CISCO

CCNP Certification

Courseware

Version 1.0

www.firebrandtraining.com

IMPLEMENTING TELEWORKER SERVICES

Teleworker Connectivity

- Branch offices and teleworkers can connect via cable modems or DSL
- Both utilize PPP to provide authentication and accounting ability since neither is native to Ethernet/ATM
  - Cable modems
    - Governed by Data over Cable Service Interface Specification (DOCSIS)
    - Uses PPPoE
  - DSL
    - Uses PPPoA

IPSEC

IPSEC

- IPSEC
  - Supports IP traffic only
  - Traditionally supports unicast traffic only
  - Provides security by offering authentication and encryption
  - Can be used in the following VPN configuration
    - DMVPN (Dynamic Multipoint Virtual Private Network)
    - Virtual Tunnel interfaces (VTI)
    - GET (Group Encrypted Transport)

What is IPSec?

- A IP Security framework that includes multiple protocols and algorithms
- Provides for:
  - Authentication of every IP packet
  - Verification of data integrity for each packet
  - Confidentiality of packet payload
  - Anti-replay protection to verify each packet is unique

IPSec Components

- Security Protocols
  - Authentication Header (AH)
  - Encapsulating Security Payload (ESP)
- Key Management
  - ISAKMP, IKE, SKEME
- Security Algorithms
  - DES, 3DES, AES

IPSec Security Protocols

- Encapsulating Security Payload (ESP)
  - –IP Protocol **50**
  - –Provides **confidentiality**, integrity, origin authentication, and anti-replay
- Authentication Header (AH)
  - –IP Protocol **51**
  - –Provides integrity, origin authentication, and anti-replay
- Security protocols can operate in **tunnel mode** or **transport mode**

Transport Mode

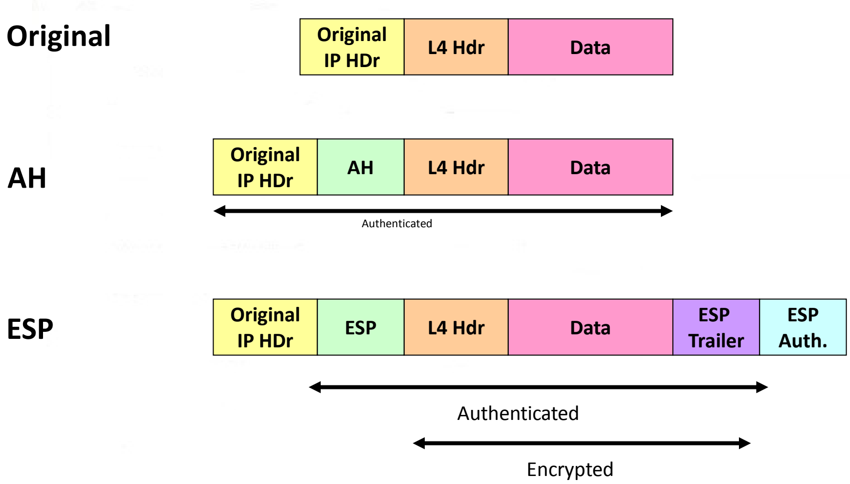

Tunnel Mode

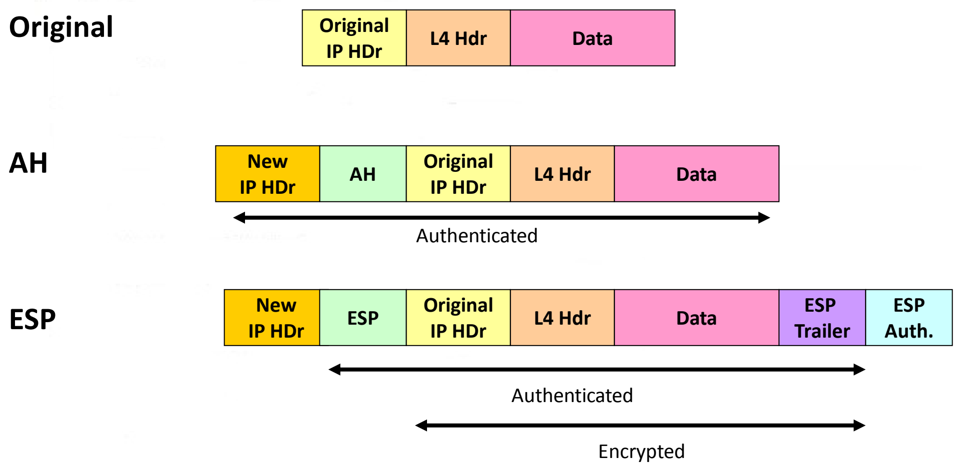

Implementing IPSec VPNs

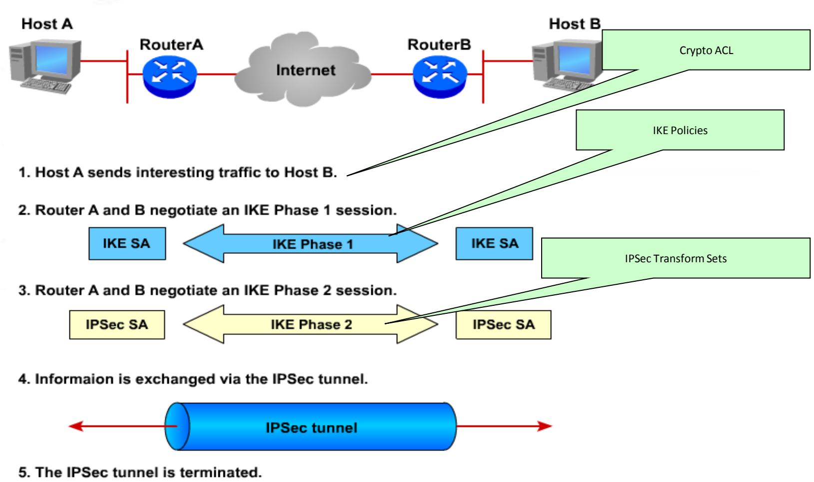

Implementing IPSec VPNs

1. Establish ISAKMP policy
2. Configure IPSec transform set
3. Configure crypto ACL
4. Configure crypto map
5. Apply crypto map to interface

1. Establish ISAKMP policy

Crypto isakmp policy 1 authentication pre-shared hash sha encryption aes 128 group 2 ! Crypto isakmp key SECRET address 172.16.0.2	Crypto isakmp policy 5 authentication pre-shared hash sha encryption aes 128 group 2 ! Crypto isakmp key SECRET address 172.16.0.1

2. Configure IPSec transform set

E0 10.1.0.0/16 A E1 172.16.0.1 E1 172.16.0.2 B E0 10.2.0.0/16

Crypto ipsec transform-set MY_VPN esp-aes 128 esp-sha-hmac	Crypto ipsec transform-set MY_VPN esp-aes 128 esp-sha-hmac

3. Configure crypto ACL

E0 10.1.0.0/16 A E1 172.16.0.1 E1 172.16.0.2 B E0 10.2.0.0/16

Access-list 101 permit ip 10.1.0.0 0.0.255.255 10.2.0.0 0.0.255.255	Access-list 101 permit ip 10.2.0.0 0.0.255.255 10.1.0.0 0.0.255.255

4. Configure crypto map


```
Crypto map VPN_TO_ROUTERB 10 ipsec-
  isakmp
  set peer 172.16.0.2
  match address 101
  set transform-set MY_VPN
```

```
Crypto map VPN_TO_ROUTERA 10 ipsec-
  isakmp
  set peer 172.16.0.1
  match address 101
  set transform-set MY_VPN
```

5. Apply crypto map to interface


```
Interface ethernet 1
 ip address 172.16.0.1 255.255.0.0
 crypto map VPN_TO_ROUTERB
!
Ip route 10.2.0.0 255.255.0.0
  172.16.0.2
```

```
Interface ethernet 1
 ip address 172.16.0.2 255.255.0.0
 crypto map VPN_TO_ROUTERA
!
Ip route 10.1.0.0 255.255.0.0
   172.16.0.1
```

GRE VPNS

GRE and GRE over IPSEC

WHY GRE?

- GRE can encapsulate and tunnel any protocol, IPSEC is limited to IP
- GRE can encapsulate unicast, multicast, and broadcast traffic, IPSEC is traditionally limited to just unicast
- Hence GRE can be used to tunnel **dynamic routing protocols**
- Major weakness within GRE: extremely limited security

GRE Tunnels

Original	**Original IP HDr**	**L4 Hdr**	**Data**		
GRE	**IP Hdr**	**GRE**	**Original IP HDr**	**L4 Hdr**	**Data**

Routing over GRE tunnel

Interface tunnel 0
 ip address 192.168.5.5 255.255.255.252
 tunnel source serial 0/0
 tunnel destination 172.17.0.1
 tunnel mode gre ip
!
Router eigrp 1
 network 192.168.5.4 0.0.0.3
 network 192.168.1.0 0.0.0.255
!
Ip route 0.0.0.0 0.0.0.0 172.16.0.254

Interface tunnel 1
 ip address 192.168.5.6 255.255.255.252
 tunnel source serial 0/0
 tunnel destination 172.16.0.1
 tunnel mode gre ip
!
Router eigrp 1
 network 192.168.5.4 0.0.0.3
 network 192.168.2.0 0.0.0.255
!
Ip route 0.0.0.0 0.0.0.0 172.17.0.254

Fa0/0
192.168.1.1

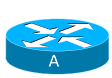

s0/0
172.16.0.1

s0/0
172.17.0.1

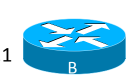

Fa0/0
192.168.2.1

Routing over GRE

Tunneling issues:

- Tunnel interface numbers do not have to match
- The tunnel source must be a local interface that is reachable over the WAN connection
- The tunnel destination must be an address of the remote router that is reachable over the WAN
- The tunnel interfaces must be on the same network to form neighborship

Static Route

Routing issues

- Need a route to connect over the WAN/internet
  - Ip route 0.0.0.0 0.0.0.0 <next-hop ip> ***!Not remote tunnel endpoint!***
- EIGRP neighbor commands are unnecessary because GRE will convert the EIGRP multicast traffic to unicast
- EIGRP autonomous system numbers must match
- EIGRP network statements must enable directly connected networks
  - LAN interface
  - Tunnel interface
  - **NOT** the WAN interface as it is connected to the internet

GRE Over IPSec

- Commonly used on Internet
- Emulates WAN to provide hub-and-spoke topology

Tunnel	**IP Hdr**	**ESP**	**IP Hdr**	**GRE**	**IP HDr**	**L4 Hdr**	**Data**	**ESP Trailer**
Transport		**IP Hdr**	**ESP**	**GRE**	**IP HDr**	**L4 Hdr**	**Data**	**ESP Trailer**

GRE over IPSEC

- GRE configuration requires
  - Tunnel source
  - Tunnel destination
  - Tunnel IP address
- IPSEC configuration requires a crypto map
  - The crypto ACL must reference GRE traffic instead of the LAN traffic

GRE over IPSEC configuration

Access-list 110 permit gre host 10.10.0.1 host 10.20.0.1

Crypto map **VPN** 10 ipsec-isakmp

set peer 10.20.0.1

set transform-set BRANCH_VPN

match address 110

Interface tunnel 0

ip address 192.168.0.1 255.255.255.0

tunnel source serial 0/0

tunnel destination 10.20.0.1

*crypto map **VPN***

interface serial 0/0

ip address 10.10.0.1 255.255.255.0

*crypto map **VPN***

- Crypto map name must match, case sensitive
- ACL must reference GRE traffic from one tunnel endpoint to the other and be referenced within the crypto map

Implement IPv6 Routing

IPv6 TOPICS

- Comparison with IPv4
- Addressing
- Address Assignment
- Routing
- Transition Methods

IPv4 versus IPv6

How IPv6 is better

IPv6 Enhancements

- No more broadcasts
- Built-in in support for anycast addresses
- IP mobility is built-in
- IPSEC is mandatory
  - Routing protocols no longer have any authentication methods as they rely on IPv6 IPSEC
- Number of Addresses
  - In IPv4 address depletion and routing table size are major concerns. In IPv6 these concerns are alleviated:
    - 128 bit totaling approximately 3.4x 10^38
    - Route summarization much more effective
    - Furthermore there is no need for NAT/PAT

IPv6 Header Enhancements

- Header length fixed at 40 **bytes,** IPv4 header 20 bytes
- No more L3 Checksum
  - Simpler and more efficient than IPv4 despite increased size
- Next Header field – specifies the next encapsulated protocol
- Flow label– to improve QOS and monitoring

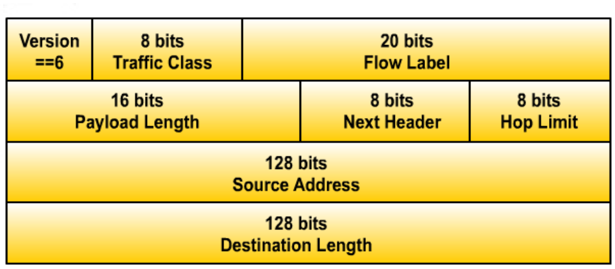

IPv6 Addressing

Rules and Examples

Addressing

- All addresses are 128 bits
- CIDR notation used to denote subnet mask
- Write as sequence of eight sets of four hex digits (16 bits each) separated by colons
- Can be written shorthand:
  – Lead zeros in a quartet may be omitted
  – Contiguous all-zero groups may be replaced by "::" but only one such group can be replaced

IPv6 Addressing Example

- ***3ffe:3700:0200:00ff:0000:0000:0000:0001***
  can be written:
- ***3ffe:3700:200:ff:0:0:0:1***
  or:
- ***3ffe:3700:200:ff::1***

IPv6 Address Types

- IPv6 defines three types of addresses or scopes:
  - **Unicast**
    - Global: public addresses
    - Link local: not routable; used for router and neighbor discovery
    - Unique local: equivalent of RFC 1918 addresses (site local addresses have been deprecated); uses FD00::/8
  - **Anycast**— Address specifies a set of hosts/servers for a given organization's application. A packet sent to an anycast address is delivered to one of the hosts identified by that address, usually the closest one as defined by the routing protocol.
    - All nodes should provide uniform service
    - Suitable for load-balancing and content-delivery services
    - **(config-if)#ipv6 address** <address> **anycast**
  - **Multicast**— Same concept as IPv4 multicast.

IPv6 Addressing

- Interfaces can have multiple addresses of any sort:
  - Unicast
  - Multicast
  - Anycast
- All interfaces have link local addresses (used by routing protocols)
  - By just enabling ipv6 on an interface a link local address will be automatically generated

    (config-if)#ipv6 enable

Global Unicast Addressing

- **Global Unicast**
  - Equivalent to IPv4 public address except there is no concept of a class in IPv6
  - Addresses start with 001 in binary(2000::/3)
- Classless routing and geographic assignment lessons learned from IPv6 are being deployed from the start
- ICANN owns addresses and along with IANA assigns them as follows
  - Registry -> /12
    - Registrars then hand addresses to Tier one ISP or subsidiary registrars
  - ISP Prefix -> /32
  - Customer Prefix -> /48
    - Known as global routing or site prefix
  - Subnet Prefix -> /64
- Remaining bits -> Interface (host) ID
- An example of a prefix would be 2000:1:2:3::/64
  - Note the prefix must end in :: to represent the host id with all zeros

Link-Local (unicast) Addresses

- **Link-Local Unicast**
  - No real equivalent in IPv4
  - Start with FE80::/10
  - Used by routing protocols, neighbor discovery, and router discovery
  - Also used to denote next-hop addresses within the IPv6 routing table
  - Can be automatically created using EUI-64 variants or manually specified
    - **(config-if)#ipv6 address** <address> **link-local**

Multicasting

- **Multicast**
  - FF00::/8
  - FF02::/16 link local addresses, for example
    - **FF02::1 all hosts**
    - **FF02::2 all routers in a local segment**
    - **FF02::5 ALL OSPF router**
    - **FF02::6 DR and BDR**
    - **FF02::9 RIPng**
    - **FF02::A EIGRP**
    - **FF02::1:2 unknown DHCP servers (dhcp relay agent function)**

Other Notable IPv6 Addresses

- ::/0 is the notation for a default route
  - **(config)#ipv6 route ::/0 s0/0/0**
- ::1/128 is the loopback address
  - Equivalent of 127.0.0.1
- ::/128 is the notation for an unspecified route or address

Address Assignment

IPv6 Address Assignment

- Hosts can be configured with an IP address via 3 methods
  - Static or manual addressing
  - DHCPv6
    - Stateful, roughly the same as DHCPv4
    - Does not assign default gateway addresses however
    - **(config-if)#ipv6 address dhcp**
  - Stateless Autoconfiguration
    - No equivalent in IPv4
    - No need for stateful DHCP
    - Uses EUI-64 to generate host address
    - Uses stateless DHCP to acquire DNS information

Static or Manual addressing

(config-if)#**ipv6 address** *address/prefix-length* **[eui-64]**

- EUI-64: automatically configure a host address

(config-if)#**ipv6 address 3FFE:2E::/64 eui-64**

OR

(config-if)#**ipv6 address 2001:2:3::1/64**

OR

(config-if)#**ipv6 enable**

EUI-64

- Used to automatically generate a unique host (interface) ID
- 48 bit MAC address is split and FF-FE is added in the middle to make up 64 bit host address
- 7^{th} bit of first octet is set to 1 to specify the address as unique (Group/Local bit)
- For example the MAC address 0**0**5C:99D0:CF34 becomes 0**2**5C:99**FF:FE**D0:CF34 for the host portion of an IPv6 address

Stateless Auto-configuration

- Part of Neighbor Discovery Protocol (ICMP)
- Multicast traffic including
  - Router Solicitation (RS) messages
    - Source address is link-local unicast, destination address is FF02::2
    - From host to all routers
    - Host sends message to acquire the IP address of its router
    - To configure cisco device as a host using stateless autoconfig
      - **(config)#ipv6 address autoconfig**
  - Router Advertisement (RA) messages
    - Source address is link-local unicast, destination address is FF02::1
    - From router to all hosts
    - The router advertises its address along with one or more prefix
    - **#show ipv6 router** shows cached RA messages

Stateless Autoconfiguration

- Routers send **Router Advertisement (RA)** messages that may include:
  - Whether nodes could use address auto-configuration
  - One or more on-link **IPv6 prefixes** that nodes on the local link could use
  - **Lifetime** information for each prefix
  - Whether the router sending the advertisement should be used as a default router
  - Additional information for hosts, such as the hop limit and MTU a host should use

ICMP: Neighbor Discovery

- Address Resolution Protocol was used to resolve IP address to MAC addresses
- ARP was broadcast based and no longer exists in IPv6
- Effectively replaced by Neighbor Discovery protocol
- **#Show ipv6 neighbors** shows NDP cache
- PCs send out Neighbor Solicitation (NS) messages to resolve MAC addresses and listen for Neighbor Advertisements (NA) as a response
- Uses solicited node multicast addresses
  - FF02::1:FF:0/104
    - Last 24 bits based on IPv6 address
  - MAC address: 01005e followed by the last 23 bits of IPv6 address

DAD and Inverse NDP

Duplicate Address Detection

- NS message with one's own solicited multicast address
- If response received then there is a duplicate address

Inverse Neighbor Discovery

- Replaces inverse ARP in Frame-relay networks
- Uses inverse NS and inverse NA messages

Routing

Static routes, dynamic routing protocols and redistribution

IPv6 Routing

- Static
- RIPng
- OSPFv3
- EIGRP for IPv6
- MP-BGP4
- IS-IS for IPv6

IPv6 Routing enabled with
(config)#**ipv6 unicast-routing**

Routing Protocols

- No more network commands, all protocols enabled on a per-interface level
- No native authentication methods, all protocols use IPv6 ESP/AH (IPv6 IPSEC)
- Neighbors do not have to be on the same subnet in EIGRP and OSPFv3
- Routing table shows Local (L) routes with /128 masks to represent host addresses
  - **#show ipv6 route**
- Redistribution
  - Host routes (L for local) are not redistributed
  - Seed metric for RIPng=source IGP metric
  - No **subnets** keyword for OSPFv3 as there is no classful concepts in IPv6
  - Directly connected networks are not automatically redistributed
    - Redistribute command must have the keyword **include-connected**

Static Routing

- **(config)#ipv6 route** <prefix/mask> <next-hop IP address>
- **(config)#ipv6 route** ::/0 2000:1:2::1
- Can use any valid next-hop IP address
- If link-local address is used for the next-hop address then you must configure the exit interface and link-local address
  - **(config)#ipv6 route** 2003:12::/64 s0/0/0 FE80:1::1

RIPng

- Essentially the same as RIPv2
- Globally enable RIPng process
  - **(config)#ipv6 router rip** <name> **enable**
- Enable individual interfaces to run RIPng
  - **(config-if)#ipv6 rip** <name>
- Link local and host routes not advertised
- Seed metric into RIPng for redistribution is based on the source IGP metric

Differences between OSPF v2 and v3

- Multiple OSPF instances can run over a single link
  - Cannot select specific interface addresses via ACLs or any other method into a given OSPF process.
- Uses link-local addresses to find adjacent neighbors
- 224.0.0.5 is now **FF02::5**
- 224.0.0.6 is now **FF02::6**

Differences between OSPF v2 and v3 (cont.)

- New LSA types
  - Link LSAs (type 8): link-local flooding
  - **Intra-area prefix (type 9):** generated by ABR and sent to backbone

OSPFv3 Configuration

1. Enable IPv6 routing

(config)#**ipv6 unicast-routing**

2. Enter Router Configuration mode

(config)#**ipv6 router ospf 1**

3. Assign 32-bit Router ID

(config-router)#**router-id 1.1.1.1**

OSPFv3 Configuration

4. Enable OSPF process on a per interface basis

(config)#**interface ethernet 0**

(config-if)#**ipv6 ospf 1 area 0**

(config)#**interface serial 0**

(config-if)#**ipv6 ospf 1 area 1**

OSPFv3 Configuration

5. (optional) Configure parameters on the interface

(config-if)#**ipv6 ospf priority 255**

(config-if)#**ipv6 ospf cost 20**

6. (optional) Configure summarization

(config)#**ipv6 router ospf 1**

(config-router)#**area 1 range 2001:0DB8::/48**

OSPFv3 Example

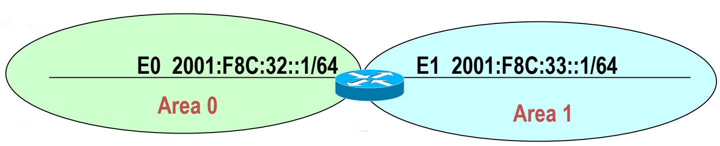


```
ipv6 unicast-routing
!
interface ethernet 0
    ipv6 address 2001:F8C:32::1/64
    ipv6 ospf 1 area 0
    ipv6 ospf 1 priority 0
!
```

```
interface ethernet 1
    ipv6 address 2001:F8C:33::1/64
    ipv6 ospf 1 area 1
!
ipv6 router ospf 1
    router-id 1.1.1.1
```

Verifying OSPFv3

```
Router#show ipv6 ospf interface
FastEthernet0/0 is up, line protocol is up
Link Local Address FE80::205:5FFF:FED3:5808, Interface ID 3
Area 1, Process ID 1, Instance ID 0, Router ID 172.16.3.3
Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State BDR, Priority 1
Designated Router (ID) 172.16.6.6, local address FE80::205:5FFF:FED3:6408
Backup Designated router (ID) 172.16.3.3, local address FE80::205:5FFF:FED3:5808
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:05
Index 1/1/1, flood queue length 0
Next 0x0(0)/0x0(0)/0x0(0)
Last flood scan length is 12, maximum is 12
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 1, Adjacent neighbor count is 1
Adjacent with neighbor 172.16.6.6 (Designated Router)
Suppress hello for 0 neighbor(s)
```

Troubleshooting OSPFv3

- **Clear ipv6 ospf process**
  - Re-establishes adjacencies
  - repopulates the OSPF database
  - runs the shortest path first (SPF) algorithm

EIGRP

- Also uses router-id
- Also does not advertise link local or host routes
- Neighbors do not have to be on the same subnet
- Shutdown by default, need to issue the **no shutdown** command
  - This command exists for OSPFv3 as well but OSPFv3 is enabled by default

Transition Methods

Coexistence with IPv4 networks

Dual-stack, Tunneling, and NAT-PT

Transition Approaches

1. Dual Stack
   -systems configured with IPv4 and IPv6 addresses
   -IPv4 and IPv6 routing protocols can be run simultaneously
2. NAT-PT
   - Translate the entire IPv4 header to IPv6 and vice versa
   - Can utilize DNS as application layer gateway
   - Deprecated
3. Tunneling
   - IPv6 packets encapsulated within IPv4
   - Configured between dual-stack routers or hosts
   - 5 types of tunnels
     - Manually Configured Tunnels (MCT)
     - GRE
     - 6to4 tunnels
     - ISATAP
     - IPv4 compatible tunnels (deprecated)

Tunneling

- Less configuration of IPv6 as compared to dual-stack but more overhead due to tunneling.
- Two categories of tunneling
  - Point to point (static)
    - GRE
    - Manually Configured Tunnels (MCT)
  - Multipoint (automatic)
    - 6to4
    - ISATAP

Tunneling

- Point to point
  - Configure tunnel source and destination
  - Support dynamic routing protocols
  - Good if there is frequent usage as there is less work per packet
- Multipoint
  - Configure only tunnel source
  - Dynamically learns tunnel destination based on destination IPv6 address or IPv6 next-hop address
  - More addressing rules
  - IPv4 address embedded into IPv6 address
  - Do not support routing protocols, therefore need static routes

Tunneling

- 6to4 tunnels
  - IPv4 protocol 41
  - Each site receives a /48 prefix comprised of
    - 2002::/**16** (address range specifically assigned to 6to4)
    - Followed by IPv4 address
      - » Use the IPv4 address specified as the tunnel source
  - Automatic, allows for multiple destinations
  - Used for the Internet
- Manual tunneling
  - IPv4 protocol 41

Configuring IPv6 Tunnels: Manual

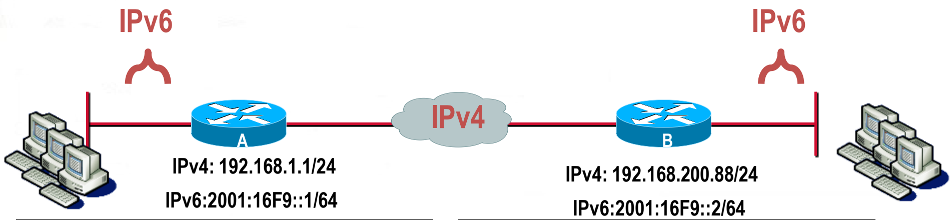


```
interface tunnel 0
  ipv6 address 2001:16F9::1/64
  tunnel source 192.168.1.1
  tunnel destination 192.168.200.88
  tunnel mode ipv6ip
```

```
interface tunnel 0
  ipv6 address 2001:16F9::2/64
  tunnel source 192.168.200.88
  tunnel destination 192.168.1.1
  tunnel mode ipv6ip
```

Configuring IPv6 Tunnels: GRE

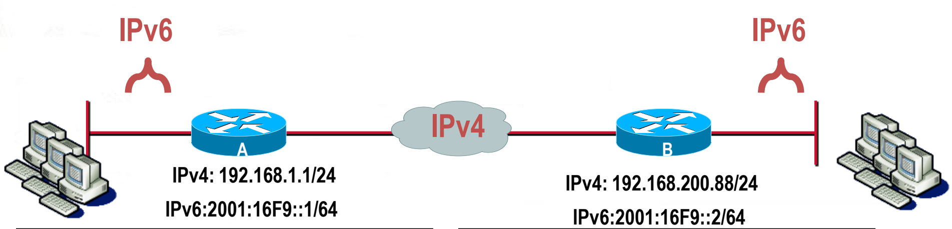


```
interface tunnel 0
  ipv6 address 2001:16F9::1/64
  tunnel source 192.168.1.1
  tunnel destination 192.168.200.88
```

```
interface tunnel 0
  ipv6 address 2001:16F9::2/64
  tunnel source 192.168.200.88
  tunnel destination 192.168.1.1
  tunnel mode gre ip
```

Configuring 6to4 Tunnels

- Point-to-multipoint model (**multiple destinations**)
- Destination is determined by extracting IPv4 address from IPv6 address
  - IPv4 address is converted to hex
  - Start with **2002::/16**
  - **/48 bit prefix with 2nd and 3rd quartet derived from IPv4 address of tunnel source**

192.168.99.1 = C0A8:6301

```
interface tunnel 0
    ipv6 address 2002:C0A8:6301::1/64
    tunnel source ethernet 0
    tunnel mode ipv6ip 6to4
interface ethernet 0
    ip address 192.168.99.1 255.255.255.0
```