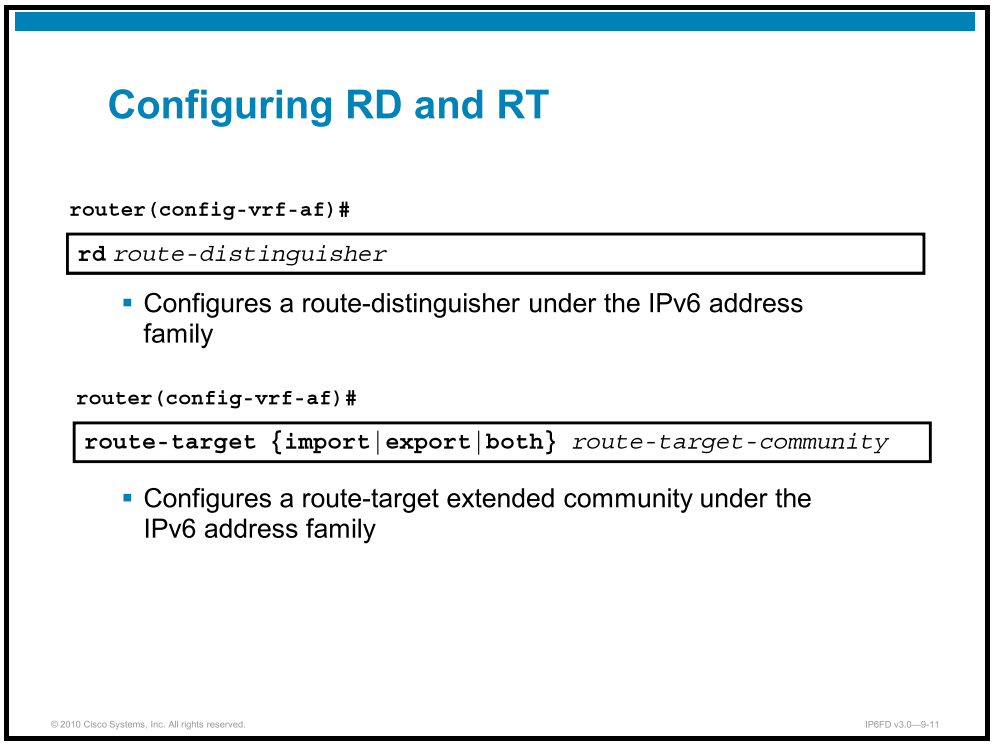

The figure shows commonly used Cisco IOS Software commands for configuration of the route target and route distinguisher; for example:

```
vrf definition VPN A
 rd 100:100
 route-target export 100:100
 route-target import 100:100
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
```

Note The route target and route distinguisher can be applied globally, under the VRF configuration, or separately, under each address family.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Configuring VRF Interfaces

```
router(config-if)#
vrf forwarding vrf-name
```

- Associates VRF with the customer interface

```
hostname 6PE1
!
interface Serial2/0
description Customer VRF interface
vrf forwarding VPN A
ip address 172.16.4.1 255.255.255.0
ipv6 address 2003:410:210:1::/64 eui-64
!
```

© 2010 Cisco Systems, Inc. All rights reserved.

The figure shows the configuration to associate a VPN routing and forwarding instance with an interface or subinterface. Use the **vrf forwarding** command in interface configuration mode. To disassociate a VRF from an interface, use the **no** form of this command.

Configuring BGP

- BGP neighbor configuration and activation for route exchange between PE routers
- Route redistribution into MP-BGP

```
hostname 6PE2
!
router bgp 64512
 no synchronization
 neighbor 10.1.1.1 remote-as 64512
 neighbor 10.1.1.1 update-source Loopback0
 no auto-summary
 !
 address-family vpnv6
 neighbor 10.1.1.1 activate
 neighbor 10.1.1.1 send-community both
 exit-address-family
 !
!
 address-family ipv6 vrf VPN A
 redistribute connected
 redistribute static
 no synchronization
 exit-address-family
!
```

© 2010 Cisco Systems, Inc. All rights reserved.

The figure displays a sample MP-BGP configuration for the 6PE2 PE router. Apart from the definition of BGP peers, the activation of the peer router 6PE1 is required. In the example, static and connected routes are also redistributed into the VPN A VRF.

https://t.me/learningnets

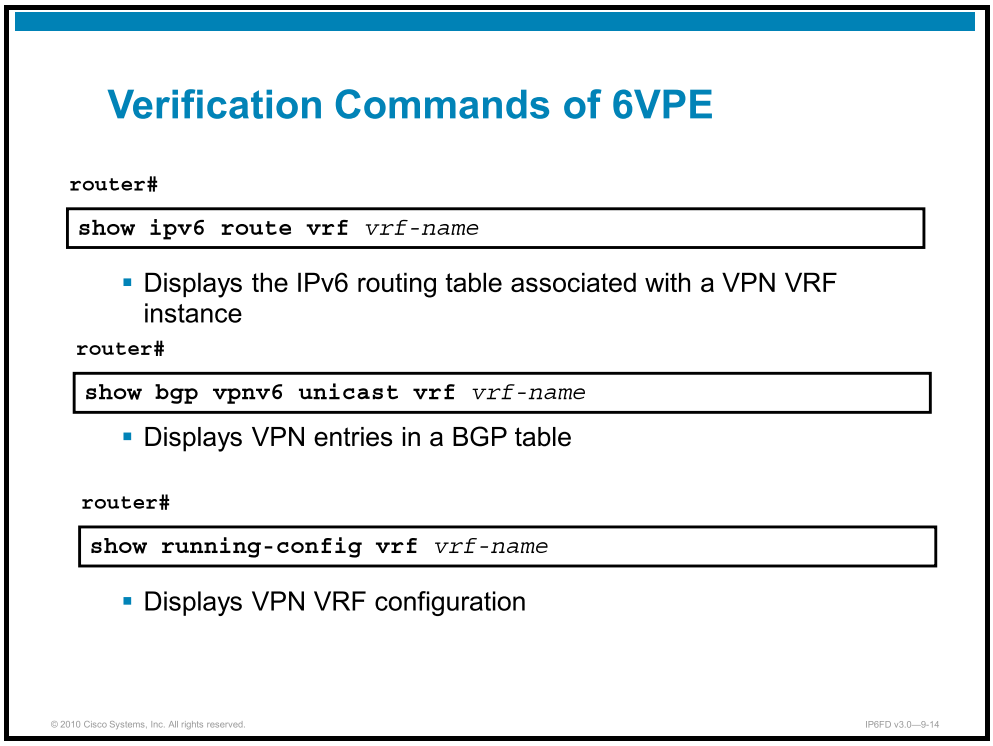

The figure shows commonly used commands for configuration verification:

```
6PE2#show ipv6 route vrf VPN A
IPv6 Routing Table - VPN A - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
       IA - ISIS interarea, IS - ISIS summary
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
C   2003:410:210:1::/64 [0/0]
     via Serial2/0, directly connected
L   2003:410:210:1:2D0:63FF:FE54:7000/128 [0/0]
     via Serial2/0, receive
B   2003:420:210:1::/64 [200/0] (line 1)
     via 10.1.1.4%Default-IP-Routing-Table, indirectly connected
S   2003:430:210:1::/64 [1/0] (line 2)
     via Serial2/0, directly connected
B   2003:440:210:1::/64 [200/0] (line 3)
     via 10.1.1.4%Default-IP-Routing-Table, indirectly connected
L   FF00::/8 [0/0]
     via Null0, receive
```

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in the lesson.

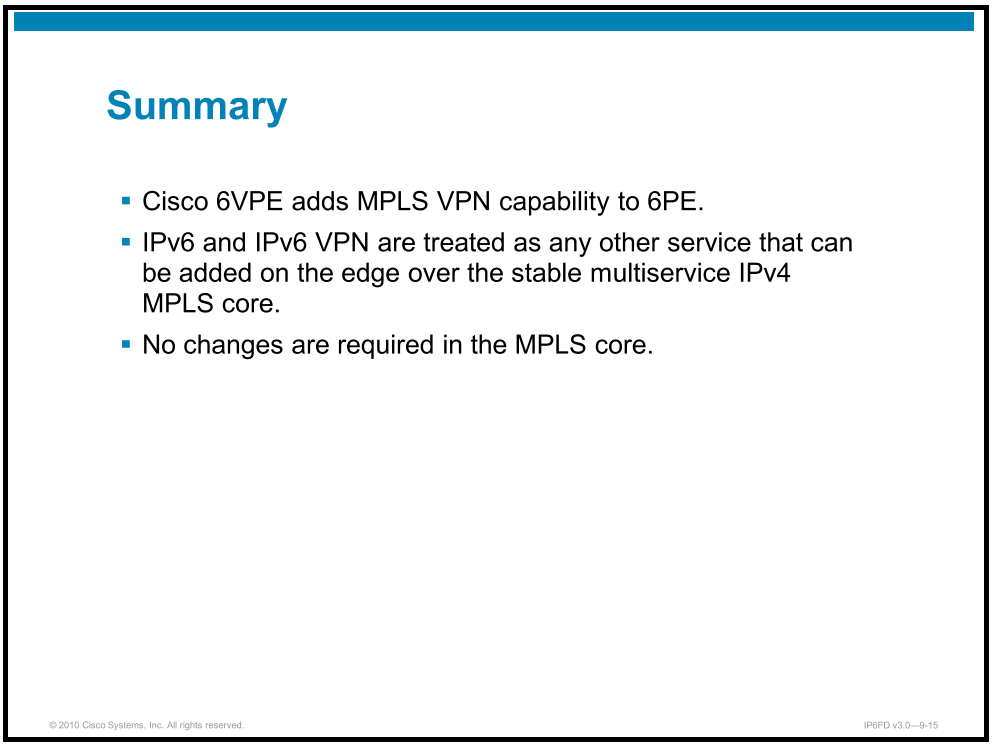

References

For additional information, refer to these resources:

- IETF. *RFC 4271, A Border Gateway Protocol 4 (BGP-4)*. January 2006. http://tools.ietf.org/html/rfc4271.
- IETF. *RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs)*. February 2006. http://tools.ietf.org/html/rfc4364.
- IETF. *RFC 4659, BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*. September 2006. http://tools.ietf.org/html/rfc4659.
- *IPv6 and 6VPE Support in MPLS VPN* at http://www.cisco.com/en/US/docs/net_mgmt/ip_solution_center/5.2/mpls_vpn/user/guide/ipv6.html#wp1002608

https://t.me/learningnets

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 4

Understanding IPv6 Broadband Access Services

Overview

IP version 6 (IPv6) is designed to help service providers provide new services to their customers. As the number of devices per broadband user exponentially increases worldwide, cable, DSL, Ethernet to the home, wireless, and other always-on access technologies can benefit from the huge address range of IPv6.

Objectives

Upon completing this lesson, you will be able to understand IPv6 broadband access services and DSL-based access in particular. This includes being able to meet these objectives:

- Describe the operation of 6rd
- List customer link encapsulations
- Explain FTTH access architecture
- Explain cable access architecture
- Explain wireless access architecture
- Explain DSL access architecture

https://t.me/learningnets

IPv6 Rapid Deployment

This topic describes mechanisms for rapid deployment of IPv6, based on RFC 5569.

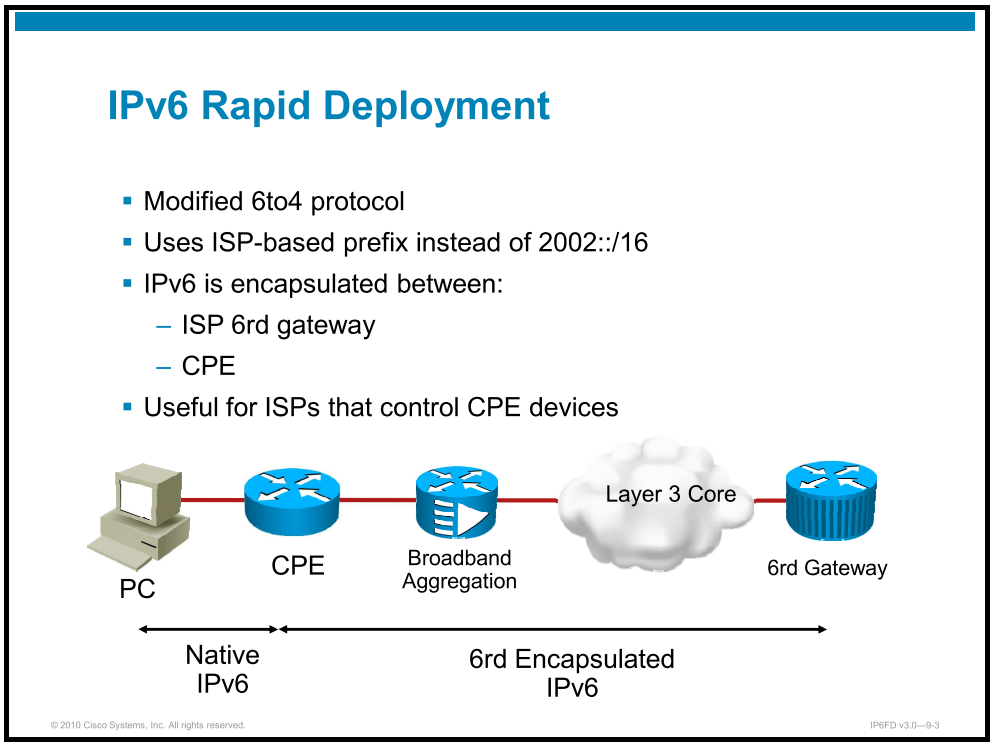

IPv6 might not be deployed at an ISP because of high initial costs and lack of customer demand. When demand finally comes, the ISP might want to rapidly deploy IPv6, with as little cost as possible. IPv6 rapid deployment (6rd) is specified in RFC 5569 and gives an ISP the chance to deploy IPv6 support without changing its core network. This type of deployment works like 6to4 encapsulation. No Multiprotocol Label Switching (MPLS) is necessary.

This type of deployment uses an ISP-based prefix, so clients are bound to ISPs, and IPv6 traffic can be accurately routed directly to the end ISP.

Encapsulation is used only inside the ISP network, to traverse IP version 4 (IPv4)-only parts of the ISP network.

Because both tunnel endpoints reside inside the ISP network, the ISP must control the user endpoint of the tunnel. Therefore, the ISP must control customer premises equipment (CPE) if 6rd is to be completely transparent to the user.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

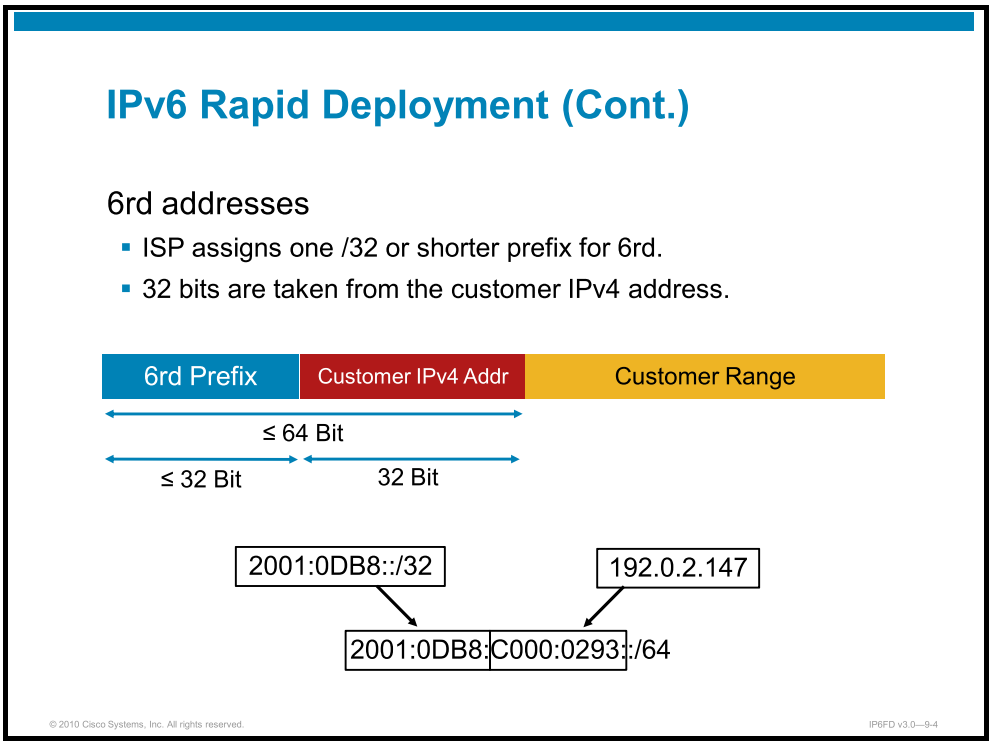

The 6rd addresses are composite addresses that consist of the following parts:

- First, as many as 32 bits form the ISP-specific 6rd prefix. If the ISP has enough addresses, then 6rd requires fewer than 32 bits. Otherwise, 32 bits are taken.
- The next field is the encoded end user IPv4 address. For 6to4 compatibility reasons, the entire 32-bit address is used, although fewer bits (for example, the last 24 bits) can be used if the routing code is modified to allow it and to use a few (for example, 8) preconfigured first bits.

In a worst-case scenario, this approach leaves the customer with the remaining 64 bits, allowing the customer to use just one IPv6 subnet per IPv4 address.

https://t.me/learningnets

IPv6 Rapid Deployment (Cont.)

Outbound packet flow

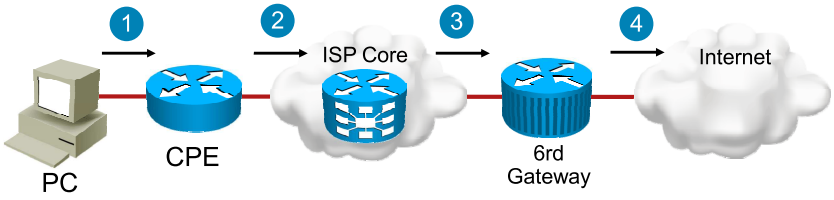

1. Native IPv6 traffic is sent to the default gateway on the client LAN.
2. IPv6 traffic is encapsulated into IPv4 and sent to the 6rd gateway.
3. IPv4 traffic is routed by the ISP core to the 6rd gateway.
4. IPv6 is de-encapsulated and sent to the Internet.

© 2010 Cisco Systems, Inc. All rights reserved.

Outbound packet flow in 6rd follows these steps:

1. Native IPv6 traffic is sent from client nodes to the CPE. The CPE, in this case, acts as a default gateway for all IPv6 traffic.
2. IPv6 traffic is encapsulated in IPv4 and sent to a 6rd gateway. There can be multiple gateways, for redundancy reasons. The 6rd gateway address is the ISP IPv4 address and must be configured on all CPE nodes.
3. IPv4 traffic is routed over the ISP core, which does not support IPv6, to the 6rd gateway.
4. The 6rd gateway de-encapsulates IPv6 from IPv4 packets and routes them to the IPv6 Internet.

https://t.me/learningnets

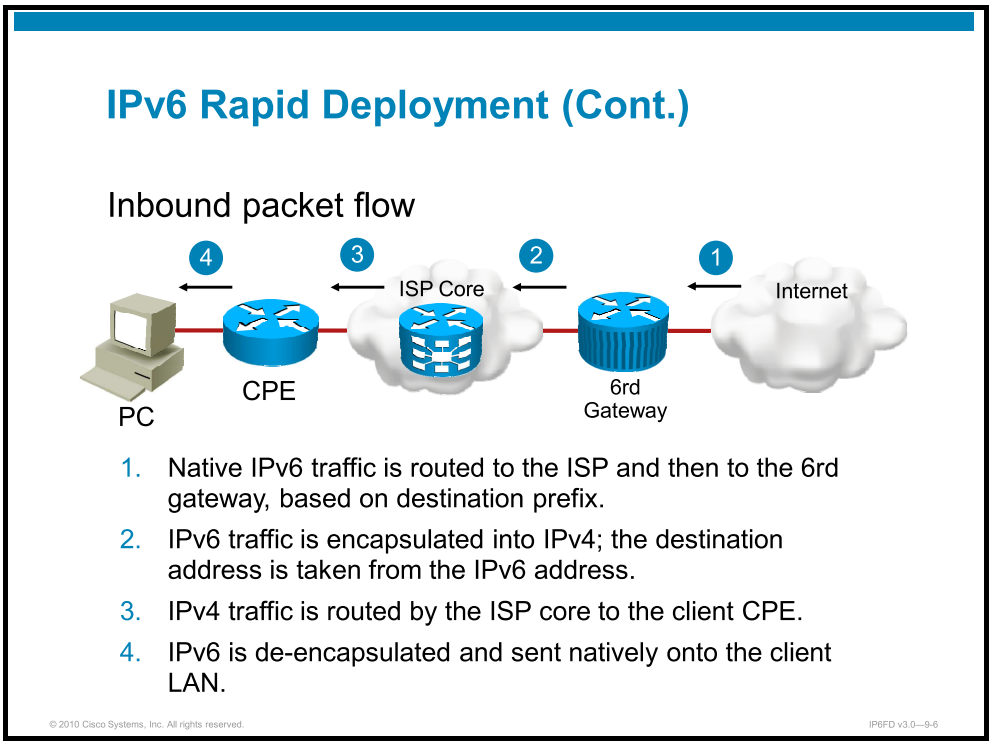

Inbound packet flow in 6rd follows these steps:

1. Native IPv6 traffic that is destined to the 6rd prefix is routed towards the ISP and then to the 6rd gateway.
2. The 6rd gateway takes the IPv6 packet, looks at the destination address, and encapsulates it into IPv4, based on the destination IPv6 address.
3. IPv4 traffic is routed across the IPv4-only core, to the CPE.
4. The CPE de-encapsulates IPv6 traffic and sends it to one of the local segments.

https://t.me/learningnets

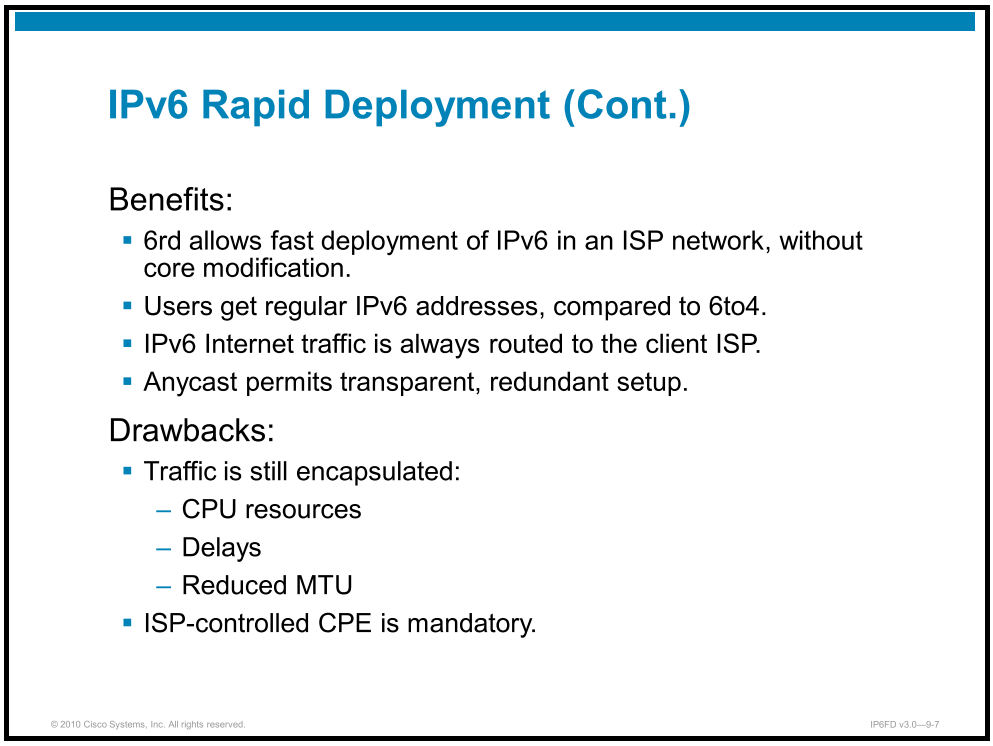

The 6rd approach has some benefits and some drawbacks:

- Benefits:
  - Client-to-Internet IPv6 can be deployed with minimal changes in the network configuration. (The 6rd gateways must be configured and CPEs must have IPv6 and 6rd support.)
  - Users get regular IPv6 addresses rather than special ranges such as 2002::/16.
  - IPv6 Internet traffic is always routed to the correct ISP, and public 6rd gateways are not necessary.
  - No single point of failure exists; redundant 6rd gateways are possible.
- Drawbacks:
  - Traffic is still encapsulated across the ISP network, incurring the following:
    - Higher demand for CPU resources
    - Higher delays for user traffic
    - Reduced maximum transmission unit (MTU) for all IPv6 traffic
  - For complete transparency for the user, CPE must be controlled by the ISP.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Customer Link Encapsulations

This topic describes IPv6 link encapsulation options.

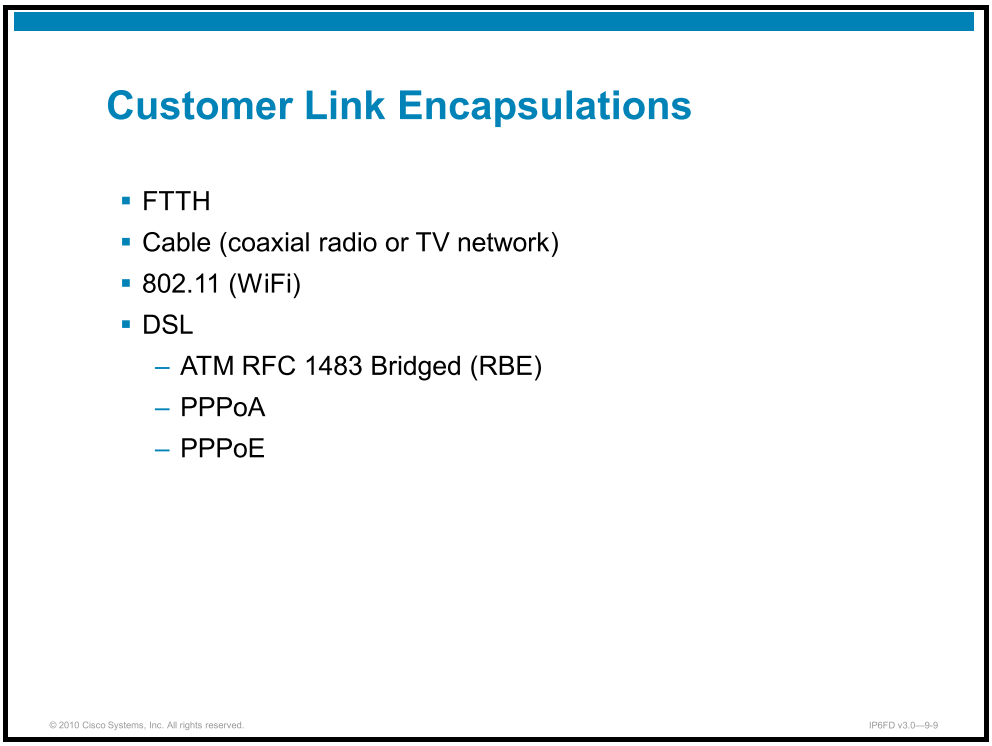

Options for broadband access services include the following:

- Ethernet and similar (802.11)
- DSL-based access

PPP over Ethernet (PPPoE), PPP over ATM (PPPoA), and routed bridge encapsulation (RBE) access methods are good encapsulation candidates to offer IPv6 connectivity in various access-service ISP designs. Those methods can make use of the IPv6 extensions to the authentication, authorization, and accounting (AAA) function. In the future, ISPs offerings will also consider Ethernet and Ethernet-like access-based services for IPv6.

The following topics describe DSL-based access services in more detail.

https://t.me/learningnets

FTTH Access Architecture

This topic describes fiber to the home (FTTH)-based access architecture and services in the IPv6 ISP offering.

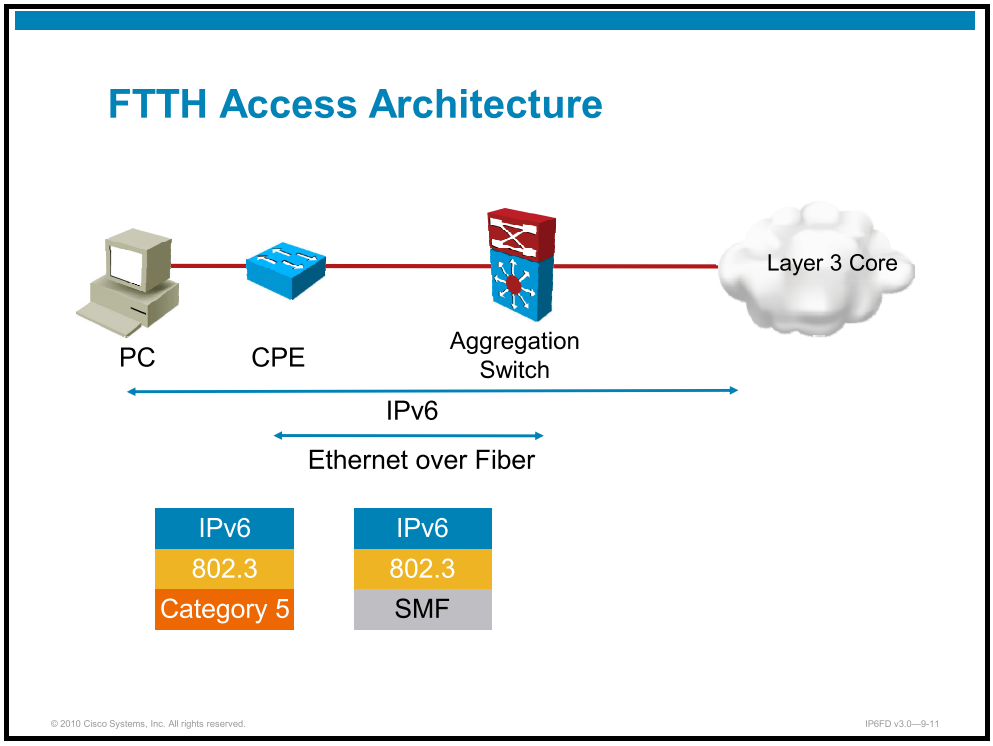

The figure displays typical components in FTTH access architecture. End hosts are directly connected to the ISP network via Ethernet switches. There is no encapsulation if ISP supports native IPv6.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Cable Access Architecture

This topic describes cable-based access architecture and services in the IPv6 ISP offering.

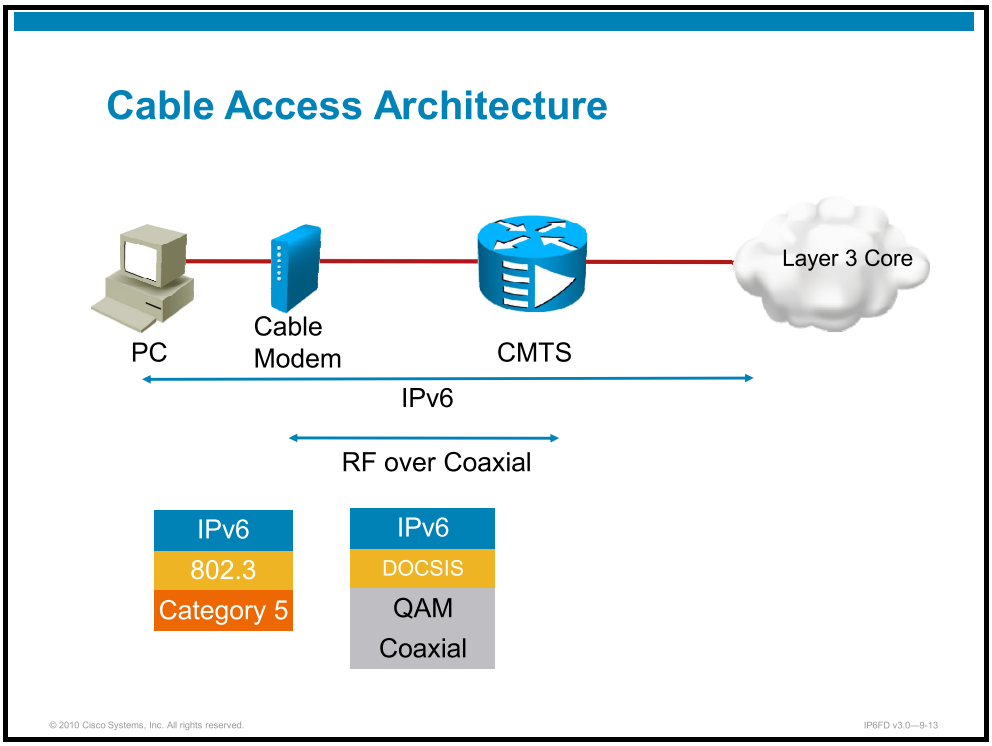

The figure displays typical components in cable network access architecture. End hosts are connected to the ISP network via a coaxial cable network. The CPE is a modem that is connected to a cable modem termination system (CMTS) over the cable network. Ethernet is bridged across, and IPv6 can be encapsulated in Ethernet like it is in IPv4.

https://t.me/learningnets

Wireless Access Architecture

This topic describes wireless-based access architecture and services in the IPv6 ISP offering.

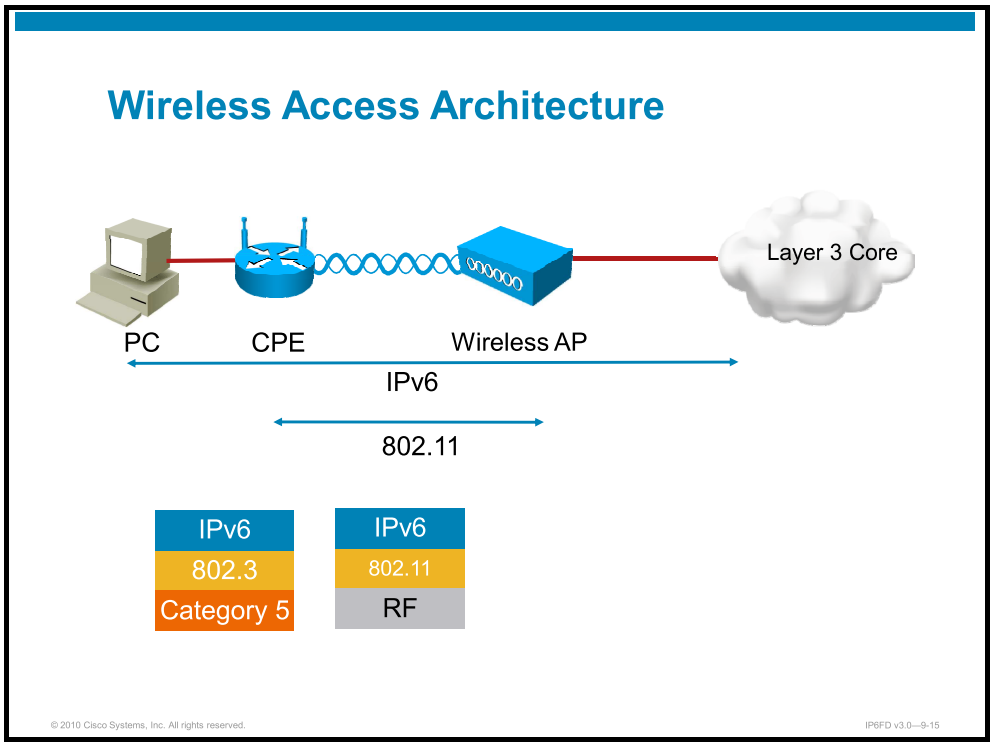

The figure displays typical components in a wireless access architecture. End hosts are connected to the ISP network via wireless equipment. There are no technical restrictions for using native IPv6 over wireless, but most Cisco equipment in small office, home office (SOHO) range (Cisco 800 Series Routers, Linksys routers) does not support IPv6 over wireless and so is inappropriate for the CPE role in an IPv6 deployment over wireless.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

DSL Access Architecture

This topic describes DSL-based access architecture and services in the IPv6 ISP offering.

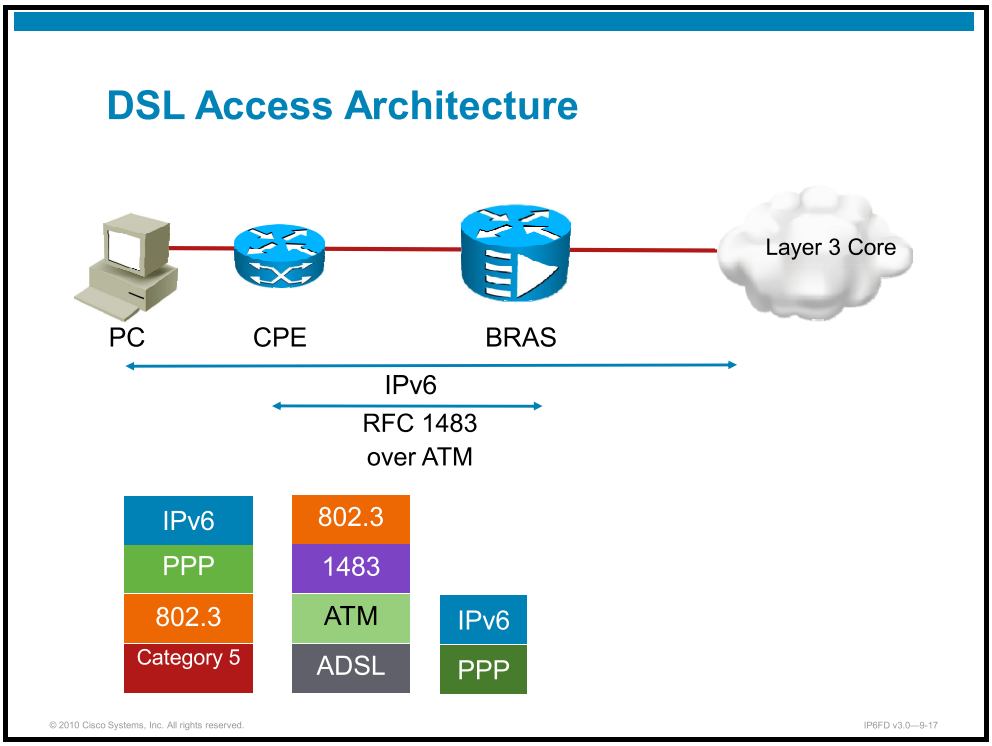

The figure displays typical components in DSL access architecture. End hosts are directly connected to the ISP network via CPE devices, and PPP sessions are terminated on the IPv6-enabled DSL access multiplexer (DSLAM).

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in the lesson.

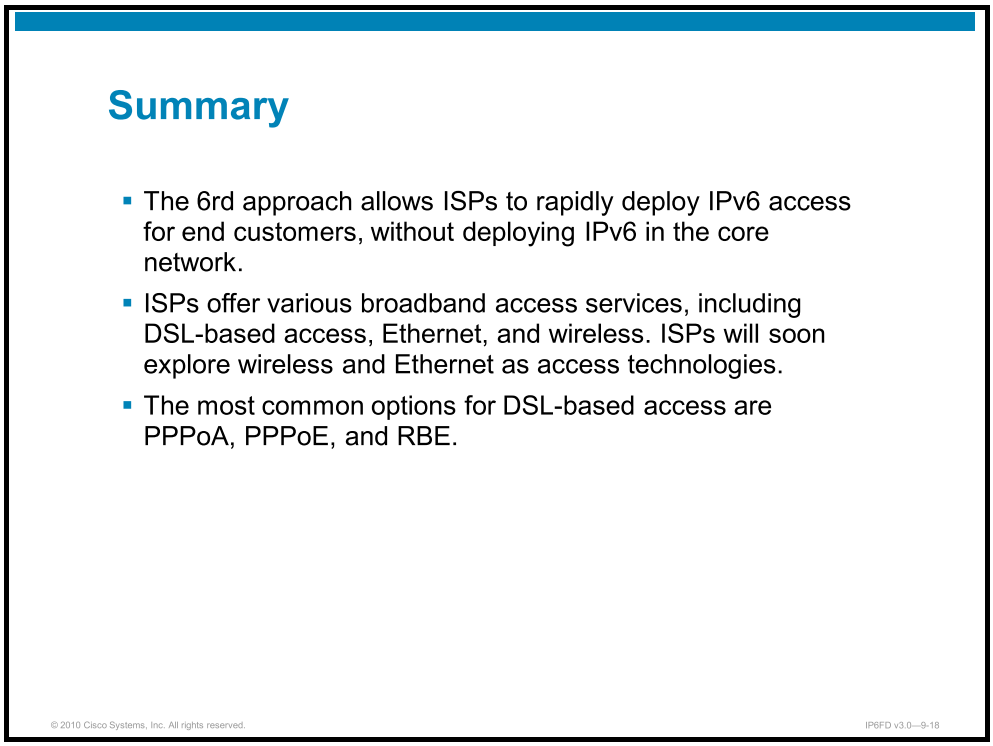

References

For additional information, refer to these resources:

- IETF. *RFC 2684, Multiprotocol Encapsulation over ATM Adaptation Layer 5.* September 1999. http://tools.ietf.org/html/rfc2684.
- IETF. *RFC 2516, A Method for Transmitting PPP over Ethernet (PPPoE).* February 1999. http://tools.ietf.org/html/rfc2516.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

This topic summarizes the key points that were discussed in this module.

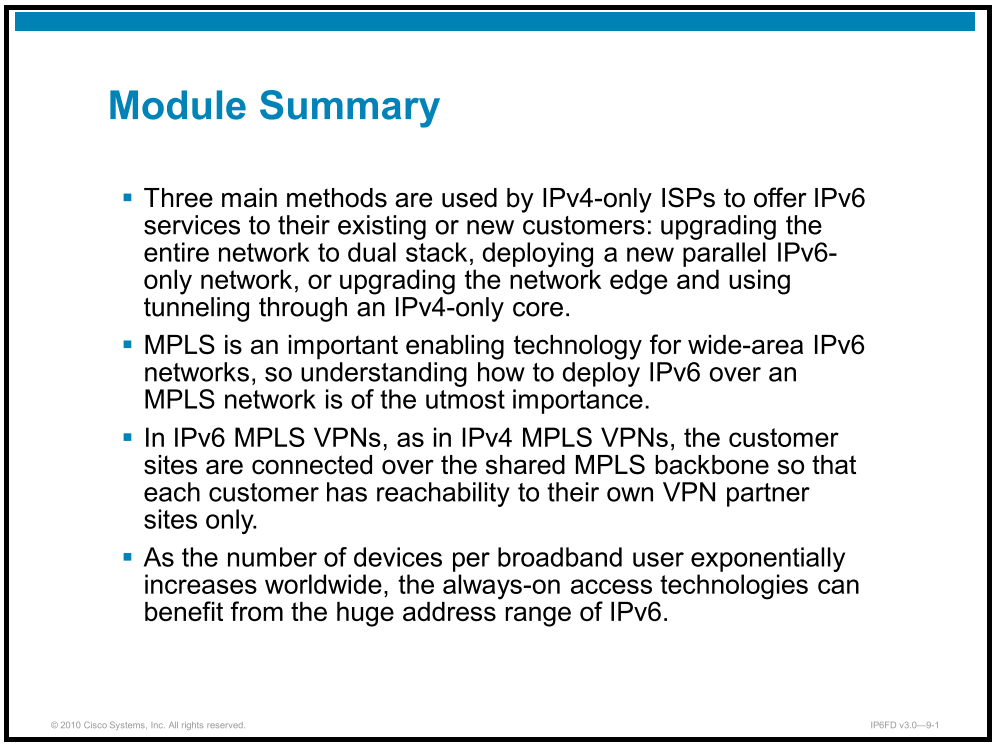

https://t.me/learningnets

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

Q1) What is one way in which IPv4 and IPv6 differ substantially in designs for the enterprise edge connection to the public network (Internet)? (Source: "Identifying IPv6 Service Provider Deployment")

A) IPv4 traffic probably is run natively through the edge device, whereas IPv6 traffic almost certainly is encapsulated in an IPv4/GRE tunnel.
B) IPv6 outperforms IPv4, and most traffic will quickly transition to IPv6 as a result.
C) IPv4 traffic is NAT-translated at the network edge, breaking the peer-to-peer model, whereas IPv6 traffic is probably routed, preserving the peer-to-peer model.
D) IPv4 and IPv6 traffic are inspected by an edge device that is capable of deep inspection of IPsec-encrypted traffic.

Q2) Tunnel techniques have several advantages over NAT-PT translation as a transition mechanism. What is a primary advantage to tunneling? (Source: "Identifying IPv6 Service Provider Deployment")

A) Tunneling preserves the end-to-end, peer-to-peer model of the network, whereas NAT-PT breaks the end-to-end model.
B) Tunneling makes it possible to fully utilize the Layer 2 MTU size, as opposed to NAT-PT, which suffers from the attachment of additional IPv6 extension headers.
C) NAT-PT supports IPsec via UDP encapsulation, whereas a tunneled IPv6 packet does not arrive intact at the destination.
D) IPv6-in-IPv4 tunneling delivers IPv6 packets intact at the ultimate destination, thus allowing IPsec to work, whereas NAT-PT breaks IPsec.

Q3) IPv6 transport schemes for IPv6 for MPLS-enabled ISPs vary in their implementations, but all have a single theme in common. What is the single theme in common? (Source: "Identifying IPv6 Service Provider Deployment")

A) Core P routers must be dual-stacked to carry IPv6 traffic.
B) Edge PE routers must all have the same topology if IPv6 support is enabled—they must all be upgraded to be Cisco 6PE routers.
C) All MPLS packets must use two-deep label stack encapsulation.
D) Core P routers do not need to be updated in most cases—upgrades are limited to the edge PE routers within the provider network.

Q4) IPv6 can be carried over MPLS networks by using simple protocol 41 tunneling in many cases. In the most common case, where are the tunnel endpoints? (Source: "Understanding Support for IPv6 in MPLS")

A) within the customer network environment, not on any ISP device
B) edge-to-edge within the ISP network, from the ingress PE to the egress PE
C) from the ingress ISP-provided router to the far-end CPE
D) from the ISATAP host to the ISP-provided egress Cisco 6PE router

https://t.me/learningnets

Q5) What is a Cisco 6PE device? (Source: “Understanding Support for IPv6 in MPLS”)

A) an ISP router that is IPv6-capable and translates IPv6 packets into MPLS packets

B) an ISP router that is dual-stacked and runs IPv6 facing the customer and MPLS facing the core of the MPLS network

C) a customer-provided router that encapsulates IPv6 packets into IPv4 MPLS packets

D) an ISP router that encapsulates IPv6 packets in IPv4 GRE tunnels for transport across the MPLS network

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answer Key

Q1) C

Q2) A

Q3) D

Q4) A

Q5) B

https://t.me/learningnets

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Module 10

IPv6 Case Studies

Overview

The growth of the Internet and the adoption of networking over the past 20 years are pushing IP version 4 (IPv4) to the limits of its addressing capacity and its ability for continued growth. To sustain the evolution of the Internet, the Internet Engineering Task Force (IETF) developed a next-generation protocol, IP version 6 (IPv6). This module describes the factors that led to IPv6 development and compares IPv4 with IPv6.

Module Objectives

Upon completing this module, you will be able to describe case studies for enterprise, service provider, and branch networks. This ability includes being able to meet these objectives:

- Describe the most common planning and implementation approaches, distinguish among enterprise network designs, and adopt a high-level implementation strategy for the process of transitioning from IPv4-only to dual-stack architectures
- Describe and identify the most common planning and implementation approaches, as they pertain to moving to IPv6 in service provider core networks
- Describe and identify the most common planning and implementation approaches, as they pertain to moving to IPv6 in branch networks

https://t.me/learningnets

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1

Planning and Implementing IPv6 in Enterprise Networks

Overview

Many enterprise IT departments are considering a shift to IP version 6 (IPv6) throughout their infrastructure. The dual-stacking approach is the most common way to integrate IPv6 functionality into IP version 4 (IPv4) environments. This shift requires much planning and careful design efforts, together with precise and timely staging. IT departments assess only limited and controllable parts of the infrastructure during the initial phases, before adopting a complete transition approach for the entire network. Being aware of these challenges and the impact that they impose on production networks is important for every IPv6 technology engineer and implementer. This lesson will give you an overview of industry best practices, specific concerns, and common solutions that pertain to planning and implementing IPv6 in enterprise networks. The solutions that are presented in this lesson follow Cisco validated designs and best practices.

Objectives

Upon completing this lesson, you will be able to describe the most common planning and implementation approaches, distinguish among enterprise network designs, and adopt a high-level implementation strategy for the process of transitioning from IPv4-only to dual-stack architectures. This ability includes being able to meet these objectives:

- Define an enterprise network
- Identify specific concerns with campus networks that use IPv6
- Identify specific concerns with WAN networks that use IPv6

https://t.me/learningnets

Enterprise Network Definition

This topic defines an enterprise network and its characteristics.

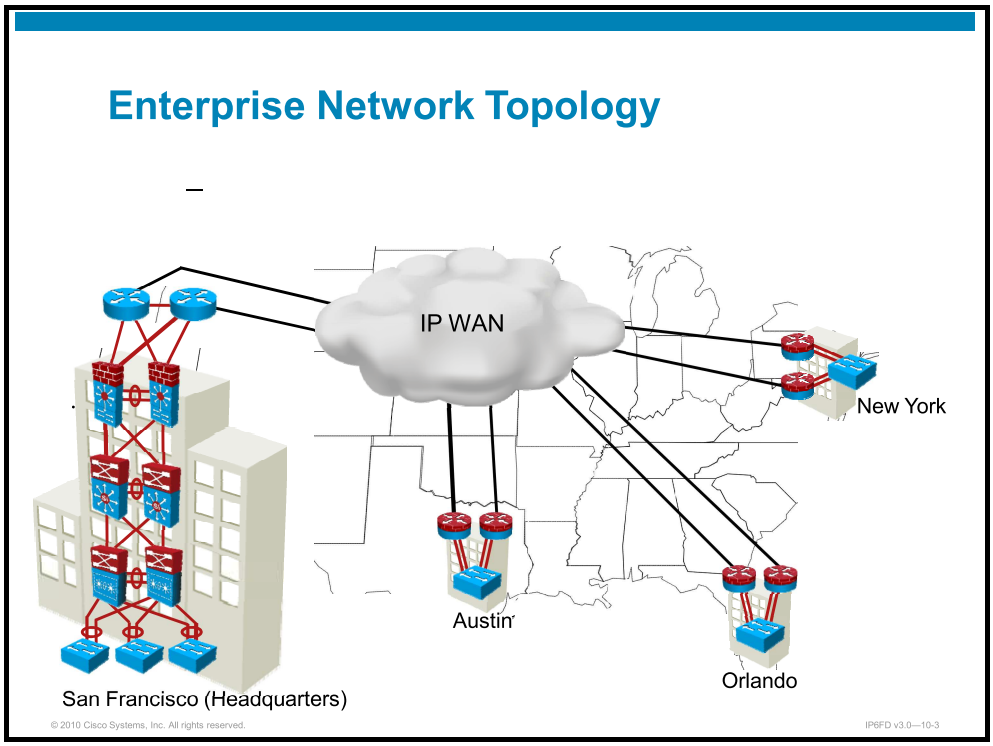

The enterprise Cisco Validated Design network has been a driver of company productivity and competitive advantage for many years. The enterprise network addresses many requirements and challenges, including global enterprise availability, collaboration and real-time communication applications, continuous evolution of answers to security threats, and adaptation flexibility. The network also successfully answers other expectations for anywhere, anytime access; increased capacity; and manageable complexity.

The sample enterprise network, Enterprise6, consists of an enterprise headquarters campus in San Francisco and three other locations in Austin, Texas; Orlando, Florida; and New York City, New York. All these locations are interconnected through high-speed, highly available links on the IP WAN.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

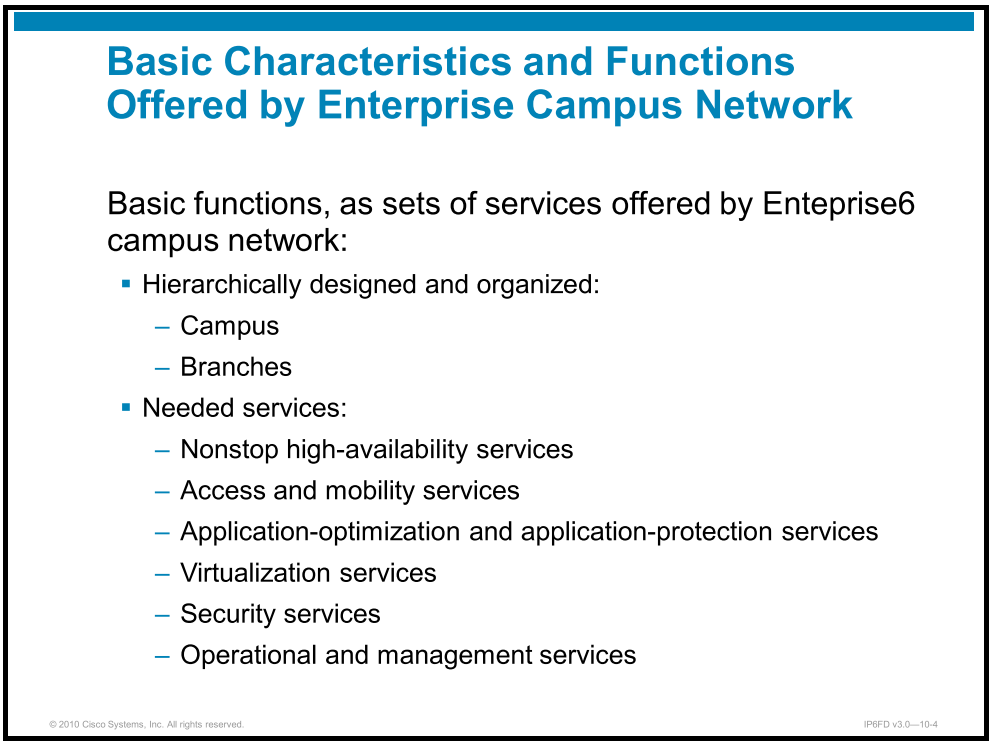

Enterprise networks are organized hierarchically. Boundaries of functional areas limit the propagation of failures and allow for easier management. An enterprise network consists of one or more campus networks that are located at the principal locations and of remote branches.

The campus network consists of integrated elements that comprise a set of services. These services are used by a group of users and end-station devices that all share the same high-speed switching communications fabric. The services include packet-transport services (both wired and wireless), traffic identification and control (security and application optimization), traffic monitoring and management, and overall systems management and provisioning. These basic functions are implemented to provide and directly support the higher-level services the enterprise IT organization provides for the end-user community.

These functions include the following:

- Nonstop high-availability services
- Access and mobility services
- Application-optimization and application-protection services
- Virtualization services
- Security services
- Operational and management services

https://t.me/learningnets

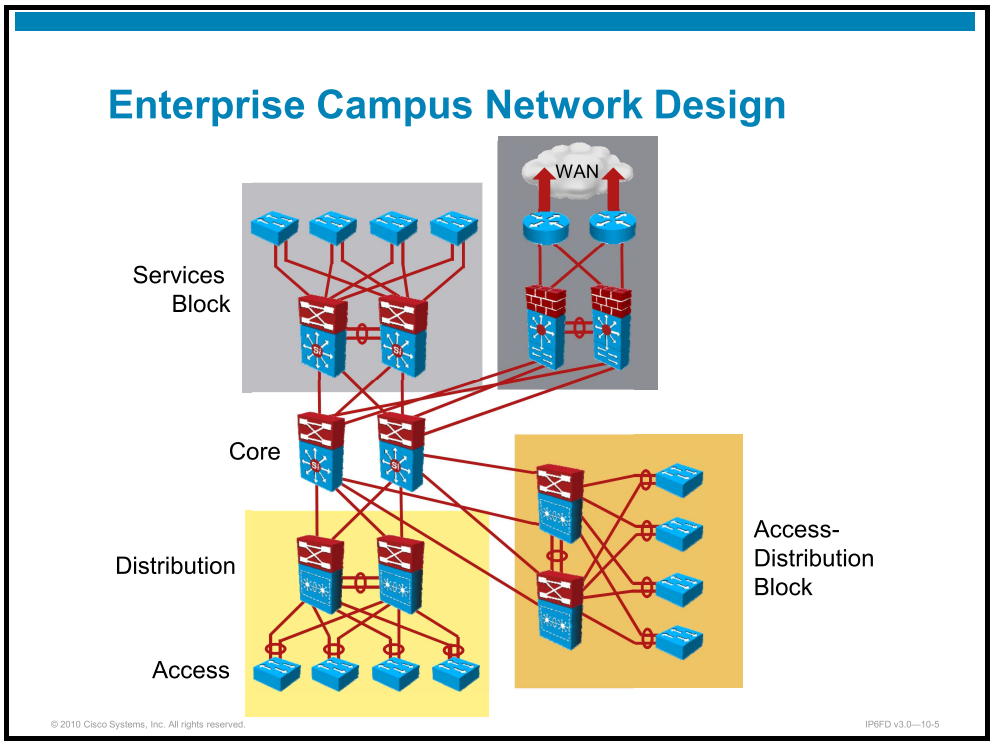

The enterprise campus is that portion of the enterprise computing infrastructure that provides access to network communication services and resources for end users and devices that are spread over one geographic location. In the example, the campus spans the headquarters building in San Francisco. The enterprise campus network also acts as the core, or backbone, of the network and provides interconnectivity between other portions of the overall network.

The campus core interconnects the campus access, the data center, and WAN portions of the network. The Enteprise6 campus network comprises a set of capabilities and services.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

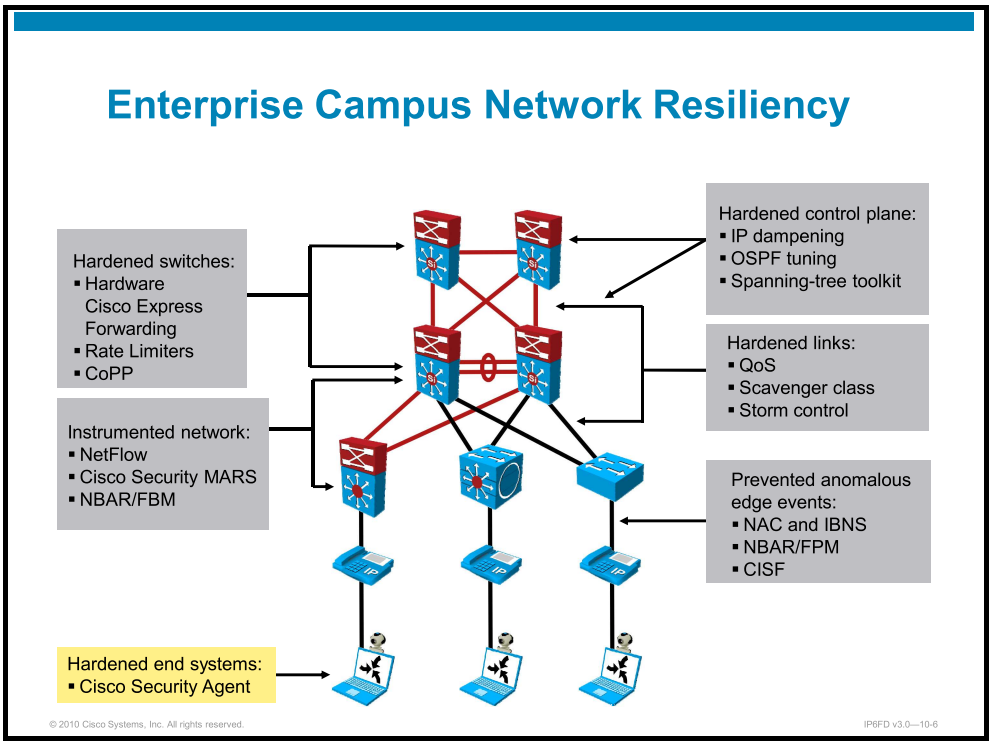

The enterprise employs a high level of resiliency on its campus network, through a few approaches. These approaches include hardened individual components, switches, and links; additional rate-limiting capabilities of software and hardware functions; explicit controls on the behavior of edge devices; and instrumentation and management tools that provide feedback to the network operation teams.

https://t.me/learningnets

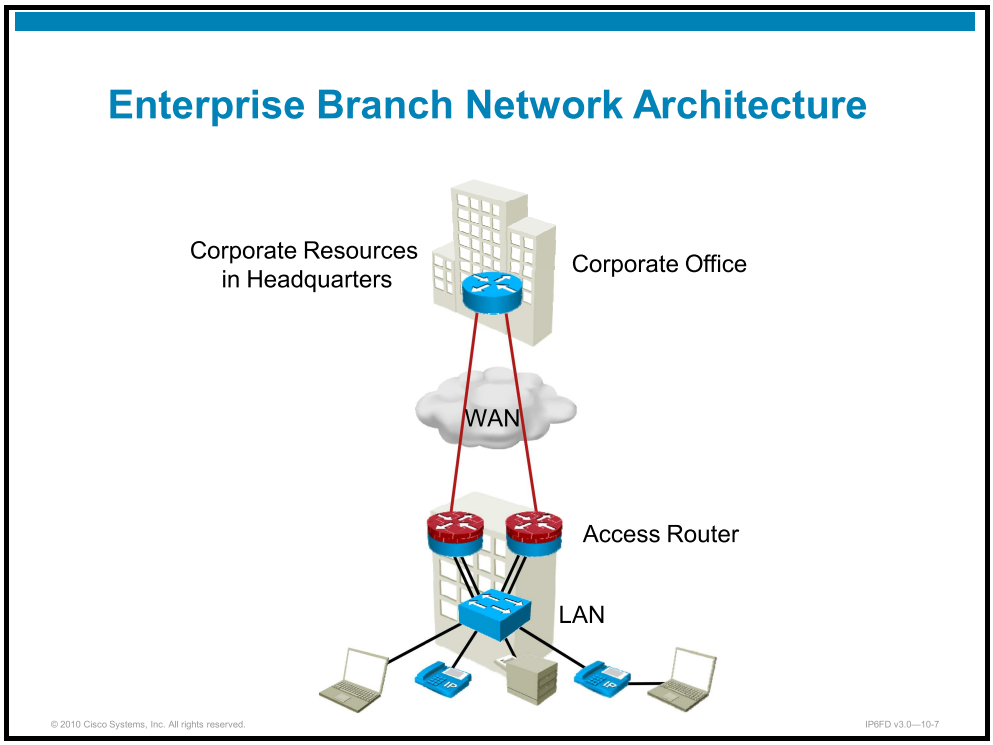

The branch offices follow the dual-tier enterprise branch-architecture framework on the network infrastructure layer.

The dual-tier profile consists of two Cisco 2900 Series or 3900 Series Integrated Services Routers that are connected to an external switch. Dual WAN links and box redundancy provide a greater level of high availability for all enterprise branch locations. For Enterprise6, all resources are located at the headquarters. WAN and LAN services are not integrated in this enterprise branch-architecture profile. The integrated services routers (ISRs) serve to terminate WAN connections, and a desktop switch performs the LAN connectivity. In this way, the enterprise offers high availability with distributed LAN connectivity via desktop switches and offers WAN connectivity via branch routers.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Implementing IPv6 in an Enterprise Campus Network

This topic describes specific concerns in campus networks that use IPv6.

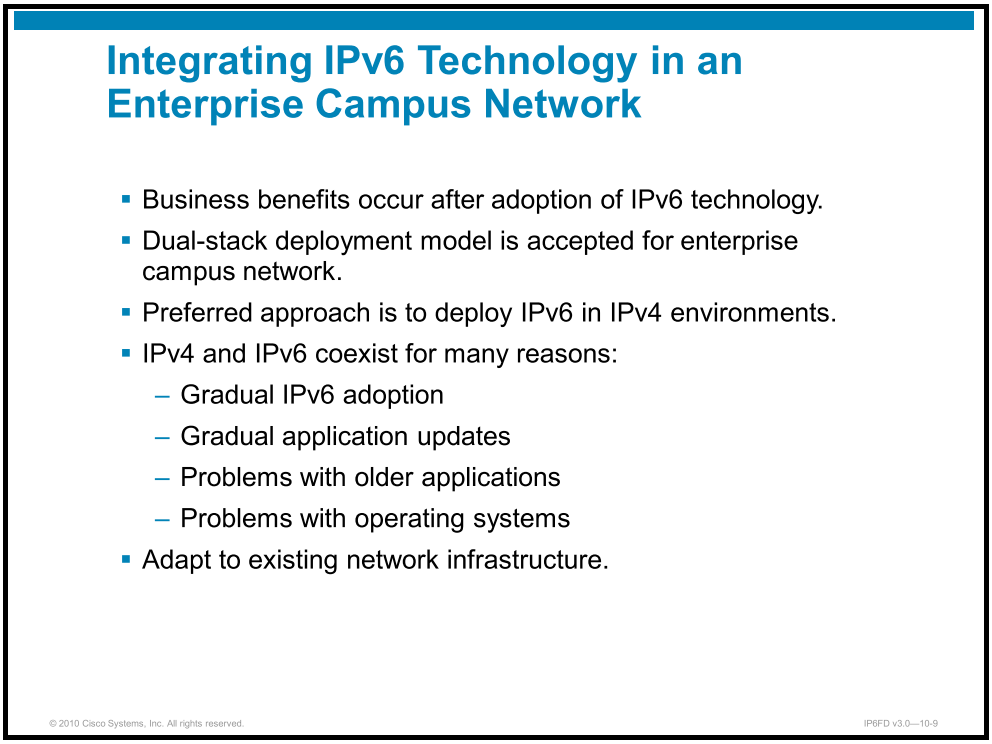

The integration of IPv6 technology and related capabilities in the campus network underscores the main principle of IPv6 integration planning: Upgrades to the infrastructure and applications can be performed in parallel.

A completely IPv6-capable infrastructure, by itself, delivers no business value. The value is achieved only when the enterprise begins to experience business benefits as it adapts its applications and operations to take advantage of the new capabilities that such an infrastructure makes possible.

IPv4 and IPv6 will need to coexist for many years in the enterprise, for several reasons:

- Most enterprises adopt an IPv6 infrastructure gradually, rather than all at once.
- If an enterprise completely turns off support for IPv4 in the network, then its IPv4 applications, websites, and services no longer work. By continuing to support IPv4, the enterprise can gradually upgrade applications to IPv6 and test them—a solution that is less expensive than upgrading and testing all applications at the same time.
- Adding IPv6 support to older applications for which the company does not own the source code might not be possible.
- Upgrading or replacing old but stable operating systems or platforms that are used for dedicated applications might not provide return on investment (ROI).

The dual-stack deployment model that is accepted for the sample Enteprise6 campus network is based on the dual-stack transition mechanism.

https://t.me/learningnets

Dual-stack is the preferred, most versatile way to deploy IPv6 in existing IPv4 environments. IPv6, along with the associated features to make IPv6 routable, highly available, and secure, can be enabled wherever IPv4 is enabled. In some cases, IPv6 is not enabled on a specific interface or device because of the presence of legacy applications or hosts for which IPv6 is not supported. Conversely, IPv6 may be enabled on interfaces and devices for which IPv4 support is no longer needed.

Deploying IPv6 in the campus by using the dual-stack deployment model offers several advantages over using other models. The primary advantage is that the dual-stack model does not require tunneling within the campus network.

The primary drawback is that network equipment upgrades might be required if the existing network devices are not IPv6 capable.

The dual-stack approach, as a hybrid model strategy, is to employ two or more independent transition mechanisms that have the same deployment goals. Flexibility is the main aspect of the hybrid approach, in which any combination of transition mechanisms can be used to best fit a given network environment. The hybrid model adapts as much as possible to the characteristics of the existing network infrastructure.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

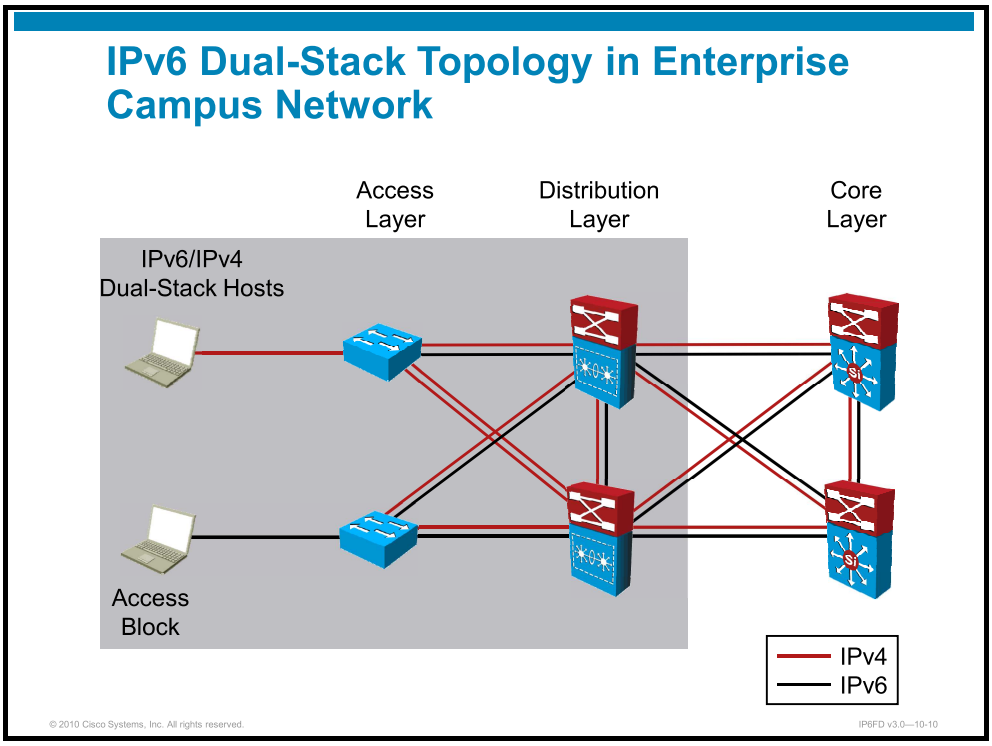

The high-level overview of IPv6 dual-stack topology in an enterprise campus network represents the three-tier topology, with access, distribution, and core layers supporting both IPv4 and IPv6.

https://t.me/learningnets

IPv6 Campus Network Address Prefix Size

64 Bits	Less than 64 Bits	More than 64 Bits
Recommended by RFC 3177 and IAB/IESG Consistency makes management easy Required for SLAAC, SEND, and privacy extensions and for Microsoft DHCPv6 server Significant address-space loss	More hosts per broadcast domain Considered bad practice 64 bits offer more space for hosts than the media can support efficiently	Address space conversation Special cases: • /126 valid for point-to-point links • /127 not valid for point-to-point links • /128 loopback Complicates management Must avoid overlap with specific addresses: • Router Anycast (RFC 3513) • Embedded RP (RFC 3956)

© 2010 Cisco Systems, Inc. All rights reserved.

Several link prefix size considerations exist:

- **/64:** On VLAN interfaces, using a /64 prefix is recommended because this prefix is easy and consistent for address management. This prefix is required for Stateless Address Autoconfiguration (SLAAC), SEND, and privacy extension use and for Microsoft DHCP version 6 (DHCPv6) services.
- **Less than 64:** There are no real use cases in which a site needs more addressing on a link than a /64 prefix can provide, so this size and is considered a bad practice.
- **Greater than 64:** Some opinions in the IPv6 community assert that a /64 prefix for point-to-point links is a waste and even a security attack vector. Many deployments use a /64 prefix on VLANs (or links on which hosts reside), /126 on point-to-point links, and /128 on loopbacks.

https://t.me/learningnets

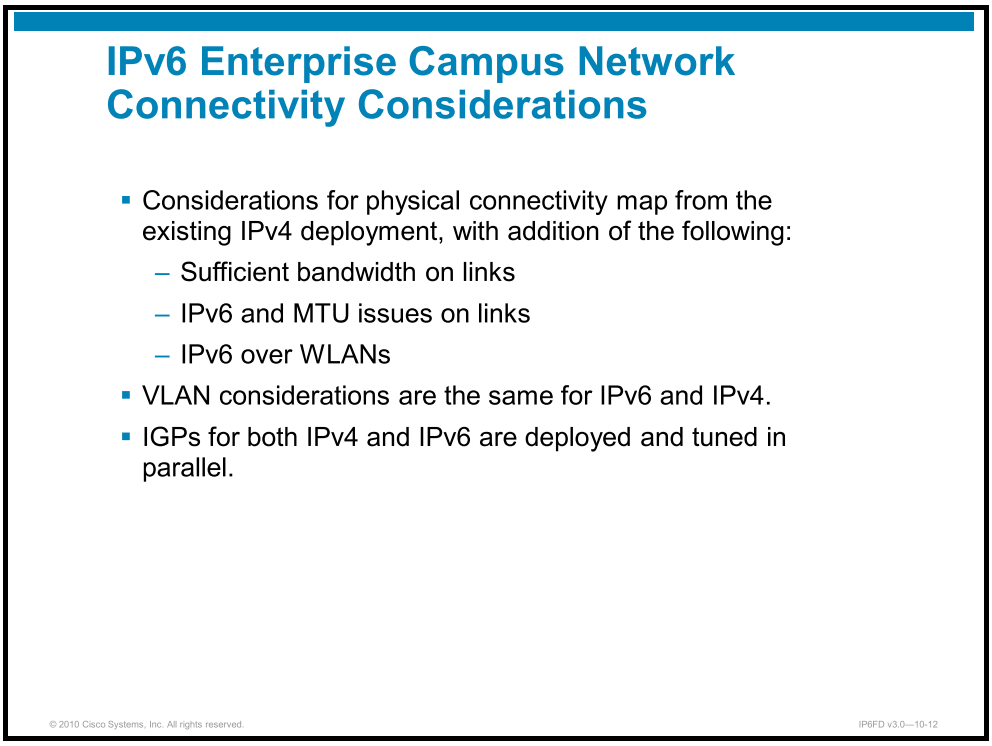

Considerations for physical connectivity with IPv6 are the same as with IPv4, with the addition of three elements:

- Ensuring that there is sufficient bandwidth for both existing and new traffic.
- Understanding how IPv6 manages the maximum transmission unit (MTU) on a link
- Operating IPv6 over wireless LANs (WLANs)

Sufficient bandwidth for both IPv4 and IPv6 traffic is an important factor for the deployment of any new technology, protocol, or application. IPv6 issues with the MTU and fragmentation in IPv6 on the link are more closely observed in RFC 2460 and RFC 1981. IPv6 should also operate correctly over WLAN access points (APs), much as it operates over Layer 2 switches. Specific concerns in WLAN environments include managing WLAN devices via IPv6 and controlling IPv6 traffic via APs or controller-based quality of service (QoS), VLANs, and access control lists (ACLs). IPv6 must be supported on the AP and controller devices if it is to take advantage of these more-intelligent services on the WLAN devices.

VLAN considerations for IPv6 are the same as for IPv4. When using dual-stack configurations like those in the Enteprise6 campus network, both IPv4 and IPv6 traverse the same VLAN. The use of IPv6 on data VLANs that are trunked along with voice VLANs behind the IP phones is also considered and supported in the enterprise campus network.

The interior gateway protocol (IGP) for IPv4 in the enterprise campus network is Enhanced Interior Gateway Routing Protocol (EIGRP), and the configuration for IPv6 can be either EIGRP or Open Shortest Path First (OSPF) version 3 (OSPFv3). Both IPv4 and IPv6 need to be tuned, and a top priority should be that the IGPs provide a stable, scalable, and fast converging network.

https://t.me/learningnets

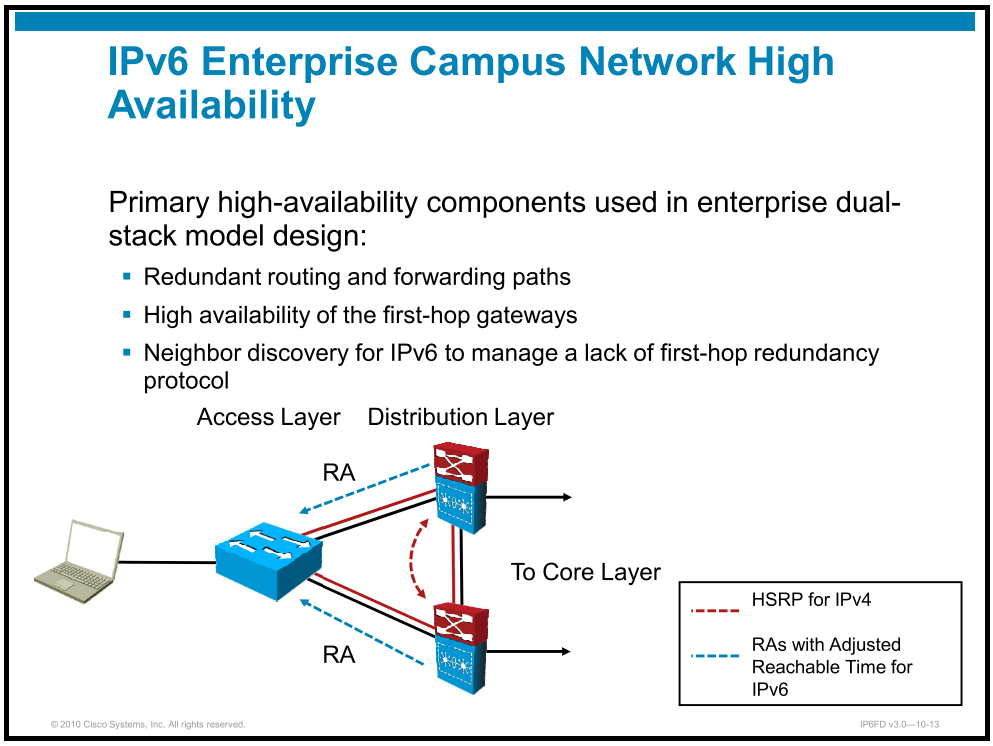

These primary components are used in the enterprise dual-stack model design:

- Redundant routing and forwarding paths
- High availability of the first-hop gateways
- Neighbor discovery for IPv6

The redundant routing and forwarding paths are accomplished by using EIGRP for IPv6 or by using OSPFv3 with the functionality of Cisco Express Forwarding.

In the dual-stack model design, the distribution layer switches are the first Layer 3 devices connected to the hosts in the access layer. Traditional campus designs use first-hop redundancy protocols such as Hot Standby Routing Protocol (HSRP), Gateway Load Balancing Protocol (GLBP), or Virtual Router Redundancy Protocol (VRRP) to provide first-hop redundancy.

Crude first-hop availability, for environments in which HSRP or GLBP for IPv6 are not yet used, is provided by the neighbor discovery method in IPv6. Neighbor discovery for IPv6 (as defined in RFC 2461) implements the use of neighbor unreachability detection (NUD). NUD is a mechanism that allows a host to determine whether a router (neighbor) in the host default gateway list is unreachable. Hosts receive the NUD value, which is known as the reachable time, from the routers on the local link via regularly advertised route advertisements. The default reachable time is 30 seconds.

NUD is used when a host determines that the primary gateway for IPv6 unicast traffic is unreachable. A timer is activated; when the timer expires (reachable time value), the neighbor begins to send IPv6 traffic to the next available router in the default gateway list. Under default configurations, using the next gateway in the default gateway list will take no longer than 30 seconds.

The NUD value should be adjusted only on link and VLANs on which hosts reside. Switches that support a real first-hop redundancy protocol such as HSRP or GLBP for IPv6 do not need to have the reachable time adjusted.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

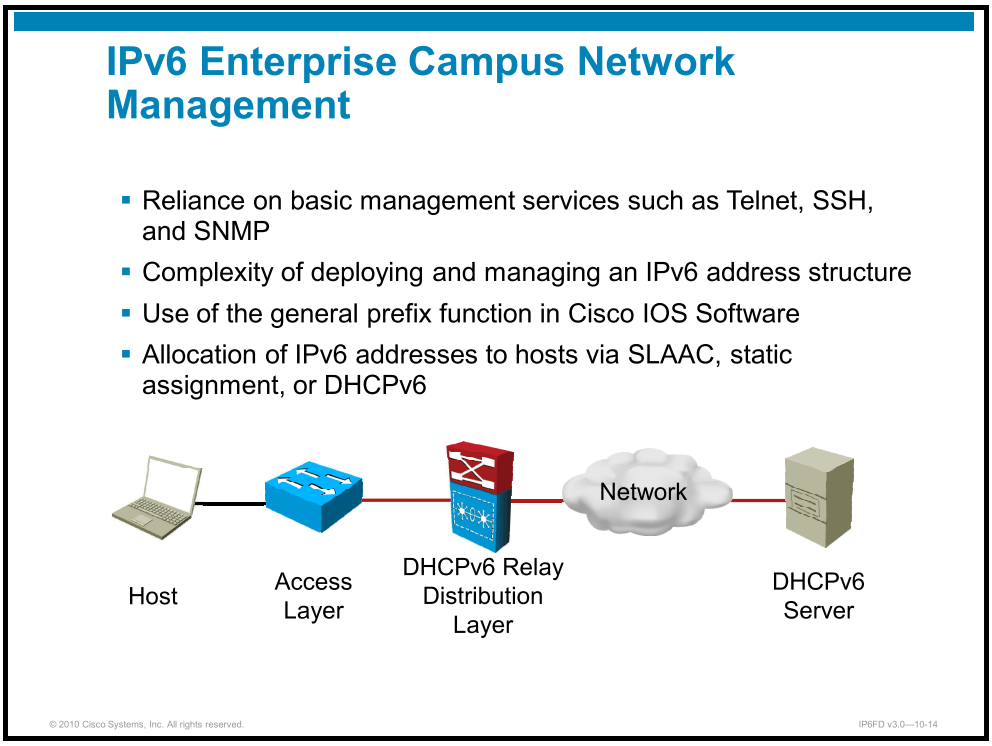

Management tools and instrumentation for IPv6 are still under development. In the enterprise campus network, you will rely on basic management services such as Telnet, Secure Shell (SSH), and Simple Network Management Protocol (SNMP). SNMP over IPv6 transport is supported on the most recent versions of Cisco IOS Software, depending on the Cisco Catalyst platform.

Analyzing IPv6, even at an elementary level, is enough to understand the size and the potential complexity of deploying and managing the IPv6 address structure. Deploying large hexadecimal addresses on many network devices should be automated—or at least more user friendly—than it is today. Industry efforts are in place to provide recommendations and solutions to the address-management issue.

One way to simplify the deployment of address prefixes on a campus switch is to use the general prefix feature. This feature lets you define a user-friendly prefix name (or names) in the global configuration of the switch. That user-friendly name can be used, on a per-interface basis, to replace the usual IPv6 prefix definition on the interface.

You can assign IPv6 address allocation to hosts in the campus access layer via SLAA, static assignment, or DHCPv6. DHCPv6 assignment is the predominant method for enterprise campus administrators. With DHCPv6 Relay Agent support, you can allocate and manage addresses much like you do with DHCP in IPv4.

But having DHCPv6 Relay Agent support in the network is only part of the equation. The client must support DHCPv6, and there must be a DHCPv6 server. Cisco has tested four DHCPv6 servers for campus networks:

- Cisco Network Registrar
- Cisco DHCPv6 server in Cisco IOS Software
- Windows 2008 Server
- Dibbler

Enterprise6 campus network devices will be served by a Windows 2008 DHCPv6 server.

https://t.me/learningnets

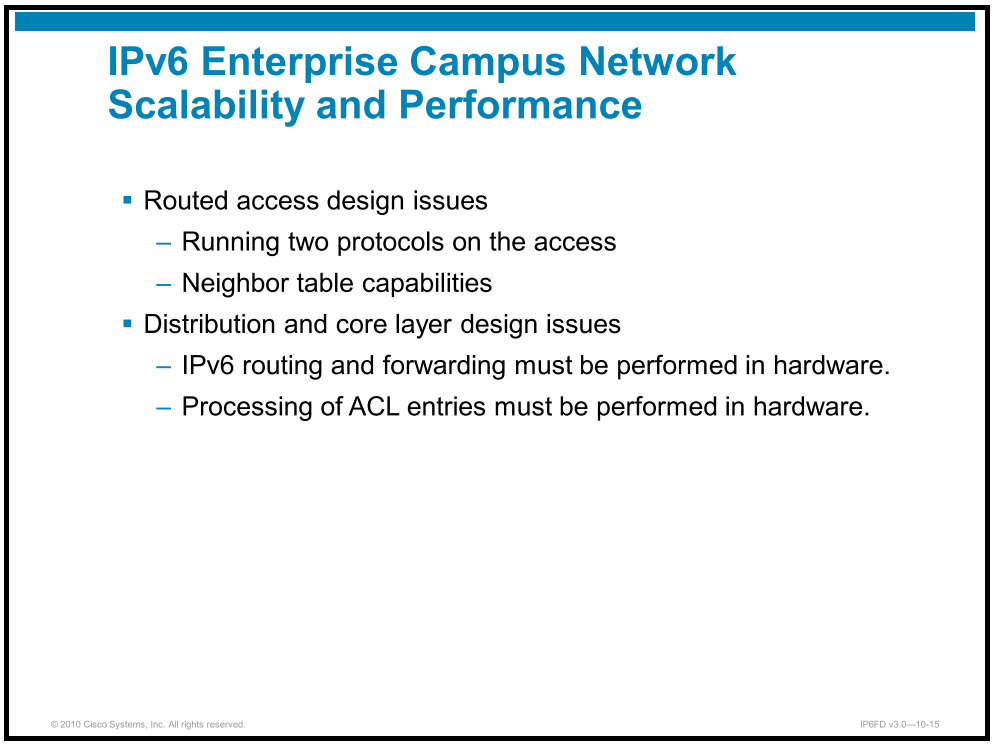

Understanding the link, memory, and CPU utilization of the existing Enteprise6 campus network before deploying the IPv6 dual-stack model is very important. If any of these aspects are already stressed, then adding IPv6 (or any new technology, feature, or protocol) into the design is not recommended.

IPv6 implementations often experience a change in traffic utilization ratios on the campus network links. As IPv6 is deployed users leverage it as the transport for applications that were historically IPv4-only, IPv4 traffic utilization is often reduced. An increase in overall network utilization usually derives from controlling traffic for routing.

One of the primary scalability considerations of routed-access design is running two protocols on the routed access or distribution layer switch. In the routed access or distribution layer, the switch must track both IPv4 and IPv6 neighbor information. Like Address Resolution Protocol (ARP) in IPv4, a neighbor cache exists for IPv6. With IPv4, there is usually a one-to-one mapping of an IPv4 address to a MAC address; with IPv6, multiple IPv6 addresses that the host might have (such as link-local, unique-local, and multiple global addresses) can be mapped to a single MAC address in the neighbor cache of the switches.

For scalability and performance issues of the distribution and core layer, two other considerations are valid for the dual-stack model:

- IPv6 routing and forwarding must be performed in hardware.
- The processing of ACL entries must be performed in hardware. IPv6 ACLs in the distribution layer are primarily used for QoS and security, to protect the control plane of the switch from attack.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

IPv6 in an Enterprise WAN Network

This topic describes specific concerns with branch networks that use IPv6.

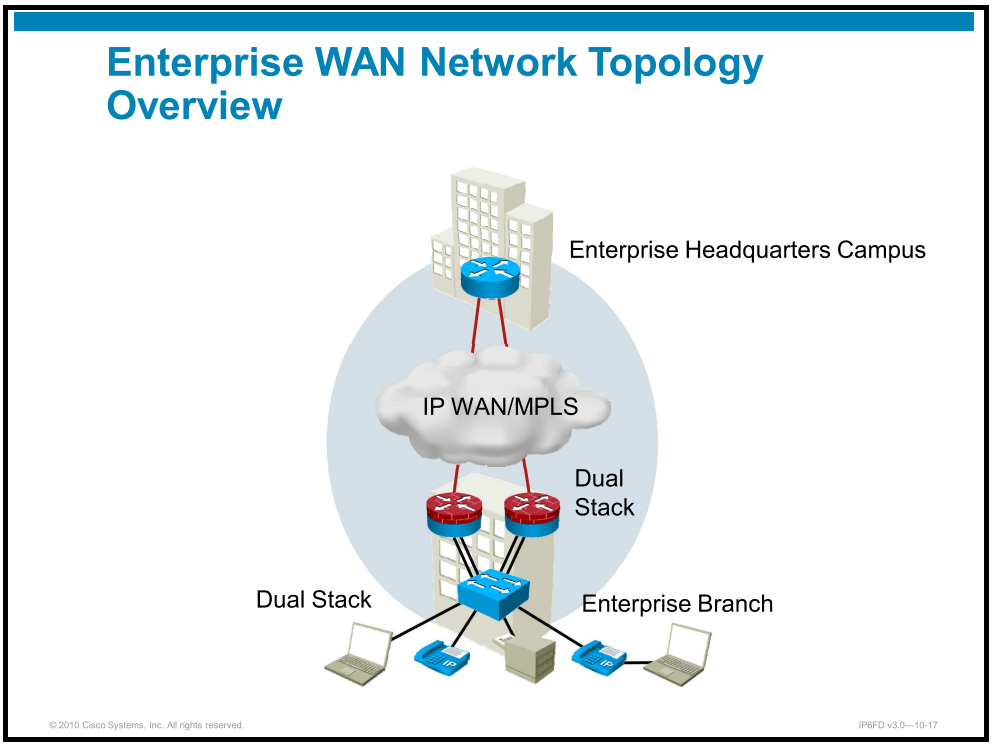

Enterprise WAN network deployment for branch locations follows typical dual-stack and dual-tier deployment methodologies and guidelines.

IPv4 and IPv6 dual-stack devices are employed at the headquarters and branch locations.

All Enteprise6 branch locations use the dual-tier approach to separate the WAN and LAN networks between dedicated devices.

https://t.me/learningnets

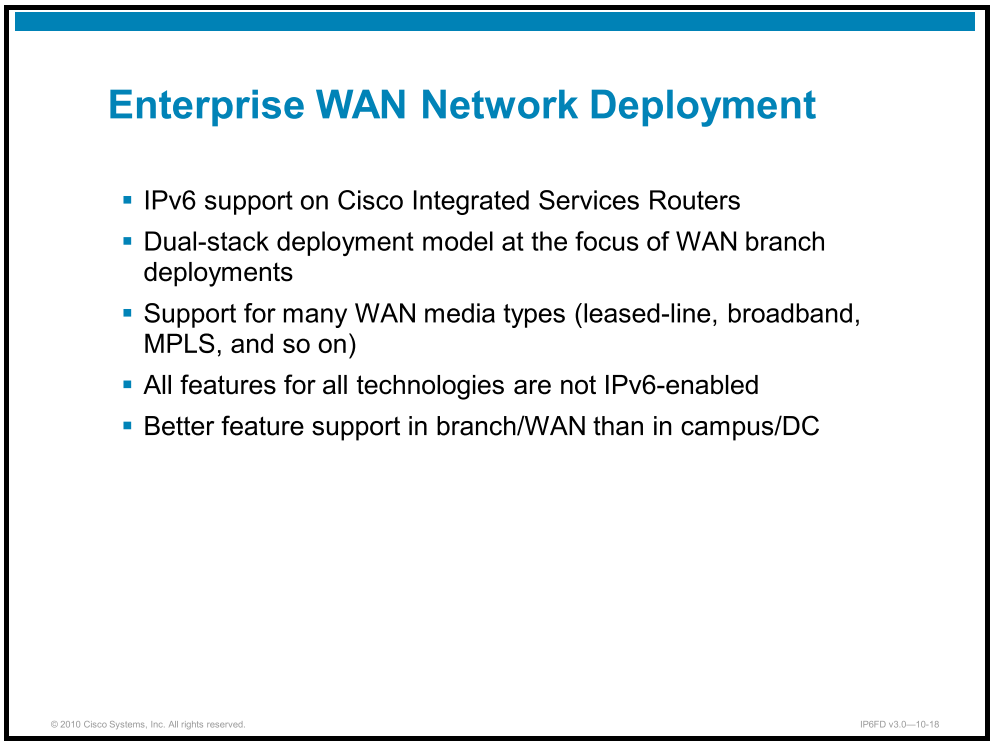

Cisco offers IPv6 support on its complete portfolio of routing platforms, including Cisco Integrated Services Routers. Most modern branch deployments follow the dual-tier deployment model because of its benefits and advantages and its direct, simple deployment approach. Many administrators and network architects choose the dual-tier, dual-stack approach for branch WAN network deployments because of its support for many WAN media types, such as Frame Relay, leased lines, broadband, Multiprotocol Label Switching (MPLS), and so on. However, not all IPv4-compatible technology features are IPv6-enabled, and the actual feature support for branch and WAN deployments is much better than for campus and data center deployments.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

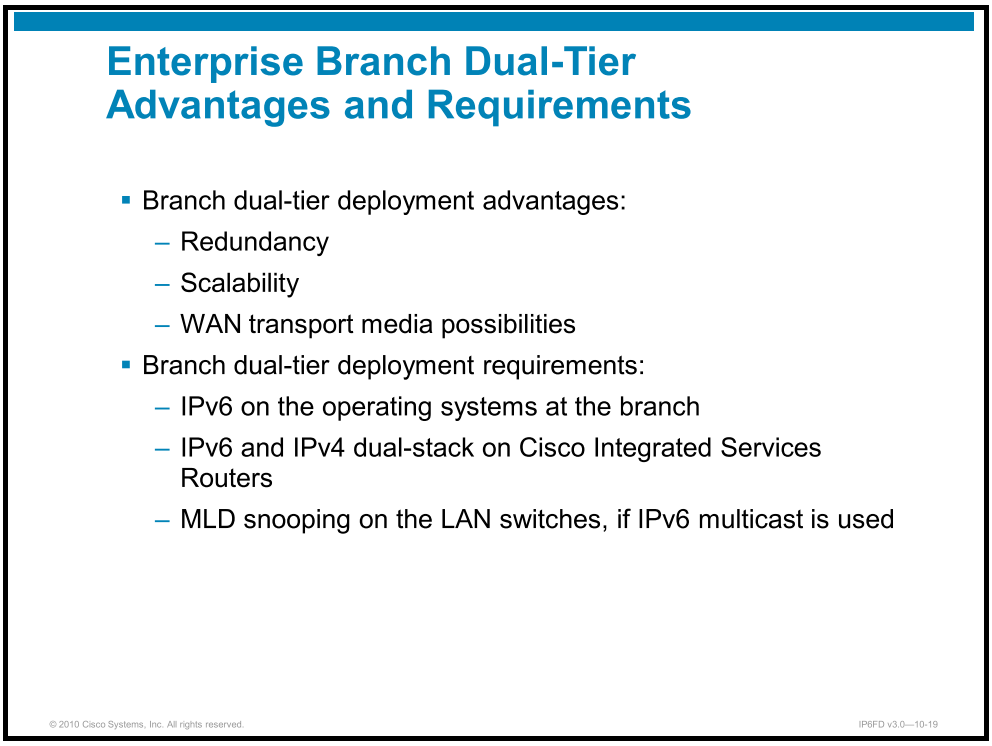

The dual-tier approach separates the LAN and WAN components, to offer fault-tolerance. Two WAN routers connect to the IP WAN and MPLS or Frame Relay cloud, and they are also redundantly connected to the LAN switch.

Scalability in the dual-tier model approach is achieved by the separation of the LAN and WAN components in the branch network. If additional LAN switches are required at the branch, then the Catalyst switches can be used together by using the Cisco StackWise topology.

With the dual-stack deployment model, both IPv4 and IPv6 are used all the way from the hosts in the branch LAN, across the WAN, and into the headquarters network. This approach eases the operational aspects of deploying IPv6 in the branch because no special tunnel considerations are needed.

Both dual-tier routers provide security services in the dual-stack deployment model, and no IPsec tunnels are used.

These requirements are needed when deploying the dual-stack branch WAN network:

- IPv6 support on the operating systems on all machines in the branch
- IPv6 and IPv4 dual-stack support on the Cisco Integrated Services Routers
- MLD-snooping support on the LAN switches, if using IPv6 multicast

https://t.me/learningnets

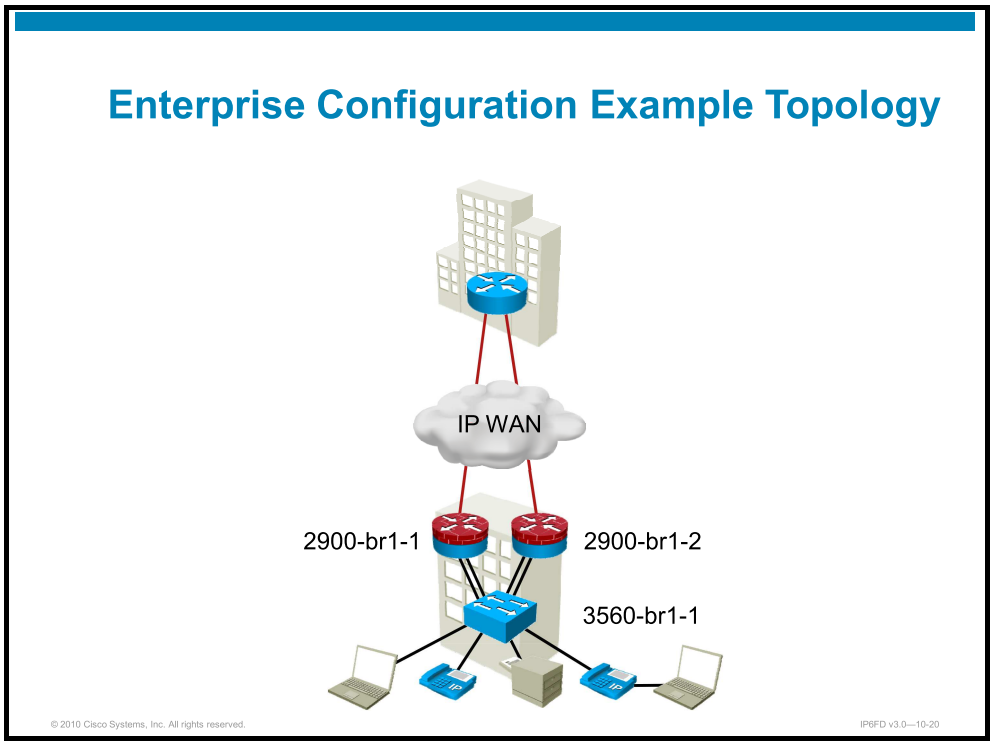

The topology in the figure serves as the basis for these configuration examples.

Dual-Stack WAN Configuration

Enterprise dual-stack deployment consists of two Cisco 2900 Series Integrated Services Routers, which provide WAN connectivity, and one Cisco Catalyst 3560 Series Switch, which provides connectivity for LAN devices.

The WAN consists of one serial connection from each of the two branch routers. No tunnels are used in the dual-tier profile. The serial interfaces on the branch routers are enabled for dual-stack operation and do not have any dependency on IPv4 to pass IPv6 packets across the WAN. This scenario is optimal, compared to running tunnels.

The IPv4 WAN configuration is not shown in the listing.

The IPv6 WAN configuration for an ISR is as follows:

```
interface Serial0/1/0
ipv6 address 2001:DB8:CAFE:1262::BAD1:1010/64
```

Dual-Stack LAN Router Configuration

The part of the configuration that is shown is the LAN IPv6 configuration for a Cisco 2900 Series Integrated Services Router.

The configuration shows the trunk links between the router and a Catalyst 3560 Series Switch. The DHCP configuration for IPv6 is also present. The IPv6 DHCP pool is used for the data-only VLAN 100.

HSRP version 2 (HSRPv2) for IPv6 is used between the two routers. Router 2900-br1-1 is configured to be the active HSRP router. This router tracks the serial interfaces for HSRP. If the serial link goes down, the router triggers HSRP to switch to standby mode. GLBP for IPv6 is also supported and can be used instead of HSRP.

The figure does not show the IPv4 LAN configuration.

https://t.me/learningnets

The following is the IPv6 LAN configuration for ISR 2900-br1-1:

```
ipv6 unicast-routing
ipv6 cef
!
ipv6 dhcp pool DATA OS
dns-server 2001:DB8:CAFE:10:20D:9DFF:FE93:B25D
dns-server 2001:DB8:CAFE:10:51A1:5B1:4A85:B3DA
domain-name cisco.com
!
interface FastEthernet0/0.100
 description DATA VLAN for PCs
 encapsulation dot1Q 100
 ipv6 address 2001:DB8:CAFE:2100::BAD1:1010/64
 ipv6 nd other-config-flag
 ipv6 dhcp server DATA OS
 standby version 2
 standby 201 ipv6 autoconfig
 standby 201 priority 120
 standby 201 preempt delay minimum 30
 standby 201 authentication ese
 standby 201 track Serial0/1/0.10 90
!
interface FastEthernet0/0.200
 description Voice VLAN for IP Phones
 encapsulation dot1Q 200
 ipv6 address 2001:DB8:CAFE:2200::BAD1:1010/64
 standby version 2
 standby 202 ipv6 autoconfig
 standby 202 priority 120
 standby 202 preempt delay minimum 30
 standby 202 authentication ese
 standby 202 track Serial0/1/0.10 90
!
interface FastEthernet0/0.300
description PRINTER VLAN
encapsulation dot1Q 300
ipv6 address 2001:DB8:CAFE:2300::BAD1:1010/64
standby version 2
standby 203 ipv6 autoconfig
standby 203 priority 120
standby 203 preempt delay minimum 30
standby 203 authentication ese
standby 203 track Serial0/1/0.10 90
```


https://t.me/learningnets

Dual-Stack LAN Switch Configuration

The following part of the LAN IPv6 configuration is for Catalyst 3560 Series Switch 3560-br1-1. The configuration shows the VLANs and trunk configurations, as well as the management IPv6 address.

```
3560-br1-1
!
vtp domain ese-branch
vtp mode transparent
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree portfast bpduguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
vlan internal allocation policy ascending
!
vlan 100
 name DATA
!
vlan 200
 name VOICE
!
vlan 300
 name PRINTERS

interface FastEthernet0/1
 description To 2900-br1-1 TRUNK
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 100,200,300
 switchport mode trunk
 load-interval 30
!
interface FastEthernet0/2
 description To 2900-br1-2 TRUNK
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 100,200,300
 switchport mode trunk
 load-interval 30
!
interface FastEthernet0/4
 description phone with PC connected to phone
 switchport access vlan 100
 switchport mode access
 switchport voice vlan 200
 load-interval 30
 spanning-tree portfast
 spanning-tree bpduguard enable
```


© 2010 Cisco Systems, Inc.

https://t.me/learningnets

```
!
interface Vlan100
 description VLAN100 for PCs and Switch management
 ipv6 address 2001:DB8:CAFE:2100::BAD2:F126/64
!
```

Routing Configuration

EIGRP for IPv6 is used on both Frame Relay links to the headquarters site and on all LAN interfaces. This part of the IPv6 routing configuration is for ISR 2900-br1-1.

```
ipv6 unicast-routing
!
key chain ESE
 key 1
  key-string 7 04490A0808245E
!
interface Loopback0
 ip address 10.124.102.1 255.255.255.255
 ipv6 address 2001:DB8:CAFE:2000::BAD1:1010/128
 ipv6 eigrp 1
!
interface FastEthernet0/0.100
 description DATA VLAN for PCs
 ipv6 eigrp 1
!
interface FastEthernet0/0.200
 description Voice VLAN for IP Phones
 ipv6 eigrp 1
!
interface FastEthernet0/0.300
 description PRINTER VLAN
 ipv6 eigrp 1
!
interface Serial0/1/0
ipv6 eigrp 1
 ipv6 hold-time eigrp 1 35
 ipv6 authentication mode eigrp 1 md5
 ipv6 authentication key-chain eigrp 1 ESE
!
ipv6 router eigrp 1
 router-id 10.124.102.1
 stub connected summary
 no shutdown
 passive-interface FastEthernet0/0.100
 passive-interface FastEthernet0/0.200
 passive-interface FastEthernet0/0.300
 passive-interface Loopback0
!
```


https://t.me/learningnets

The following is the static route on switch 3560-br1-1, out VLAN 100, to the HSRP virtual IP (VIP) address that is used on 2900-br1-1 and 2900-br1-2:

```
ipv6 route ::/0 Vlan100 FE80::5:73FF:FEA0:C9
```

Security Configuration

For the IPv6 security configuration, the described dual-tier enterprise profile is protected by a comprehensive security policy and design at the headquarters site. In addition, some considerations exist for the branch sites.

The enterprise dual-tier profile does not have a Cisco IOS Firewall that is configured for IPv6.

The branch LAN interfaces can have an ingress ACL to permit traffic from the VLAN interfaces, based on source prefix or on specific applications.

This part of the IPv6 security configuration is for ISR 2900-br1-1.

```
interface FastEthernet0/0.100
 description DATA VLAN for PCs
 ipv6 traffic-filter DATA LAN-v6 in
 no ipv6 redirects
 no ipv6 unreachables
!
interface Serial0/1/0
no ipv6 redirects
 no ipv6 unreachables
!
no ip http server
!
ipv6 access-list MGMT-IN
 remark permit mgmt only to loopback
 permit tcp 2001:DB8:CAFE::/48 host
2001:DB8:CAFE:2000::BAD1:1010
 deny ipv6 any any log-input
!
 ipv6 access-list DATA LAN-v6
 remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX
2001:DB8:CAFE:2100::/64
 permit icmp 2001:DB8:CAFE:2100::/64 any
 remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX
2001:DB8:CAFE:2100::64
 permit ipv6 2001:DB8:CAFE:2100::/64 any
 remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE
LINK-LOCAL PREFIX
 permit icmp FE80::/10 any
 remark PERMIT DHCPv6 ALL-DHCP-AGENTS REQUESTS FROM HOSTS
 permit udp any eq 546 any eq 547
 remark PERMIT ALL PIM PACKETS FROM OTHER BRANCH ROUTER
 permit 103 FE80::/16 any
 remark DENY ALL OTHER IPv6 PACKETS AND LOG
 deny ipv6 any any log-input
!
banner login ^C
```


© 2010 Cisco Systems, Inc.

https://t.me/learningnets

```
Unauthorized access to this device and/or network is
prohibited.
^C
line vty 0 4
 ipv6 access-class MGMT-IN in
 transport input ssh
!
```

This part of the IPv6 security configuration is for switch 3560-br1-1.

```
interface Vlan100
 ipv6 address 2001:DB8:CAFE:2100::BAD2:F126/64
!
ipv6 access-list MGMT-IN
 permit tcp 2001:DB8:CAFE::/48 host
2001:DB8:CAFE:2100::BAD2:F126
 deny ipv6 any any log-input
!
banner login ^C
Unauthorized access to this device and/or network is
prohibited.
^C
!
line vty 0 4
 ipv6 access-class MGMT-IN in
 transport input ssh
!
```

QoS Configuration

This configuration for Cisco 2900 Series Integrated Services Router 2900-br1-1 is meant to show which interfaces have QoS policies applied.

```
interface FastEthernet0/0.100
 description DATA VLAN for PCs
 service-policy input BRANCH-LAN-EDGE-IN
 service-policy output BRANCH-LAN-EDGE-OUT
!
interface FastEthernet0/0.200
 description Voice VLAN for IP Phones
 service-policy output BRANCH-LAN-EDGE-OUT
!
interface FastEthernet0/0.300
 description PRINTER VLAN
 service-policy input BRANCH-LAN-EDGE-IN
 service-policy output BRANCH-LAN-EDGE-OUT
!
interface Serial0/1/0.10 point-to-point
 description TO FRAME-RELAY PROVIDER
 frame-relay interface-dlci 10
  class QOS-BR1-MAP
!
map-class frame-relay QOS-BR1-MAP
```


https://t.me/learningnets

```
 service-policy output WAN-EDGE-FRTS
!
policy-map WAN-EDGE-FRTS
 class class-default
  shape average 1460000 14600 0
  service-policy BRANCH-WAN-EDGE
```


© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

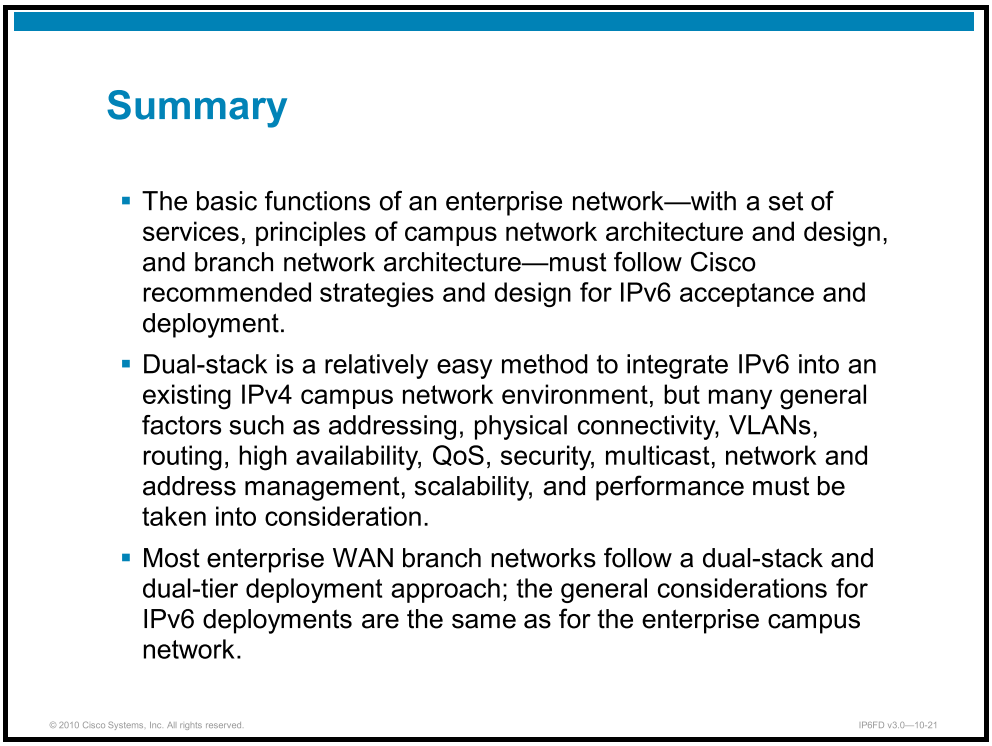

Resources

For additional information, refer to this resource:

- Cisco Systems, Inc. *Cisco Validated Designs.* http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns741/networking_solutions_products_genericcontent0900aecd80601e22.html.

https://t.me/learningnets

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Planning and Implementing IPv6 in Service Provider Networks

Overview

Trends in service provider networking are moving towards an acceptance of IP version 6 (IPv6). This trend is providing many home and business customers with native IPv6 connectivity. Service providers have many considerations when planning and implementing IPv6 in their core and access networks. These considerations range from technology to introduction issues. Feature testing and staging, backed up by special practice experience and trusted advice, are vital when planning the move and addition of IPv6. This lesson will provide you with a practical overview of industry best practices and most common scenarios of implementing IPv6 in service provider core networks.

Objectives

Upon completing this lesson, you will be able to describe and identify the most common planning and implementation approaches that pertain to moving to IPv6 in service provider core networks. This ability includes being able to meet these objectives:

- Identify the design principles for a service provider network
- Identify deployment considerations for IPv6 for Internet access services
- Identify deployment considerations for native IPv6 support in the service provider core network
- Identify deployment considerations for IPv6 over MPLS in the service provider core network

https://t.me/learningnets

Service Provider Network Design

This topic defines the service provider network and its characteristics.

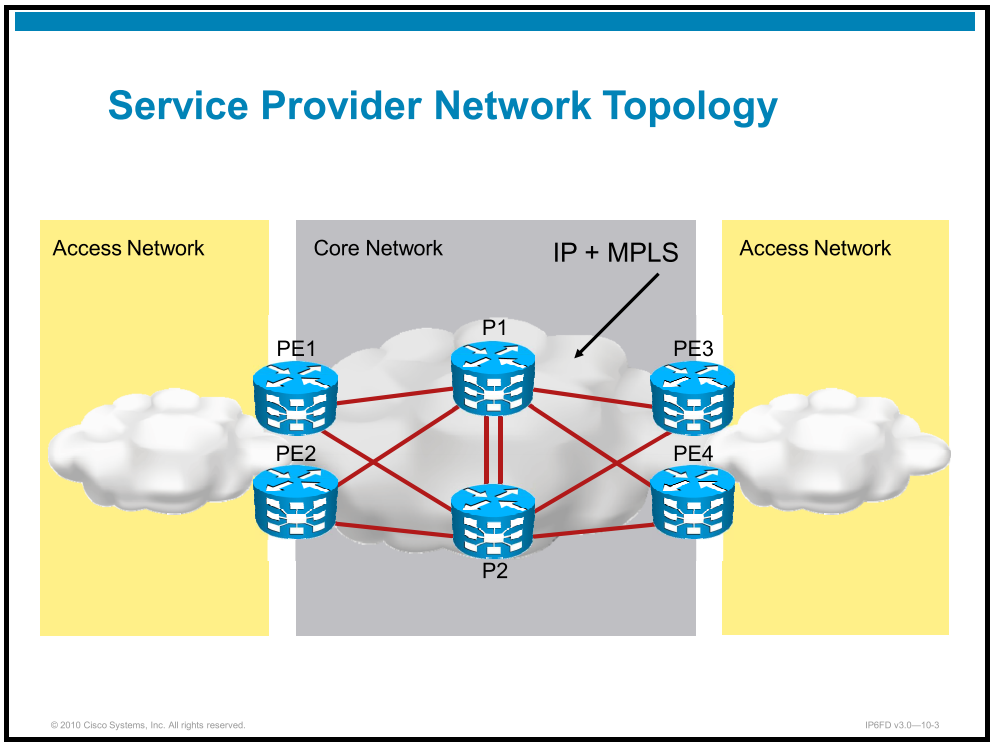

The service provider network consists primarily of two parts: the access network, with its main function of connecting various customers, and the core network, which is used for fast switching of traffic, interconnecting customers and their locations, and providing access to the Internet. Traditionally, the service provider access and core networks are both IP version 4 (IPv4)-enabled. With the acceptance of IPv6, you must take into account many additional considerations to successfully introduce the new IPv6-enabled services.

The service provider network must suffer no negative impact from the adoption of new IPv6 capabilities and must still successfully answer other expectations for anywhere, anytime access; increased capacity; and complexity.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

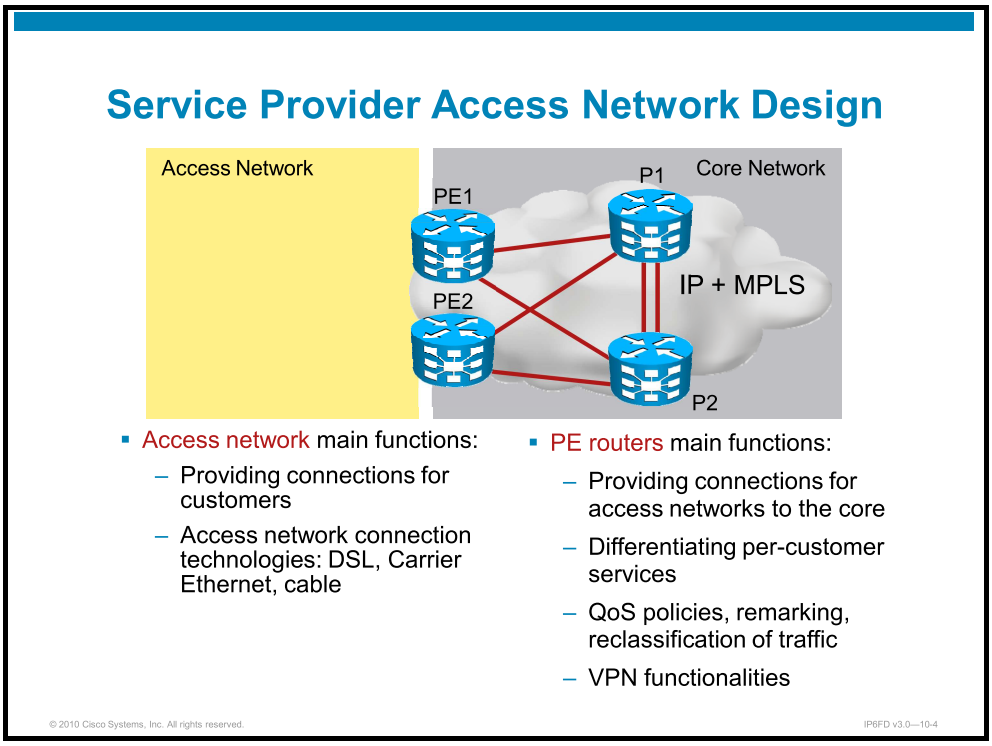

Modern service provider access networks provide myriad capabilities for connecting customers, businesses, and enterprises to the service provider core network. These access network connection technologies include DSL, Carrier Ethernet, cable, WiMax, and many others.

The main complexity of the service provider network design lies at the provider edge (PE) router implementation and architecture principle. Connecting customers to the core service provider network, the PE router differentiates among customer services and specifics. The PE router also supports different quality of service (QoS) policies, including reclassification and remarking of traffic, for different customers. Add virtual private network (VPN) functions (except for VLANs), and the service provider PEs clearly have an exceptionally important role in the overall architecture.

https://t.me/learningnets

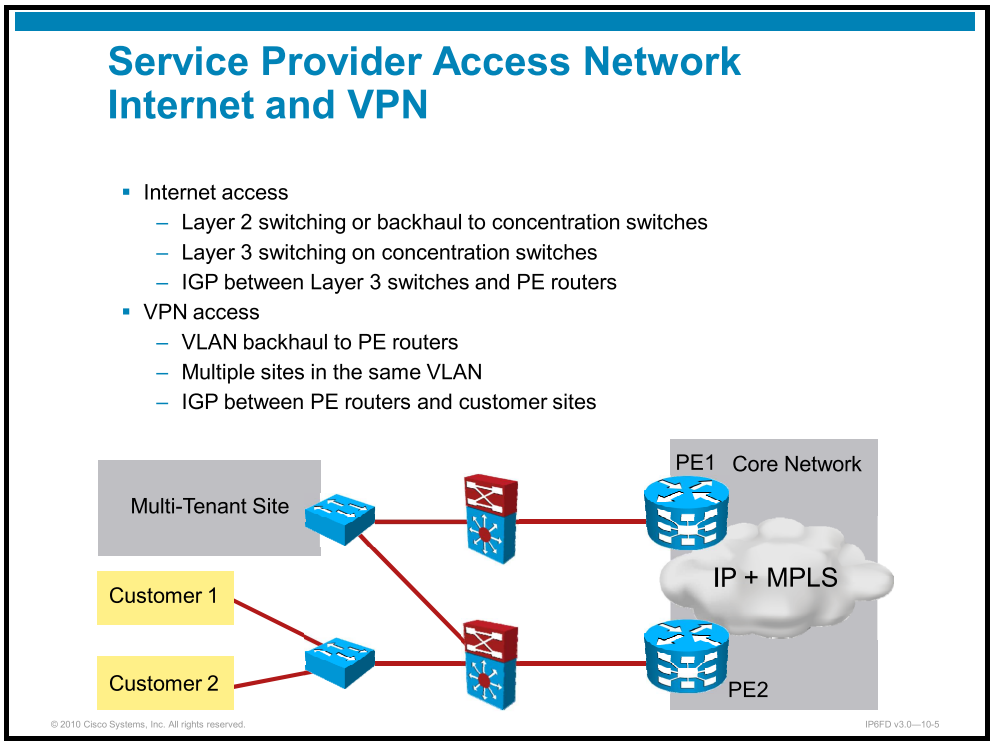

In the access network, most service providers consider IPv6 for Internet access only. Using some form of Carrier Ethernet technology and optical fibers in the access network, Layer 2 switching is usually deployed very close to the customer premises equipment (CPE). This approach separates customers, something that is recommended for the overall service provider core network design. In this design, every customer connection terminates on a Layer 3 switch for the common separation goal. Layer 3 switching usually occurs even before the PE routers on the Layer 3–enabled switch devices. As these Layer 3–enabled switches become part of the Layer 3 network with the PE routers, the switches need to exchange reachability and routing information. Consequently, a service provider will employ some type of interior gateway protocol (IGP) between the Layer 3 switches and PE routers. An important design note: The IGP that is used in the service provider access network should be different from the one that the service provider uses for its core network. The main design challenge is to completely separate and isolate the access and core parts of the service provider network. Border Gateway Protocol (BGP) is used to route customer-specific traffic through the service provider network. The design goal is to achieve a complete, stable core network.

VPN access originates on the PE routers, so VLANs extend from appropriate customers to the PE routers. In some cases, when a service provider access network serves many locations that belong to the same customer, a direct communication can be established between locations in the same VLAN.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Service Provider Core Network Design

- Service provider core network high-speed transport technologies:
  - Gigabit Ethernet, 10 Gigabit Ethernet, SDH, WDM
- P routers main function
  - Very high-speed traffic forwarding

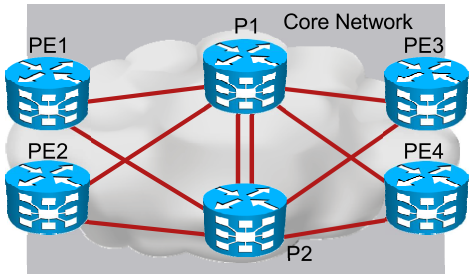

© 2010 Cisco Systems, Inc. All rights reserved.

In the core network, service providers use a high-speed transport technology, such as Gigabit Ethernet, 10 Gigabit Ethernet, Synchronous Digital Hierarchy (SDH), or wavelength-division multiplexing (WDM). Core routers are relieved of all details regarding services and additional functions, for one main reason: for fast traffic forwarding and to use their processing strength for vast amounts of packets that are switched through the service provider core network.

https://t.me/learningnets

Service Provider Core Network Functional Design

- Service provider core network overall functional design:
  - BGP enabled on PE routers
  - BGP carrying customer routes
  - Separate core IGP
  - BGP next hops and core links in core IGP
  - BGP scalability with route reflector
- IPv6 native transport in service provider core network
  - All core routers (P, PE) run BGP
- Cisco 6PE transport in service provider core network
  - Only PE routers run BGP
  - MPLS for transport between BGP next hops
  - LDP or Cisco MPLS TE across the core

© 2010 Cisco Systems, Inc. All rights reserved.

The main goals when designing the service provider core network are resiliency and stability. Events in the service provider access network should have no impact on the resilience and high-availability functional design of the service provider core network.

Customer routing information is transferred through the service provider core network by using BGP rather than the core IGP. Separating the IGP in the service provider core from the service provider access network is crucial. The IGP in the service provider core is used mostly for routing logical loopbacks, WAN links, and traffic engineering.

Because of the separation of these two routing protocols, BGP must be enabled on all PE routers, and route reflector functionality must be employed on some of the provider (P) routers or on a dedicated device to emphasize the service provider core network scalability.

With an IPv6-enabled service provider core network, two possibilities emerge:

- IPv6 native transport in the service provider core network
- Cisco IPv6 PE Router over Multiprotocol Label Switching (MPLS) (6PE) transport in the service provider core network

With the IPv6 native transport in the service provider core network, an IPv6-enabled P router examines the IPv6 packet headers and makes proper routing decisions that are based on the destination IPv6 address. IPv6 native transport is primarily employed by service providers that offer Internet access. In such cases, BGP is enabled on all core routers, both P and PE.

With the 6PE transport in the service provider core network, P routers are unaware of IPv6. P routers make forwarding and switching decisions according to MPLS labels. The 6PE transport scenario employs BGP on PE routers only; MPLS is used to connect BGP next hops. BGP next-hop connections are achieved by using Label Distribution Protocol (LDP) or Cisco MPLS Traffic Engineering (TE). If service provider network engineers establish an MPLS TE tunnel directly between two PE routers, then LDP is not needed. The main disadvantage to this solution is the possibility of a tunnel failure, which can halt traffic flow between the two involved PE routers.

https://t.me/learningnets

Native IPv6 Deployment in Service Provider Access Networks

In this topic, you will identify the deployment specifics for connecting two types of IPv6-enabled customers to the service provider core network.

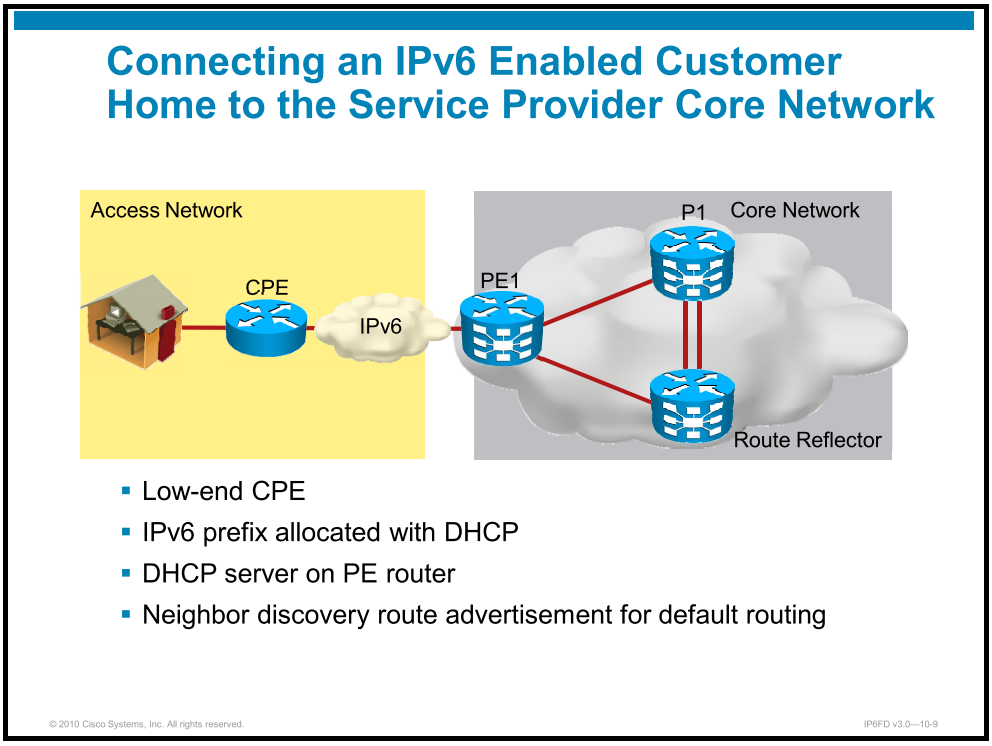

The first of the two scenarios for connecting IPv6-enabled customers to the service provider core network is an IPv6-enabled customer home.

When deploying IPv6 in the service provider access network for customer and telecommuter homes, the CPE that is used in most introductions is a low-end Cisco 800, 1800, or 1900 Series Router.

In IPv4, these devices have only one IPv4 address. In IPv6, telecommuter home routers are given an IPv6 prefix. Prefix lengths vary from /64, which is enough for one LAN network address, to /48, depending on the service provider. For small sites with few LAN networks, a prefix length of /56 or even /60 is viable and reasonable.

The remote site address is usually obtained from the service provider core network.

In IPv4, these are a few of the many options for obtaining an IP address:

- DHCP
- PPP
- Static addressing

In IPv6, prefix allocation is always accomplished by DHCP. The DHCP option for this scenario is IP address prefix delegation, which enables the CPE to obtain not only one IP address but a complete prefix. In this service provider IPv6-enabled access network scenario, the DHCP server for IPv6 resides on the PE router.

https://t.me/learningnets

Another employed IPv6 function is neighbor discovery, which is the IPv6 equivalent of Address Resolution Protocol (ARP) in IPv4. Neighbor discovery in IPv6 also supports route advertisements as a way of notifying devices that they can send traffic to the routing device that is sending out route advertisements. With these route advertisements, the issue of the default gateway is straightforwardly solved in IPv6. The main advantage of route advertisement messages is not only default gateway information but also prefix delegation.

After an IPv6 prefix is obtained via DHCP, a connected route to that prefix network, pointing to a null interface, is added to the routing table on the CPE router. With the IPv6 prefix obtained, the CPE device becomes an owner of that addressing space and processes all traffic that is destined for networks that belong to the space. The CPE router also installs a static default route (::/0) that points to the local loopback address of the PE router.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

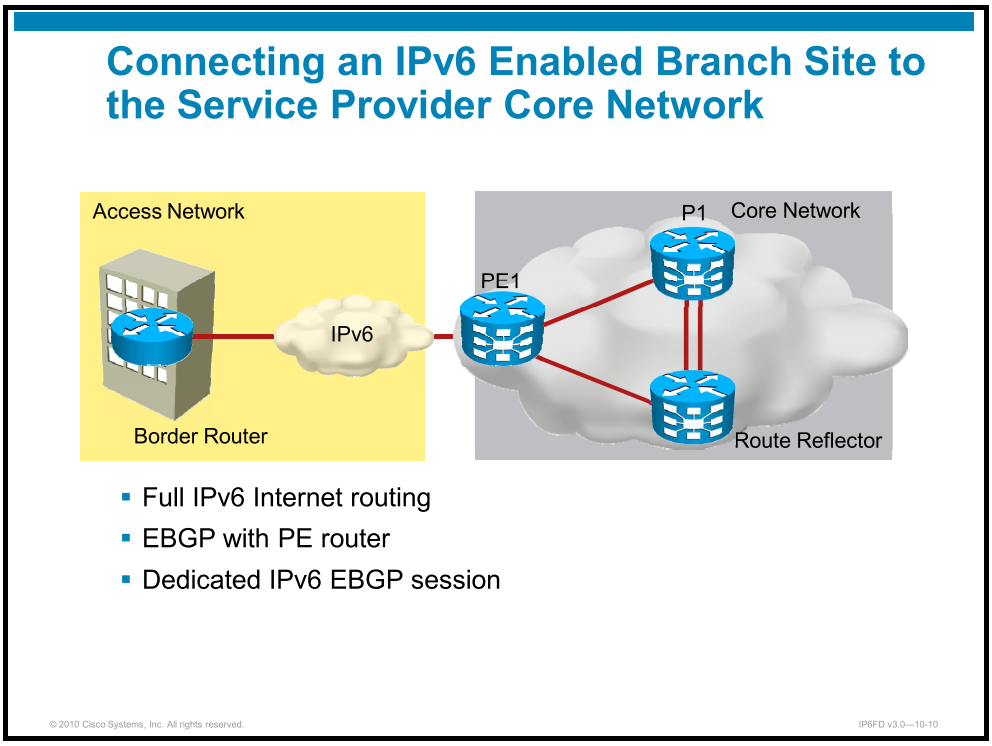

The second of the two scenarios for connecting IPv6-enabled customers to the service provider core network is an IPv6-enabled branch site.

When deploying IPv6-enabled branch sites, the service provider uses complete IPv6 internet routing. External Border Gateway Protocol (EBGP) is used on the border and PE routers for this task. One or two EBGP sessions can occur between the border and PE routers.

The service provider can have two sessions, one for IPv4 and another for IPv6, or can use only one session for both IPv4 and IPv6. When connecting many customer branch sites, two dedicated EBGP sessions is less complex.

On the border router, IPv6 addresses are needed on the loopback, local access interface, and WAN interface that connect to the PE router. Establishing BGP sessions with the IPv6 address of the loopback interface is not recommended. The border router BGP configuration is concerned with adding the IPv6 address of the PE router as a neighbor, enabling a new IPv6 address family, and advertising the network behind the border router through the EBGP process. An IPv6 route that points to the null interface for the advertised network is also needed, like in the configuration for IPv4 EBGP.

On the PE router, the IPv6 border router neighbor is also added to the BGP routing process, and an IPv6 address is configured on the WAN link.

https://t.me/learningnets

Native IPv6 Deployment in the Service Provider Core Network

This topic identifies deployment considerations for native IPv6 support in the service provider core network.

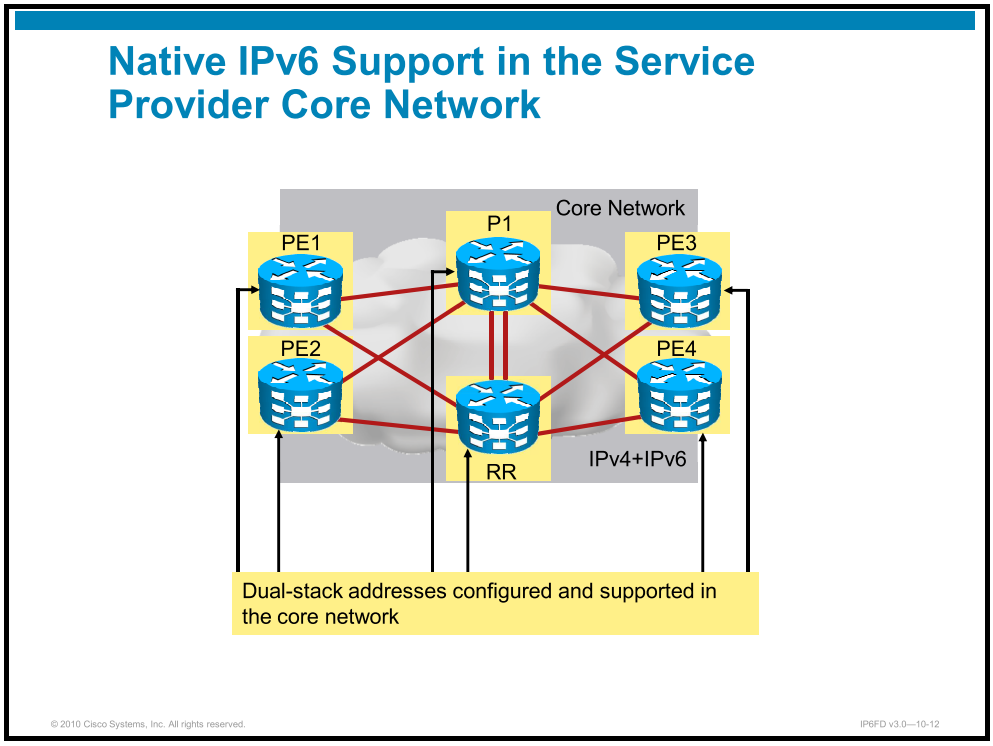

For native IPv6 support in the service provider core network, both IPv4 and IPv6 addresses need to be configured and supported on all routers in the core network.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

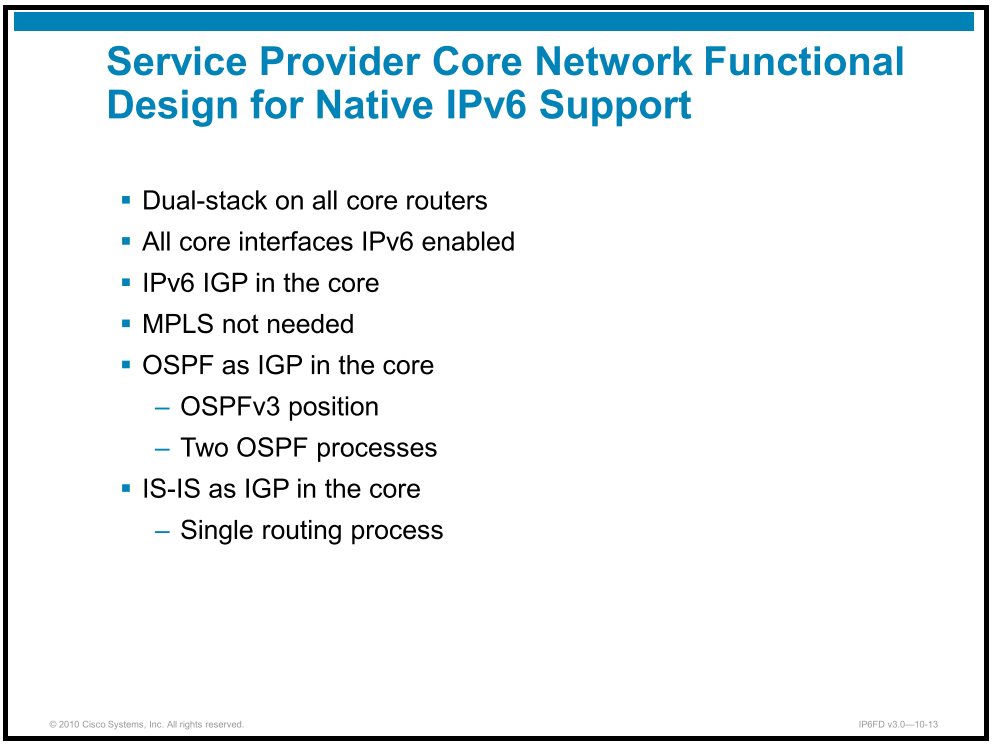

In this IPv6 deployment scenario, all core routers, including the P routers, must have the dual-stack model deployed. IPv6 addresses must be configured on all core interfaces, and two IGPs run in the core network. One IGP is used for IPv4, and the other is used for IPv6. One exception to this rule might be the use of Intermediate System-to-Intermediate System (IS-IS) as an IGP. All core routers (both PE and P) also must execute the BGP routing protocol process and be interconnected with route reflectors.

With native IPv6 deployment in the core network, a single BGP session between the IPv4 endpoints (instead of one session for IPv4 and another for IPv6) can be used.

This solution is recommended for native IPv6 support and can be used in networks in which MPLS is not deployed.

The configuration on the PE routers in the service provider core network follows a typical dual-stack approach. IPv4 and IPv6 addresses are configured on the loopback and physical interfaces that are used for connectivity to the P routers and the access network.

Open Shortest Path First (OSPF) is used as an IGP with two routing processes, one for IPv4 and another for IPv6. Both OSPF processes are enabled for the appropriate PE router interfaces.

https://t.me/learningnets

Service Provider Core Network Functional Design for Native IPv6 Support (Cont.)

- All core routers run BGP:
  - Two parallel BGP sessions: one for IPv4 and one for IPv6
  - Single BGP session between IPv4 endpoints
- Router RR serves as route reflector for both IPv4 and IPv6 address families.
- IPv6 and IPv4 addressing are configured on all P core router interfaces.

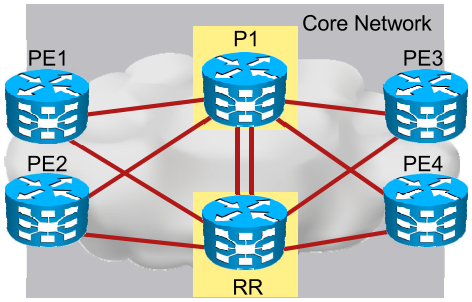

© 2010 Cisco Systems, Inc. All rights reserved.

BGP configuration on the PE is accomplished with the addition of templates for peer policies and peer sessions, allowing for additional scalability.

In the peer session template, you define parameters for the BGP session with the BGP neighbors. In the peer policy template, you define parameters that concern the address family. Peer policy templates are also inheritable.

In the IPv6 address family, BGP configuration network information and redistributed static routes to customer locations are added.

When IPv6 prefixes are sent through an IPv4 BGP session, a problem with the next hop information exists. The solution to this problem is a route map, in which you manually set the IPv6 next-hop value. If you have two distinct BGP sessions, one for IPv4 and another for IPv6, this problem is not observed. In the latter scenario, the number of sessions with router reflectors is increased by a factor of two.

On many service provider core networks, migrating to dual-stack architectures is combined with adopting IS-IS instead of OSPF as an IGP. OSPF version 2 (OSPFv2), or OSPF for IPv4, is a well-known and widely used IGP. For IPv6 routing, OPSF version 3 (OSPFv3) is a completely different protocol. When you consider IPv6 networks with large areas that have a few hundred routers, the presence of two topological databases on each device and two Shortest Path First (SPF) algorithms that run in parallel after each topology change is important.

However, IS-IS enabled devices run only one routing protocol, one topology database, and one SPF calculation after each change. Not considering the possible issues with IS-IS in dual-stack service provider core implementations, the configuration is straightforward, enables IS-IS, and adds the routing process on applicable interfaces for both IPv4 and IPv6.

https://t.me/learningnets

Route reflector router configuration for BGP sessions also follows well-defined best practices. In the peer policy configuration, you enable the sending of both the communities and designate all neighbors as route reflector clients. You add each route reflector neighbor in the core network to a BGP session, which inherits the parameters of the BGP session template. For each of the two IPv4 and IPv6 address families, you activate the route reflector neighbors and add the defined peer policy template. This approach allows you to have devices that are route reflectors for IPv6 but not for IPv4.

Enabling native IPv6 is also necessary on the P routers. Configure IPv4 and IPv6 addresses on all active interfaces. Two OSPF router processes are also configured: one for IPv4 and another for IPv6.

https://t.me/learningnets

6PE Deployment in the Service Provider Core Network

This topic identifies deployment considerations for IPv6 over MPLS in service provider core networks.

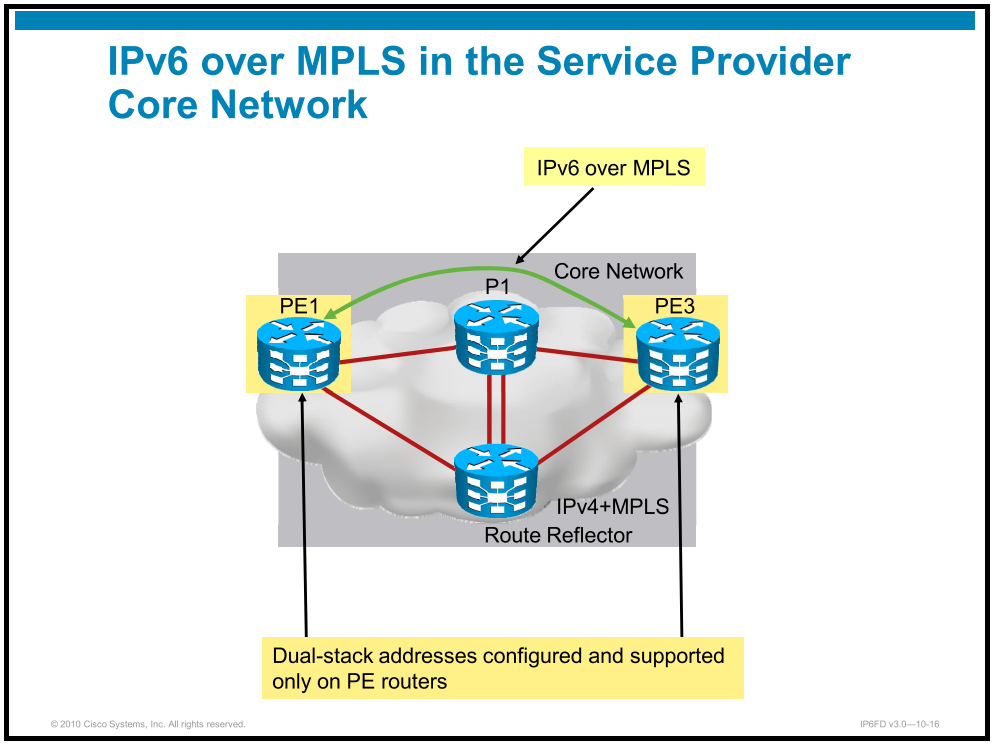

For IPv6 over MPLS support in the service provider core network, only IPv4 and MPLS are configured and supported on all routers in the core network.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Service Provider Core Network Functional Design for IPv6 over MPLS Support

- IPv6 packets transported over MPLS LSP
- BGP to propagate IPv6 prefixes and labels
- Two MPLS labels:
  - Top label LDP or MPLS-TE towards BGP next hop
  - Bottom label BGP for IPv6 prefix
- Core is IPv4 and BGP only
- IPv6 on PE routers and route reflectors
- IPv4 in the service provider core network
- IPv4 IGP in the service provider core network

© 2010 Cisco Systems, Inc. All rights reserved.

In transporting IPv6 over MPLS in the service provider core network, IPv6 packets are transported over an MPLS label-switched path (LSP). BGP is enabled on PE routers only; P routers are IPv4-enabled only. BGP is used to propagate IPv6 prefixes and their labels. PE routers have no virtual routing and forwarding instances (VRFs) but do have a global IPv6 table.

This functional design employs two MPLS labels: a top label, toward the BGP next-hop that is obtained via LDP or MPLS TE, and a second label that is exchanged directly between PE routers for the IPv6 prefix.

Consequently, you do not have IPv6 in the core network, which is IPv4 and MPLS only.

You enable IPv6 on PE routers and route reflectors. Configure only IPv6 routing, not IPv6 interfaces, on the route reflectors.

When enabling the service provider core network for IPv6 over MPLS support, you enable IPv6 on PE routers and route reflectors only. IPv6 addresses are configured on the PE router loopbacks and provider edge-customer edge (PE-CE) links only. All core links have IPv4 addressing only.

The routing protocol in the core service provider network is an IGP such as OSPF or IS-IS.

LDP also functions in the core network and is used for transporting packets.

IBGP is used between IPv4 endpoints and label sending is enabled on IBGP neighbors in the IPv6 address family.

https://t.me/learningnets

Service Provider Core Network Functional Design for IPv6 over MPLS Support (Cont.)

- BGP next-hops label distribution
- Core IBGP between IPv4 endpoints
- Sending labels on IBGP neighbors in IPv6 address family
- BGP on route reflector router with both address families IPv4 and IPv6
  - Label sending on IPv6 peer policy template
  - BGP neighbors accessible via IPv4 only
  - MPLS label for IPv6 prefixes in BGP
- Label sending for IPv6 address family on route reflector

© 2010 Cisco Systems, Inc. All rights reserved.

On PE routers, BGP implementation label sending is added for the IPv6 peer policy template. That template is used when defining BGP neighbors for the IPv6 address family. IBGP neighbors are accessible only via IPv4.

With IPv6 prefixes in BGP, MPLS labels also occur. Checking the MPLS forwarding table for a label reveals that it is connected with a certain IPv6 prefix.

When you define the peer policy template for IPv6 on the route reflector, you also define label sending.

After checking the IPv6 prefix entry on the PE routers, you can see an MPLS label in the BGP table. After checking the Cisco Express Forwarding table, you can see that the IPv6 prefix has two labels. One label originates from the BGP, and the other is obtained by LDP from the P router. The same rules apply on the other PE router.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

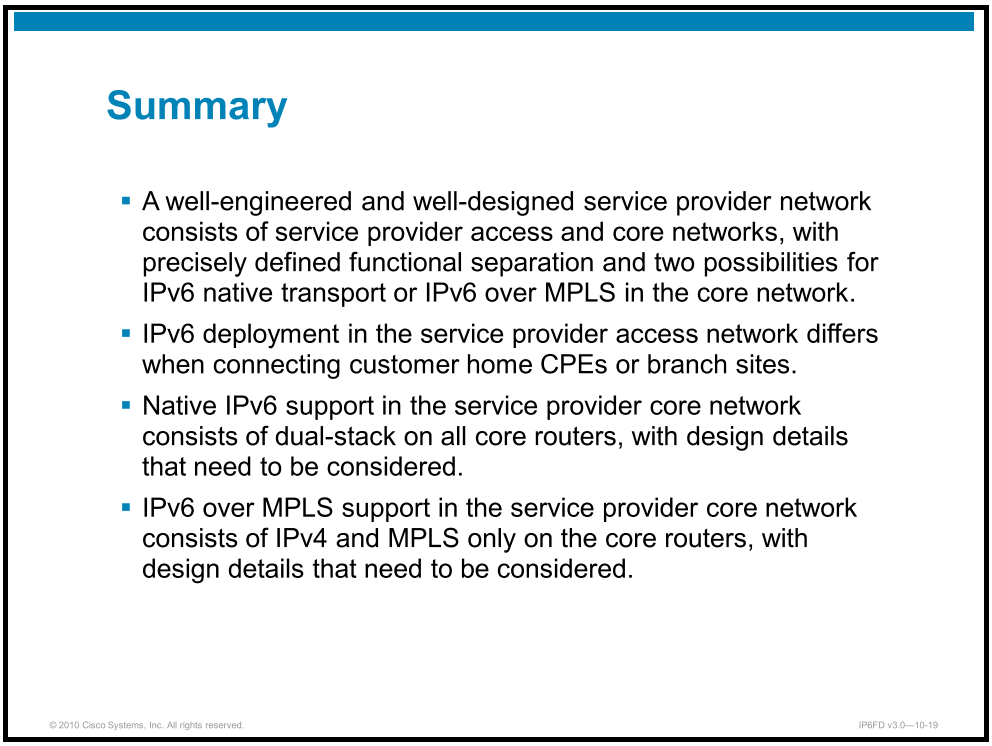

https://t.me/learningnets

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3

Planning and Implementing IPv6 in Branch Networks

Overview

The concerns and issues that pertain to the transition to IP version 6 (IPv6) in branch networks are like those that pertain to the enterprise campus network. To retain network integrity, end-to-end functionality, and all-around effectiveness, the punctual and methodical approach that is valid for the campus enterprise network planning and implementation must be followed closely when planning for and deploying IPv6 in the branch. In this lesson, you will familiarize yourself with the basic Cisco branch design as it pertains to IPv6. You will learn about the most commonly deployed Cisco branch profiles and how IPv6 is integrated into these profiles.

Objectives

Upon completing this lesson, you will be able to describe and identify the most common planning and implementation approaches as they pertain to moving to IPv6 in branch networks. This ability includes being able to meet these objectives:

- Describe the general considerations when deploying IPv6 in branch networks
- Describe the single-tier profile when deploying IPv6 in branch networks
- Describe the dual-tier profile when deploying IPv6 in branch networks

https://t.me/learningnets

Branch Deployment Overview

This topic identifies general deployment considerations for IPv6 in branch networks.

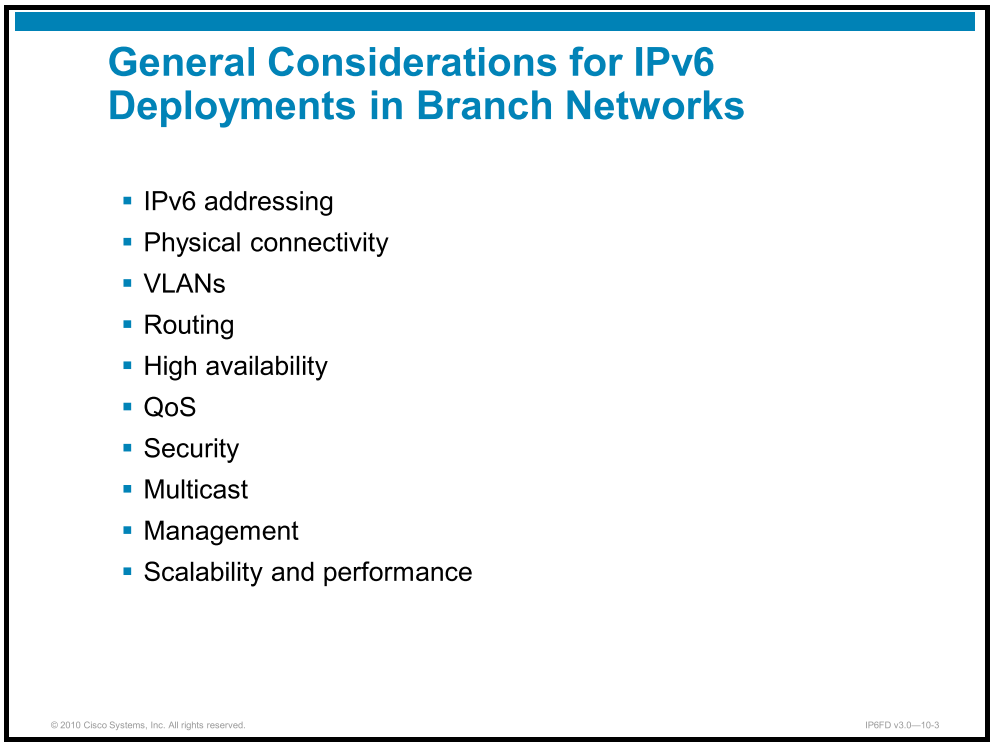

General considerations for IPv6-enabled branch networks can be compiled into these categories:

- IPv6 addressing
- Physical connectivity
- VLANs
- Routing
- High availability
- Quality of service (QoS)
- Security
- Multicast
- Management
- Scalability and performance

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

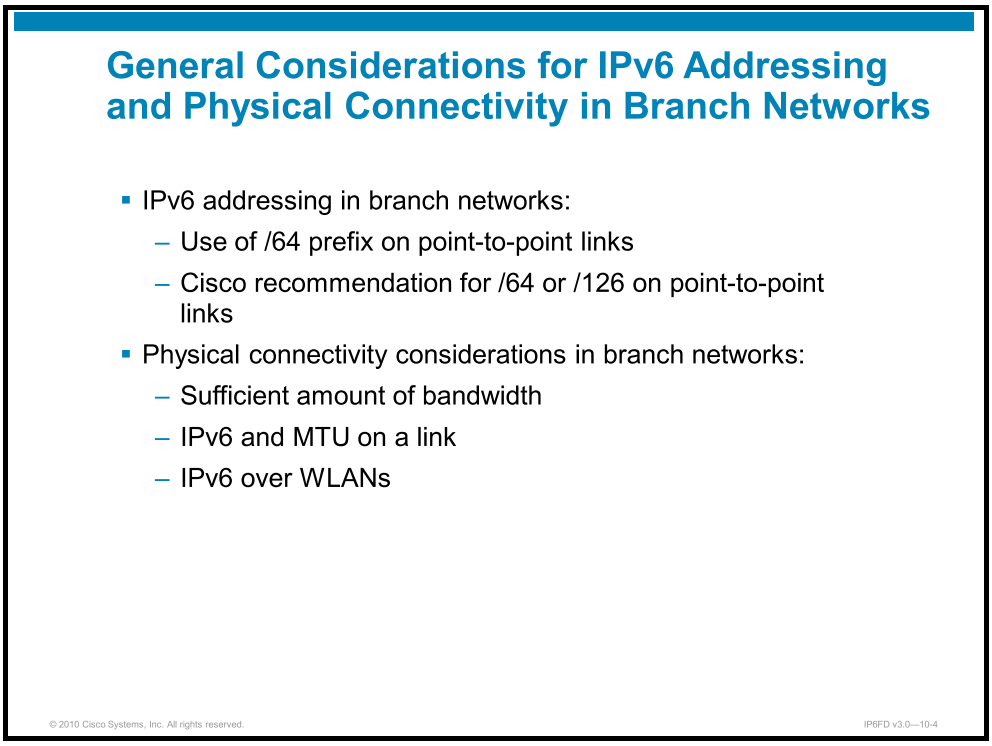

In most cases, the use of a /64 prefix on point-to-point links is acceptable. IPv6 was designed to have a large address space, and even with poor address management in place, the customer should not experience address constraints.

There is quite a bit of discussion within the IPv6 community about the practice of using longer prefixes for point-to-point links. Network administrators who want to more tightly control the address space can use a /126 prefix on point-to-point links, much like using a /30 prefix with IP version 4 (IPv4).

RFC 3627 describes why the use of a /127 prefix is harmful and should be discouraged.

In general, Cisco recommends the use of either a /64 or /126 prefix on point-to-point links. Efforts are underway within the Internet Engineering Task Force (IETF) to better document the address assignment guidelines for varying address types and prefix links.

The assignment of user IPv6 addresses in single- and dual-tier profiles is done by advertising an IPv6 prefix (via a route advertisement) on the router subinterface for the VLAN on which the PCs are located. The options for a Domain Name System (DNS) server and domain name are assigned by using DHCP for IPv6. All other VLANs use stateless autoconfiguration, with no use of options.

Considerations for physical connectivity with IPv6 are the same as with IPv4, with the addition of three elements:

- One important factor for the deployment of any new technology, protocol, or application is the assurance of sufficient bandwidth for both existing and new traffic. This issue is especially true with a branch site because the connections to the WAN are often low-speed links, and QoS can do only so much to solve bandwidth problems.
- Understanding how IPv6 manages maximum transmission unit (MTU) on a link is another consideration. Another aspect of MTU relates to the branch single-tier profile. When IP Security (IPsec) is used with Generic Routing Encapsulation (GRE) or manual tunnels, you must account for how to adjust the MTU value on the routers to ensure that the router is not forced to perform fragmentation of the IPv4 traffic because of the IPsec header and the additional tunnel overhead.

https://t.me/learningnets

- IPv6 should operate correctly over WLAN access points (APs), much like IPv6 operates over Layer 2 switches. However, considerations to IPv6 with WLAN environments include managing WLAN devices, such as APs and controllers, via IPv6 and controlling IPv6 traffic via AP or controller-based QoS, VLANs, and access control lists (ACLs). IPv6 must be supported on the AP and controller devices to take advantage of these more intelligent services on the WLAN devices.

Cisco supports the use of IPv6-enabled hosts that are directly attached to Cisco IP phone ports. These IP phone ports are switch ports and operate much like plugging the host directly into a Cisco Catalyst Layer 2 switch.

In addition to the previous considerations, Cisco recommends a thorough analysis of the existing traffic profiles, memory, and CPU use on the hosts and network equipment, as well as of the service level agreement (SLA) language, before the implementation of any of the previously described IPv6 models.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

General Considerations for IPv6 VLANs, Routing, and High Availability

- Same VLAN considerations for IPv4 and IPv6
- EIGRP for IPv4 and IPv6 in the branch network designs
- EIGRP authentication for both IPv4 and IPv6
- Redundant WAN connections
- Redundant routing and forwarding paths
- High availability of the first-hop gateways

© 2010 Cisco Systems, Inc. All rights reserved.

VLAN considerations for IPv6 are the same as for IPv4. When dual-stack configurations are used, both IPv4 and IPv6 traverse the same VLAN. The use of private VLANs is not included in any of the previously described deployment profiles.

The use of IPv6 on data VLANs that are trunked with voice VLANs behind IP phones is fully supported.

Choosing an interior gateway protocol (IGP) to run in the campus network depends on various factors, platform capabilities, IT staff expertise, and the size of the network, to name just a few. The IGP for both IPv4 and IPv6 in the considered branch network designs was Enhanced Interior Gateway Routing Protocol (EIGRP). Open Shortest Path First version 2 (OSPFv2) for IPv4 and OSPF version 3 (OSPFv3) for IPv6 can be also used.

As previously mentioned, every effort has been made to implement the current Cisco branch design best practices. Both the IPv4 and IPv6 IGPs should be tuned according to the current best practices for the branch. A top priority of any network design should be to ensure that the IGPs are tuned to provide a stable, scalable, and fast converging routing protocol.

EIGRP has been configured to provide authentication for both IPv4 and IPv6 adjacencies and updates.

Many of the high availability requirements and recommendations are met by leveraging the existing Cisco branch design best practices. These are the primary high availability components for IPv6-enabled branch networks:

- Redundant WAN connections are established in the single-tier profile. The primary WAN connection is a T1 connection to the ISP and the secondary connection is an asymmetric DSL (ADSL) connection to another ISP. However, both of these links come from only one branch integrated services router (ISR). In the dual-tier profile, each of the two branch ISRs has a Frame Relay connection to the private WAN.
- Redundant routing and forwarding paths are accomplished by using EIGRP for IPv4 and IPv6. In some cases, Equal Cost Multipath (ECMP) is used, and in other cases IPsec GRE and manual tunnels are used; one path is preferred over another, but the secondary path is available for redundancy.

https://t.me/learningnets

- High availability of the first-hop gateways applies only to the dual-tier profile. Hot Standby Router Protocol version 2 (HSRPv2) for IPv4 and IPv6 is used to provide first-hop gateway redundancy in the dual-tier profile. Cisco also supports Gateway Load Balancing Protocol (GLBP) for IPv4 and IPv6.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

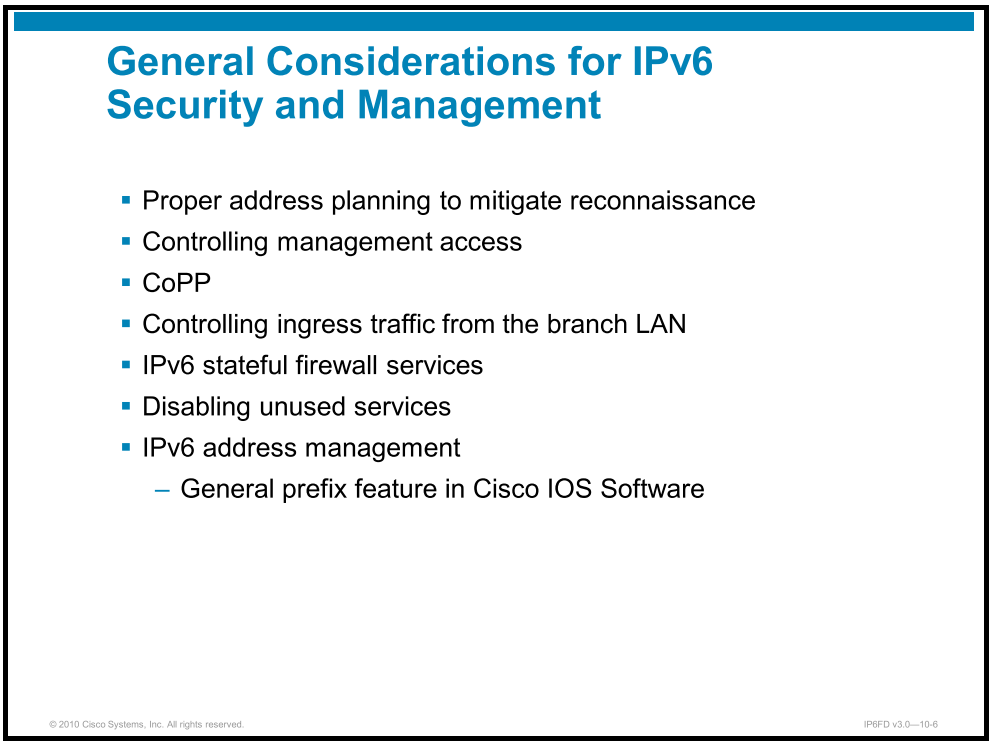

Many of the common threats and attacks on existing IPv4 campus networks also apply to IPv6. Unauthorized access, spoofing, routing attacks, viruses, worms, denial of service (DoS), and man-in-the-middle attacks are just a few issues that can plague both IPv4 and IPv6.

Many new threats to IPv6 do not affect IPv4 or operate differently in IPv4. There are inherent differences in how IPv6 processes neighbor and router advertisement and discovery, headers, and even fragmentation. Based on all of these variables and possibilities, IPv6 security is a very involved topic. Numerous efforts within both Cisco and the industry attempt to identify, understand, and resolve IPv6 security threats.

General security considerations for network device protection apply to both branch profiles.

- Reconnaissance should be made more difficult through proper address planning for campus switches:
  - Addressing of branch network devices (switches and router) should be well thought out and planned. The common recommendation is to devise an addressing plan so that the 64-bit interface identifier of the router is a value that is random across all of the devices. An example of a bad interface identifier for a device is if VLAN 2 had an address of 2001:DB8:CAFE:2::1/64 and VLAN 3 had an address of 2001:DB8:CAFE:3::1/64, where ::1 is the interface identifier of the router. This addressing is easily guessed and can allow an attacker to quickly understand the common addressing for the branch infrastructure devices. A better choice is to randomize the interface identifier of all the network devices in the branch. Using the previous VLAN 2 and VLAN 3 examples, a new address can be constructed by using an address such as 2001:DB8:CAFE:2::A010:F1A1 for VLAN 2 and 2001:DB8:CAFE:3::C801:167A for VLAN 3, where A010:F1A1 is the interface identifier of VLAN 2 for the router.
  - To ease operational management of the network devices and addressing, the security aspects should balance the randomizing of the interface identifiers with the ability to deploy and manage the devices via the randomized addresses.

https://t.me/learningnets

- Control management access to the branch routers and switches:
  - All the branch routers and switches for each profile have configurations in place to provide management access protection to the devices. All routers have loopback interfaces that are configured for management and routing purposes. The IPv6 address for the loopback interfaces uses the previously mentioned addressing approach of avoiding well-known interface identifier values.
  - To more tightly restrict access via IPv6 to a particular switch or router, an ACL is used to permit access to the management interface (line vty) by way of the loopback interface. The permitted source network is from the enterprise IPv6 prefix. To make ACL generation more scalable for a wide range of network devices, the ACL definition can permit the entire enterprise prefix as the primary method for controlling management access to the device, instead of filtering to a specific interface on the device.
- Control Plane Policing (CoPP) protects the router by preventing DoS or unnecessary traffic from negatively affecting CPU resources. Priority is given to important control plane and management traffic. The CoPP configuration is based on a variety of factors. No one deployment recommendation can be made because the specifics of the policy are determined on a case-by-case basis.
- Control ingress traffic from the branch LAN. Filter the prefixes that are allowed to source traffic. This task is usually done on ingress on the LAN or subinterface on the branch router. Controlling IPv6 traffic according to the source prefix can help to protect the network against basic spoofing.
- IPv6 stateful firewall services provide a stateful security inspection for IPv6 traffic that enters or leaves a branch network. The Cisco IOS Firewall for IPv6 has fewer features than that of IPv4. Specifically, advanced application inspection and control does not yet support IPv6.
- Many services, such as HTTP server, are supported for IPv4 and IPv6. Enabling or disabling these services generally applies to both protocols.

Management for IPv6 is still under development. Many of the traditional management tools that are used also support IPv6. The only considerations for management of the branch network are related to basic control of management services such as Telnet, Secure Shell (SSH), and Simple Network Management Protocol (SNMP). All of the IPv6-enabled devices in the two branch profiles that are described can be managed via these services (except for SNMP) over IPv6. The management of IPv6-specific MIBs, traps, and informs is supported on the Catalyst platforms by using SNMP transport over IPv4. All Cisco ISRs support SNMP over IPv6 transport.

The deployment of SNMP for IPv6 is the same as for IPv4. In the described branch profiles, SNMP version 3 (SNMPv3) AuthNoPriv is used to provide polling capabilities for the Cisco network management stations (NMS) in the headquarters data center.

Another important area of management is very address management. Anyone who analyzed IPv6, even at an elementary level, understands the size and potential complexity of deploying and managing the IPv6 address space. The process of assigning large hexadecimal addresses to many network devices should, at some point, be automated or at least made more user friendly than it is today. Several efforts are underway within the industry to provide recommendations and solutions to the address management issues.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

One way to help with the deployment of address prefixes on a Cisco ISR is by using the general prefix feature. The general prefix feature allows you to define a prefix or prefixes with a user-friendly name in the global configuration of the router. That user-friendly name can be used, on a per-interface basis, to replace the usual IPv6 prefix definition on the interface.

https://t.me/learningnets

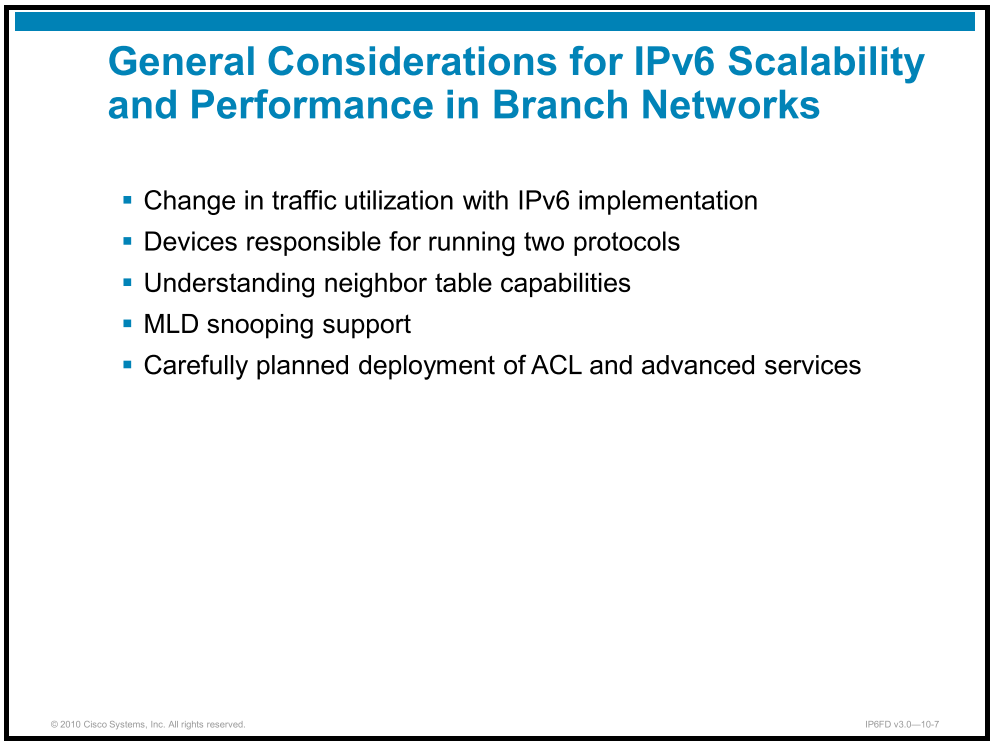

Understanding the link, memory, and CPU use of the existing branch network devices is important. If any of these aspects are already stressed, then adding IPv6 or any new technology, feature, or protocol into the design is not recommended.

IPv6 implementations often experience a change in traffic utilization ratios on the branch network links. As IPv6 is deployed, IPv4 traffic utilization is often reduced as users leverage IPv6 as the transport for applications that were historically IPv4-only. There is often a slight increase in overall network utilization; this increase usually derives from control traffic for routing and from tunnel overhead (single-tier) when manually configured tunnels are used.

One of the primary scalability considerations is that of running two protocols on the router. The branch LAN router must track both IPv4 and IPv6 neighbor information. Like ARP in IPv4, neighbor cache exists for IPv6. The primary consideration is that with IPv4, there is usually a one-to-one mapping of an IPv4 address-to-MAC address, but with IPv6, the host can have several mappings for multiple IPv6 addresses (such as link-local, unique-local, and multiple global addresses) to one MAC address in the router neighbor cache.

The number of entries can decrease to a minimum of one (link-local address) from a multitude of entries for a single host, depending on the address types that are used on the host.

It is important to understand the neighbor table capabilities of the branch network devices, to ensure that the tables are not being filled during regular network operation. Additional testing is planned to understand whether recommendations should be made to adjust timers to age out entries faster, rate-limit neighbor advertisements, and to better protect the branch routers against DoS from IPv6 neighbor discovery–based attacks.

Another consideration involves IPv6 multicast. It is important to ensure that Multicast Listener Discovery (MLD) snooping is supported in the branch LAN switch when IPv6 multicast is used, to ensure that IPv6 multicast frames at Layer 2 are not flooded to all of the ports.

Understanding the routing and forwarding capabilities of the branch routers is important. If the existing branch router is already running at high CPU and memory utilization rates for the handling of IPv4 routing tables and updates, then adding IPv6 to the existing router is a bad idea.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

It is imperative that the deployment of ACLs is carefully planned. IPv6 ACLs in the branch routers are used for QoS (classification and marking of ingress packets from the access layer), security (controlling DoS, snooping, and unauthorized access for ingress traffic in the access layer), and a combination of QoS and security (to protect the control plane of the router from attack). The router can also provide Cisco IOS stateful firewall services, intrusion detection system (IDS) or intrusion prevention system (IPS), voice services for IPv4, and new services for IPv6. Advanced services that are added to the branch router should support both IPv4 and IPv6. Performance will be affected by all of these added services plus the newly enabled IPv6 configuration.

https://t.me/learningnets

Branch Deployment Profiles: Single-Tier Profile Implementation

This topic identifies single-tier profile deployment specifics, through an example of deploying IPv6 in a branch network.

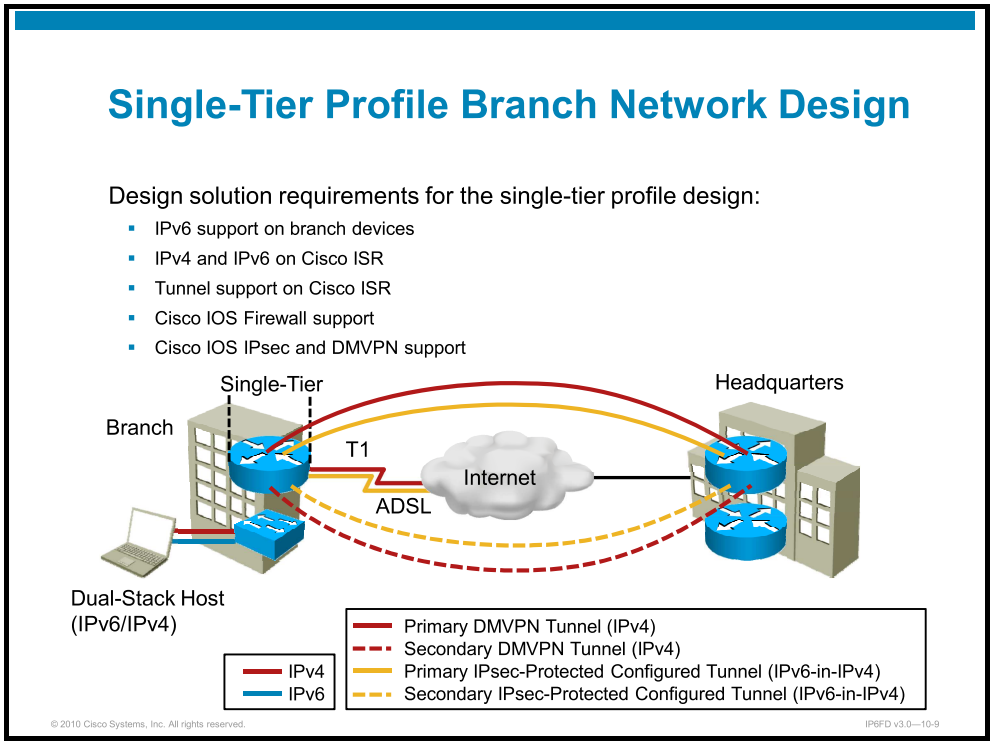

There are three Cisco best practice design specifications for most commonly deployed branch networks:

- Single-tier profile
- Dual-tier profile
- Multitier profile

Each of these branch deployment profiles has profile-defined specifics and concerns that must be addressed carefully when planning for and implementing IPv6 in the branch network.

The single-tier branch profile is a fully integrated solution. The requirements for LAN and WAN connectivity and security are met by a single ISR.

In the single-tier profile that is described in this example, one ISR is used to provide WAN connectivity, via T1, to an ISP. This T1 connection is used as the primary link to the headquarters site. For WAN redundancy, a backup connection is made via ADSL. The single-tier profile uses what is often referred to as the Internet Deployment Model.

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

IPv4 connectivity to the headquarters site is provided by IPv4 IPsec, using Dynamic Multipoint Virtual Private Network (DMVPN) technologies. IPv6 connectivity to the headquarters site is provided by using manually configured IPv6-in-IPv4 tunnels that are protected by IPv4 IPsec. The DMVPN and manually configured tunnels traverse the T1 link as the primary path and establish backup tunnels over the ADSL link. In this sample single-tier profile, IPv6 connectivity via the IPsec-protected, manually configured tunnels are required because DMVPN does not yet support IPv6. When DMVPN supports IPv6 within the design, then no additional tunnel configurations are required and IPv4 and IPv6 (dual-stack) are supported within the same DMVPN design.

All traffic that leaves the branch, including Internet-bound traffic, traverses the VPN connections to the headquarters. Cisco does not generally recommend the use of split-tunneling at the branch site. If the customer requires split-tunneling, then Cisco recommends a careful analysis and testing of the routing and security implications of such a deployment.

LAN connectivity is provided by an integrated switch module (Cisco EtherSwitch Service Module). Dual-stack (running both an IPv4 TCP/IP stack and an IPv6 TCP/IP stack) is used on the VLAN interfaces at the branch.

In addition to all the security policies that are in place at the headquarters, local security for both IPv4 and IPv6 is provided by a common set of infrastructure security features and configurations and the use of Cisco IOS Firewall. QoS for IPv4 and IPv6 is integrated into one policy.

The obvious disadvantage of the single-tier profile is the lack of router and switch redundancy. Redundancy exists for the link to the Internet and the VPN connections to headquarters. However, because there is one integrated switch and one router, if either component fails, then the site is completely disconnected from headquarters. The dual-tier profile is the solution for customers that require complete redundancy for all components such as switches, routers, and headquarters connections.

These are the solution requirements for the single-tier profile:

- IPv6 support on the operating system of the host machines in the branch
- IPv6 and IPv4 dual-stack support on the Cisco ISR
- MLD snooping support on the LAN switch or integrated network module (required if using IPv6 multicast)
- Manually configured tunnel (IPv6-in-IPv4) support on the Cisco ISR
- Cisco IOS Software release and feature set that supports Cisco IOS Firewall
- Cisco IOS Software release and feature set that supports IPsec and DMVPN

https://t.me/learningnets

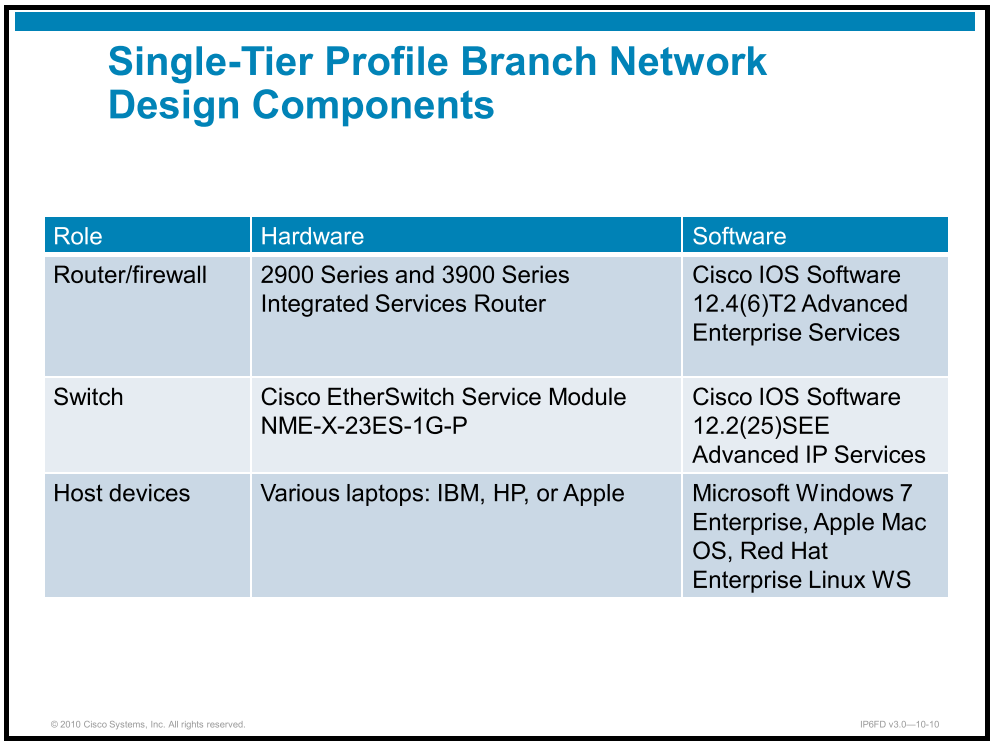

The table lists the single-tier profile network design components:

Single-Tier Profile Branch Network Design Components

Role	Hardware	Software
Router/firewall	Cisco 2900 Series and 3900 Series Integrated Services Router	Cisco IOS Software 12.4(6)T2 Advanced Enterprise Services
Switch	Cisco EtherSwitch Service Module NME-X-23ES-1G-P	Cisco IOS Software 12.2(25)SEE Advanced IP Services
Host devices	Various laptops: IBM, HP, or Apple	Microsoft Windows 7 Enterprise, Apple Mac OS, Red Hat Enterprise Linux WS

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

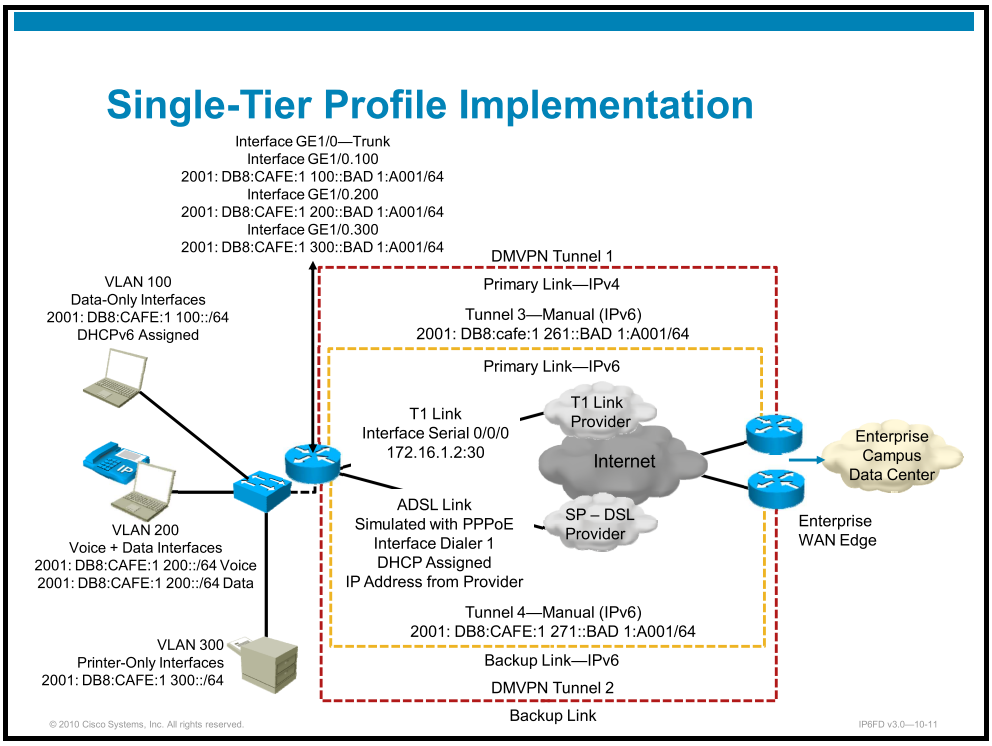

The network topology that is outlined in the figure serves as a reference for the sample single-tier profile branch network deployment.

A single router, 2900-br1-1 in this example, is used with an integrated switch, sw-br1-1, to provide WAN and LAN connectivity for the three VLANs in the branch.

The WAN consists of two connections: a T1 link to one ISP and an ADSL link to another ISP, which is used for redundancy. These tunnels are used for connectivity over the Internet to the headquarters site:

- Tunnel 1 is used as the primary DMVPN tunnel for IPv4-only traffic.
- Tunnel 2 is used as the backup DMVPN tunnel for IPv4-only traffic.
- Tunnel 3 is used as the primary manual tunnel (IPv6-in-IPv4) for IPv6-only traffic.
- Tunnel 4 is used as the backup manual tunnel for IPv6-only traffic.

All the tunnels use IPv4 IPsec for tunnel protection.

The LAN portion of this single-tier profile uses a Cisco EtherSwitch Service Module. Three VLANs are in use in this single-tier profile example:

- **VLAN 100:** This VLAN is used as the PC data VLAN. IPv4 addressing is provided by a local DHCP pool on the router. IPv6 addressing is provided by the branch router, using the prefix that is assigned to the router subinterface. DNS and domain name are provided by a local DHCP pool for IPv6.
- **VLAN 200:** This VLAN is used as the voice VLAN. IPv4 addressing is provided by a local DHCP pool on the router to include any voice-specific options like TFTP server. IPv6 addressing is provided by stateless autoconfiguration. IPv6 is enabled for planning purposes because no IPv6-enabled IP phones are in this design yet.
- **VLAN 300:** This VLAN is used as the printer VLAN. IPv4 addressing is provided by a local DHCP pool on the router. The Hewlett Packard Jet Direct cards in the branch automatically receive an IPv6 address from the router interface, via stateless autoconfiguration.

https://t.me/learningnets

Branch Deployment Profiles: Dual-Tier and Multitier Profile Implementations

This topic explains the dual-tier and multitier branch IPv6 deployment scenarios.

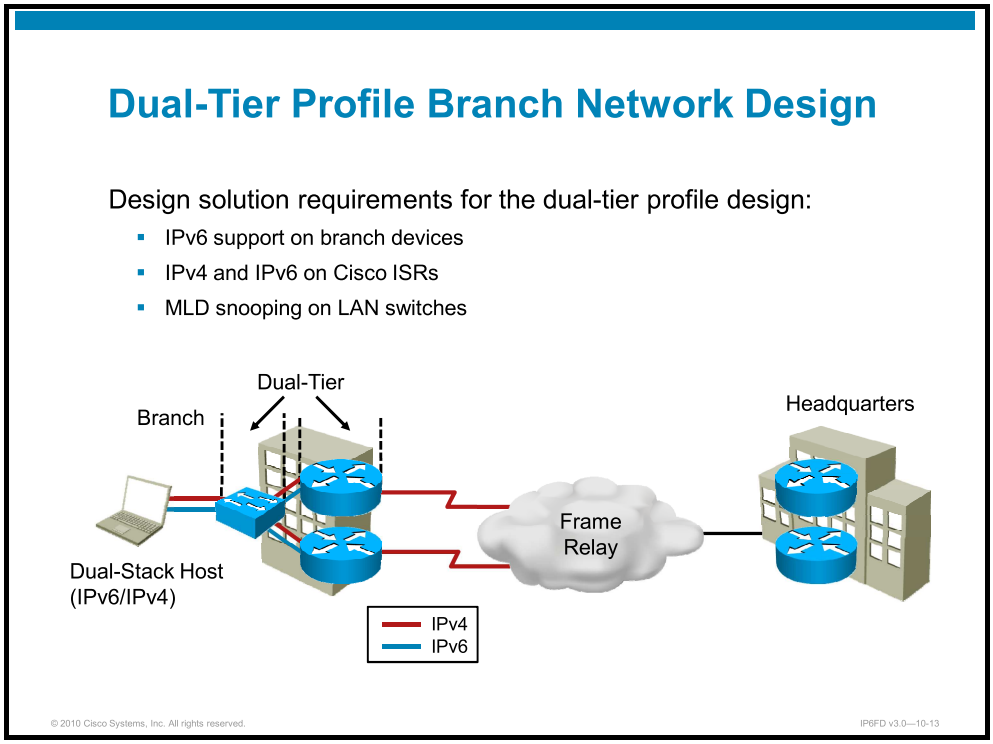

The dual-tier profile separates the routing and switching roles in the branch.

Three primary differences exist between the single-tier and dual-tier profile:

- Redundancy
- Scalability
- WAN transport

The dual-tier separates the LAN (switch) and WAN (router) components to offer fault tolerance. A single switch or multiple switches can be used to provide LAN access in the branch. Two WAN routers are redundantly connected to the Frame Relay cloud and to the LAN switch.

The dual-tier profile scales better than the single-tier profile does. In the single-tier profile, every network role that is required in the branch is performed by the ISR. This scenario is great for cost and manageability but can limit availability and scalability. The larger the branch and the more services that are enabled on the ISR, the higher the risk of overextending the performance capabilities of the ISR. This risk can be alleviated by using a more powerful ISR model, but doing so does not help with the fault-tolerance requirement. If additional LAN switches are needed at the branch, then the Catalyst switches can be used together by using the Cisco StackWise topology.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

The WAN connections in the described dual-tier model use Frame Relay instead of the Internet with IPsec VPN. IPv6 is fully supported over Frame Relay in Cisco IOS Software, so there is no need to run tunnels of any kind between the branch and headquarters. This is a great advantage for deployment and management because dual-stack is used all the way from the hosts in the branch LAN, across the WAN, and into the headquarters network. This approach greatly eases the operational aspects of deploying IPv6 in the branch because no special tunnel considerations (such as availability, security, QoS, and multicast) need to be made. The dual-tier profile uses what is often referred to as the private WAN model.

Security for the dual-tier profile is the same as for the single-tier profile, except that both routers in the dual-tier profile provide security services and no IPsec tunnels are used. Many branch deployments use the dual-tier profile.

These are the solution requirements for the dual-tier profile:

- IPv6 support on the operating system of the host machines in the branch
- IPv6 and IPv4 dual-stack support on the Cisco ISRs
- MLD snooping support on the LAN switches (required if using IPv6 multicast)

https://t.me/learningnets

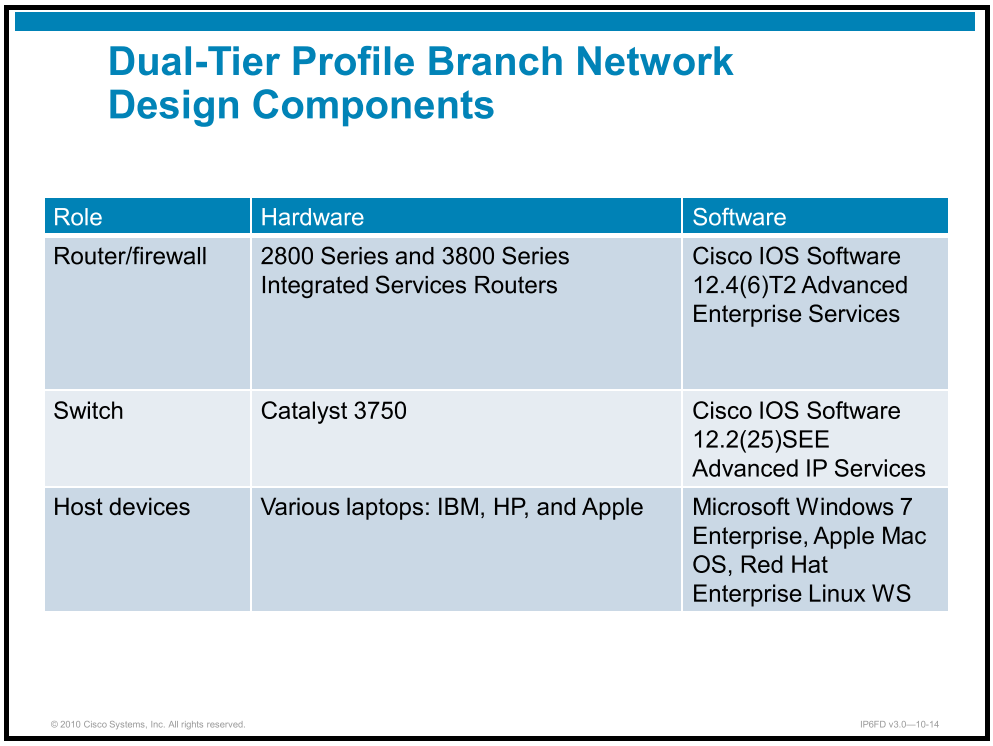

The table lists the dual-tier profile network design components:

Dual-Tier Profile Branch Network Design Components

Role	Hardware	Software
Router/firewall	2900 Series and 3800 Series Integrated Services Routers	Cisco IOS Software 12.4(6)T2 Advanced Enterprise Services
Switch	Cisco Catalyst 3750	Cisco IOS Software 12.2(25)SEE Advanced IP Services
Host devices	Various laptops: IBM, HP, and Apple	Microsoft Windows 7 Enterprise, Apple Mac OS, Red Hat Enterprise Linux WS

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

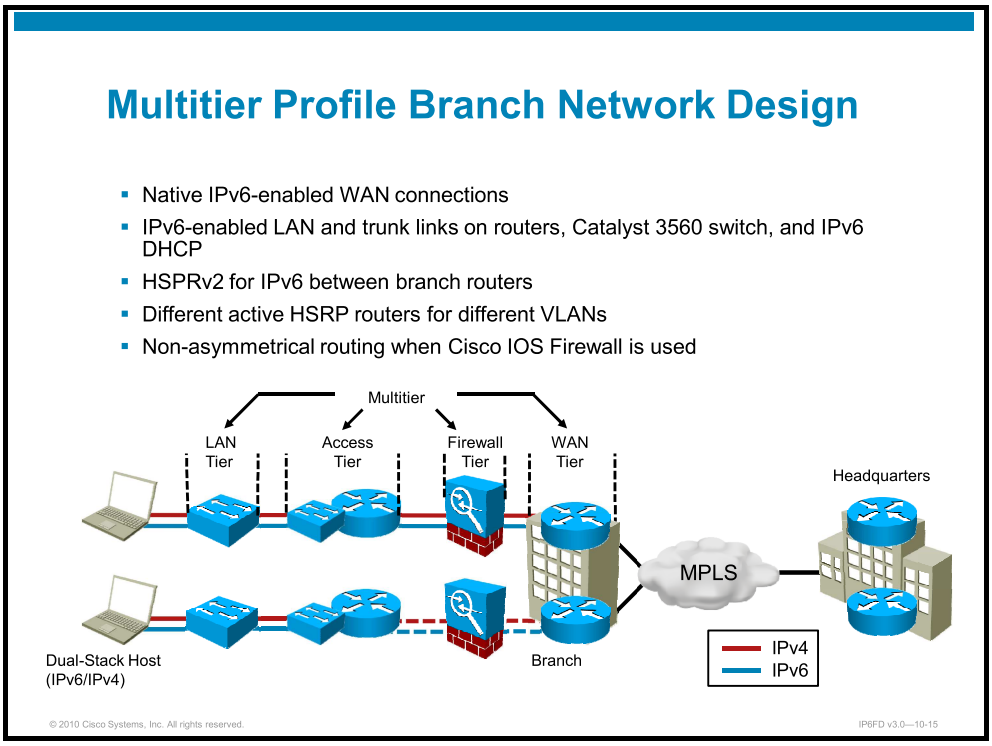

The multitier branch design profile consists of these tier components:

- LAN tier
- Access tier
- Firewall tier
- WAN tier

Connections to headquarters in the multitier profile occur over Multiprotocol Label Switching (MPLS).

Firewall services are separated from the WAN routers. Cisco ASA 5500 Series Adaptive Security Appliances provide stateful firewall services for both IPv4 and IPv6. The second Cisco ASA appliance is in stateful failover mode. In a stateful failover configuration, only one Cisco ASA appliance is active at a time.

The access tier is used for internal service and VLAN termination for the LAN tier. The access tier is like a campus distribution layer in many ways.

The LAN tier is the same as the dual-tier LAN switch. There are just more switches, to account for the larger-scale requirements that are probably found in a larger branch.

These solution requirements apply to the multitier profile:

- IPv6 support on the operating system of the host machines in the branch
- IPv6 and IPv4 dual-stack support on the Cisco ISRs
- MLD snooping support on the LAN switches (required if using IPv6 multicast)
- Cisco ASA Software Version 7.0 or later

https://t.me/learningnets

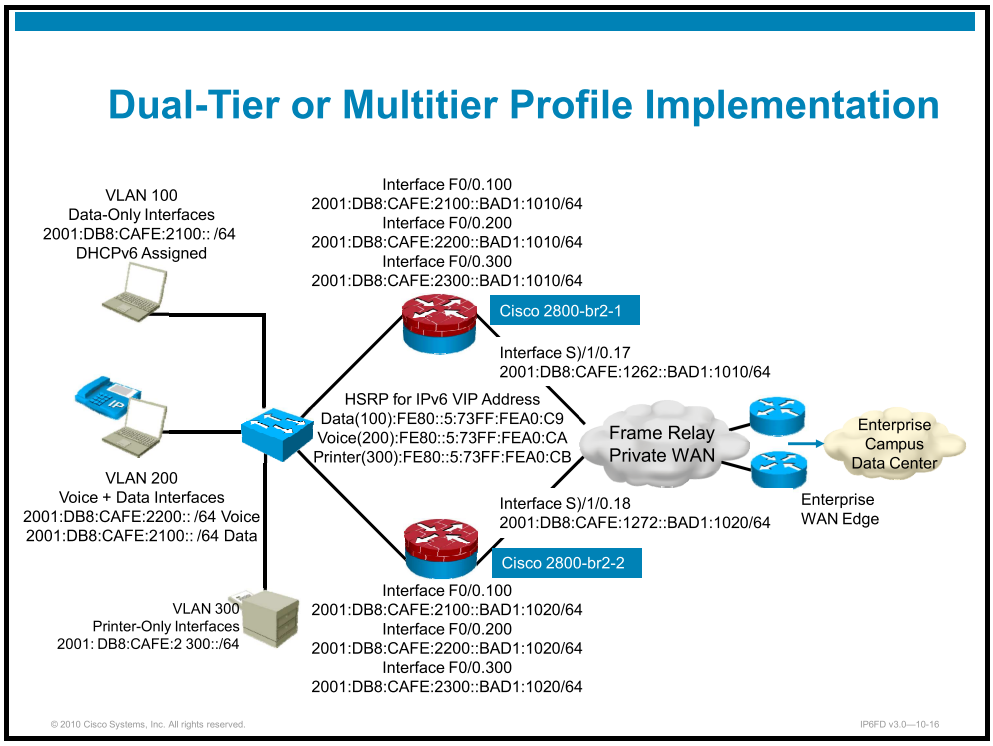

The network topology that is outlined in the figure serves as a reference for the sample dual-tier profile branch network deployment.

Two branch routers, 2900-br2-1 and 2900-br2-2, are used with a Catalyst 3560 switch, 3560-br2-1, to provide WAN and LAN connectivity for the three VLANs in the branch.

The WAN consists of one Frame Relay connection from each of the two branch routers. Unlike the single-tier profile, no tunnels are used in the dual-tier profile. The serial interfaces on the branch routers are enabled for dual-stack operation and have no dependency on IPv4 to pass IPv6 packets across the WAN. This scenario is preferred over running tunnels.

The LAN portion of the dual-tier uses a Catalyst 3560 switch. If additional switches are required for the branch, they can be trunked to the router or Cisco StackWise technology can be used. There are three VLANs in use in the dual-tier profile.

VLAN 100 is used as the PC data VLAN. IPv4 addressing is provided by a local DHCP pool on the router. IPv6 addressing is provided by the branch router, using the prefix that is assigned to the router subinterface. DNS and domain name are provided by a local DHCP pool for IPv6.

VLAN 200 is used as the voice VLAN. IPv4 addressing is provided by a local DHCP pool on the router to include any voice-specific options like TFTP server. IPv6 addressing is provided by stateless autoconfiguration. IPv6 is enabled for planning purposes; there are not yet any IPv6-enabled IP phones in this design.

VLAN 300 is used as the printer VLAN. IPv4 addressing is provided by a local DHCP pool on the router. The Hewlett Packard Jet Direct cards in the branch automatically receive an IPv6 address from the router interface, via stateless autoconfiguration.

 © 2010 Cisco Systems, Inc.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

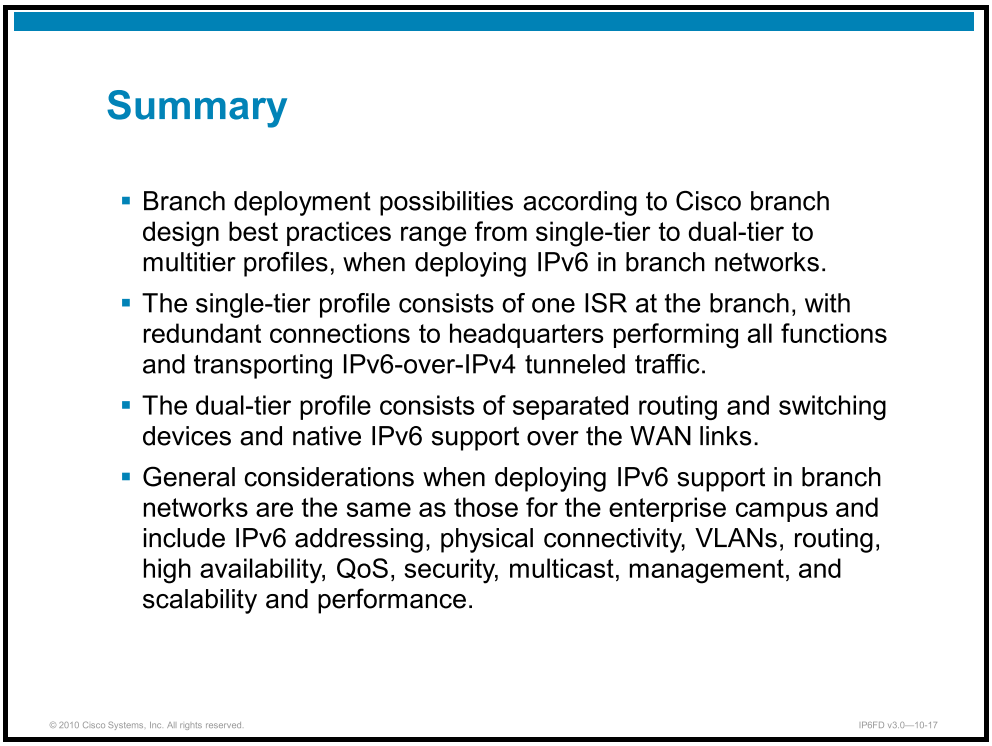

Resources

For additional information, refer to this resource:

- Cisco Systems, Inc. Cisco Validated Designs. http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns741/networking_solutions_products_genericcontent0900aecd80601e22.html

https://t.me/learningnets

© 2010 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

This topic summarizes the key points that were discussed in this module.

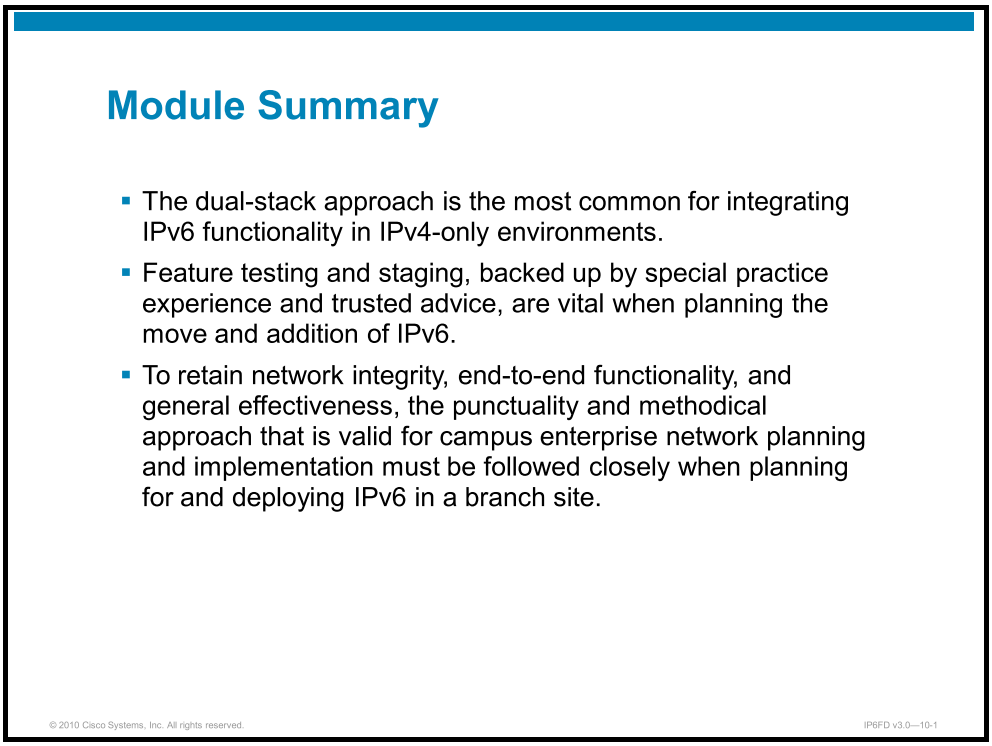

https://t.me/learningnets

© 2010 Cisco Systems, Inc.

https://t.me/learningnets