NETWORKERS 2004

IPv6

SESSION RST-1305

IPv6 NEEDS AND APPLICATIONS

IP: The Application's Convergence Layer

Cisco.com

With Millions of New Devices Becoming IP Aware, The Need for Increased Addressing and Plug and Play Networking Is Only Met with the Implementation of IPv6

RST-1305
9719_05_2004_c2

A Need for IPv6?

Cisco.com

- **IETF IPv6 WG began in early 90s, to solve addressing growth issues, but**
  - **CIDR, NAT,…were developed**
- **IPv4 32 bit address = 4 billion hosts**
  - **~40% of the IPv4 address space is still unused which is different from unallocated**
  - **BUT**
- **IP is everywhere**
  - **Data, voice, audio and video integration is a reality**
  - **Regional registries apply a strict allocation control**
- **So, only compelling reason: More IP addresses!**

RST-1305
9719_05_2004_c2

A Need for IPv6?

Cisco.com

- **Internet population**
  - ~600M users in Q4 CY ’02, ~945M by end CY ’04—only 10–15% of the total population
  - How to address the future Worldwide population? (~9B in CY ’50)
  - Emerging Internet countries need address space, e.g.: China uses nearly 2 class A (11/2002), ~20 class A needed if every student (320M) has to get an IP address
- **Mobile internet introduces new generation of Internet devices**
  - PDA (~20M in 2004), Mobile Phones (~1.5B in 2003), Tablet PC
  - Enable through several technologies, e.g.: 3G, 802.11,…
- **Transportation—mobile networks**
  - 1B automobiles forecast for 2008—Begin now on vertical markets
  - Internet access on planes, e.g. Lufthansa— train, e.g. Narita express
- **Consumer, home and industrial appliances**

IP Address Allocation History

Cisco.com

- **1981—IPv4 protocol published**
- **1985—1/16 of total space**
- **1990—1/8 of total space**
- **1995—1/3 of total space**
- **2000—1/2 of total space**
- **2002.5—2/3 of total space**

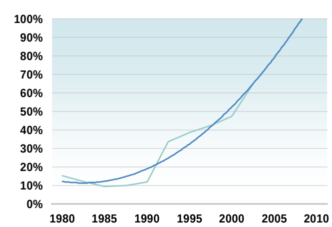

- **This despite increasingly intense conservation efforts**
  - PPP/DHCP address sharing
  - NAT (network address translation)
  - CIDR (classless interdomain routing) plus some address reclamation
- **Theoretical limit of 32-bit space: ~4 billion devices**
  practical limit of 32-bit space: ~250 million devices (RFC 3194)

Why Not NAT

Cisco.com

- Exhaustion of address space
- NAT breaks the end to end model
- Growth of NAT has slowed down growth of transparent applications
- No easy way to maintain states of NAT in case of node failures
- NAT break security
- NAT complicates mergers, double NATing is needed for devices to communicate with each other

IPv6 TECHNOLOGY

IPv6 Protocol

Cisco.com

Changes in Some Key Areas

- Simplification of header format
- Expanded address space
- Improved option support
- Mandated security

IPv6 Protocol

Cisco.com

Headers and Fields

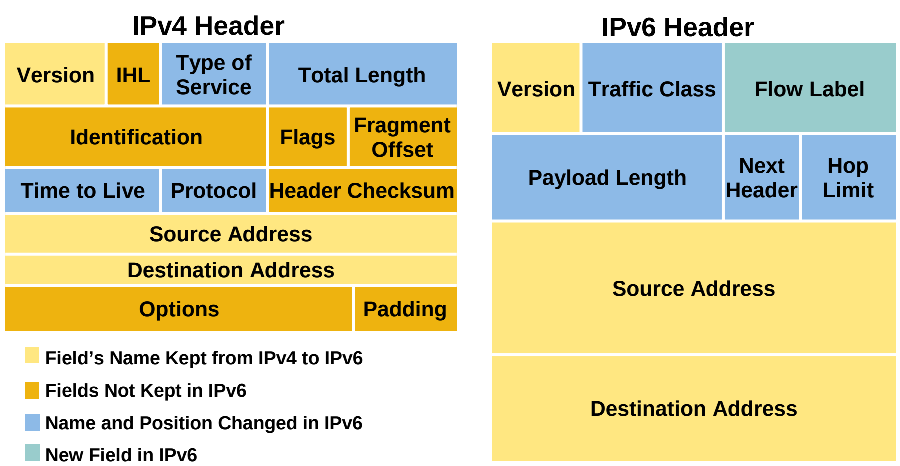

IPv6 Protocol

Cisco.com

- New field
- Flow label (RFC3697)

    Sequence of packets for which a source desires to label a flow

    Flow classifiers have been based on 5-tuple: source/destination address, protocol type and port numbers of transport

IPv6 Protocol

Cisco.com

- Flow label

    Some of these fields may be unavailable due to fragmentation, encryption or locating them past extension headers

    Looking for classifier only into IP header

    Only 3 tuple, flow label, source/destination address

ADDRESSING

RST-1305
9719_05_2004_c2

Addressing

Cisco.com

- Almost unlimited number of IP addresses
- The availability of these many addresses provides perfect platform for residential IP telephony

RST-1305
9719_05_2004_c2

Addressing

Cisco.com

Three Types of Address

1. Unicast
2. Multicast
3. Anycast

Addressing

Cisco.com

Representation

- 16-bit hexadecimal numbers
- Numbers are separated by (:)
- Hex numbers are not case sensitive

Let’s Talk a Little More on Anycast

Cisco.com

- Anycast allows a source node to transmit IP datargrams to a single destination node out of a group destination nodes with same subnet id based on the routing metrics

Anycast

Cisco.com

IPv4 Anycast Motivation and Issues

- It provided nodes a simpler way to reach any of groups application servers
- Anycasting did cause problems with stateful interactions, it requires mechanism that guides all anycast packets to the first node that responds to the request
- All anycast nodes should provide uniform service
- Suitable for load balancing and content delivery services

Addressing

Representation

- Abbreviations are possible
- Leading zeros in contiguous block could be represented by (::)

  Example:

  2003:0000:130F:0000:0000:087C:876B:140B

  2003:0:130F::87C:876B:140B
- Double colon only appears once in the address

Addressing

Prefix Representation

- Representation of prefix is just like CIDR
- In this representation you attach the prefix length
- Like v4 address 198.10.0.0/16
- v6 address is represented the same way 3ef8:ca62:12::/40

Addressing

Cisco.com

Some Special Addresses

Type	Binary	Hex
Aggregatable Global Unicast Address	0010	2
Link Local Unicast Address	1111 1110 10	FE80::/10
Unique local unicast address	1111 1100 1111 1101	FC00::/8 FD00::/8
Multicast address	1111 1111	FF00::/16

RST-1305
9719_05_2004_c2

Aggregatable Global Unicast Addresses

Cisco.com

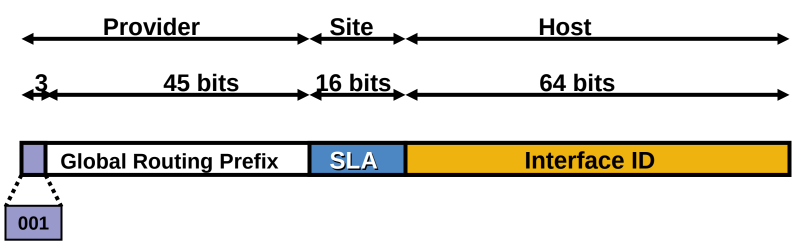

- **Aggregatable Global Unicast addresses are:**

  Addresses for generic use of IPv6

  Structured as a hierarchy to keep the aggregation

RST-1305
9719_05_2004_c2

Aggregatable Global Unicast Addresses

Cisco.com

Lowest-Order 64-Bit Field of Unicast Address May Be Assigned in Several Different Ways:

- Auto-configured from a 64-bit EUI-64, or expanded from a 48-bit MAC address (e.g., Ethernet address)
- Auto-generated pseudo-random number (to address privacy concerns)
- Assigned via DHCP
- Manually configured

RST-1305
9719_05_2004_c2 23

Aggregatable Global Unicast Addresses

Cisco.com

- Cisco uses the EUI-64 format to do stateless autoconfiguration
- This format expands the 48 bit MAC address to 64 bits by inserting FFFE into the middle 16 bits
- To make sure that the chosen address is from a unique Ethernet MAC address, the universal/local ("u" bit) is set to 1 for global scope and 0 for local scope

RST-1305
9719_05_2004_c2 24

Aggregatable Global Unicast Addresses

EUI-64

- Eui-64 address: Insert “FFE” in middle
- Invert ‘U’ bit to identify uniqueness of MAC
- Ethernet MAC address (48 bits)
- 64 bits version
- Uniqueness of the MAC
- Eui-64 address

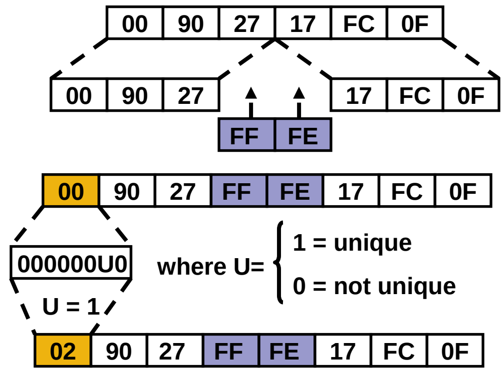

IPv6 Multicast Address

- IP multicast address has a prefix FF00::/8 (1111 1111)
- The second octet define the lifetime and scope of the multicast address

8-bit	4-bit	4-bit	112-bit
1111 1111	Lifetime	Scope	Group-ID

Lifetime	
0	if permanent
1	if temporary

Scope	
1	node
2	link
5	site
8	organization
E	global

Solicited-Node Multicast Address

Cisco.com

- For each unicast and anycast address configured there is a corresponding solicited-node multicast
- This address is link local significance only
- This is specially used for two purpose, for the replacement of ARP, and DAD

Solicited-Node Multicast Address

Cisco.com

- FF02:0000:0000:0000:0000:0001:FF00:0000/104
- FF02::1:FF00:0000/104
- Gets the lower 24 bits from the unicast address

NEIGHBOR DISCOVERY

Neighbor Discovery

- Replaces ARP, ICMP (redirects, router discovery)
- Reachability of neighbors
- Hosts use it to discover routers, autoconfiguration of addresses
- Duplicate Address Detection (DAD)

Neighbor Discovery

- Neighbor discovery uses icmpv6 messages, originated from node on link local with hop limit of 255
- Consists of ipv6 header, icmpv6 header, neighbor discovery header, and neighbor discovery options
- Five neighbor discovery messages
  - Router solicitation (icmpv6 type 133)
  - Router advertisement (icmpv6 type 134)
  - Neighbor solicitation (icmpv6 type 135)
  - Neighbor advertisement (icmpv6 type 136)
  - Redirect (ICMPV6 type 137)

Neighbor Discovery

Router Solicitation

- Host send to inquire about presence of a router on the link
- Send to all routers multicast address of FF02::2 (all routers multicast address)
- Source IP address is either link local address or unspecified IPv6 address (::)

Router Solicitation and Advertisement

1 - ICMP Type = 133 (RS)	2 - ICMP Type = 134 (RA)
Src = Link-local Address (FE80::/10)	Src = Link-local Address (FE80::/10)
Dst = All-routers multicast Address (FF02::2)	Dst = All-nodes multicast address (FF02::1)
Query= please send RA	Data= options, subnet prefix, lifetime, autoconfig flag

- Router solicitations (RS) are sent by booting nodes to request RAs for configuring the interfaces

Neighbor Solicitation

- Send to discover link layer address of IPv6 node
- For Layer 2 it is set to multicast for address resolution, unicast for node reachability
- IPv6 header, source address is set to unicast address of sending node, or :: for DAD (more later)
- Destination address is set to the unicast address for reachability and solicited node multicast for DAD

Neighbor Advertisement

Cisco.com

- Response to neighbor solicitation message
- Also send to inform change of link layer address

Neighbor Solicitation and Advertisement

Cisco.com

Neighbor Solicitation:
ICMP type = 135
Src = A
Dst = Solicited-node multicast Address
Data = link-layer address of A
Query = what is your link-layer address?

Neighbor Advertisement:
ICMP type = 136
Src = B
Dst = A
Data = link-layer address of B

A and B Can Now Exchange Packets on This Link

Redirect

Cisco.com

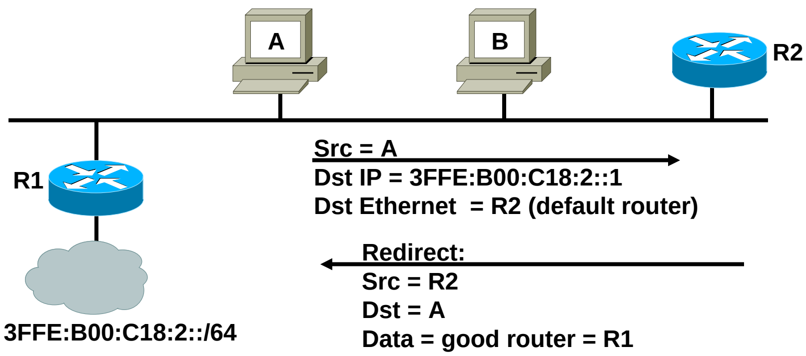

- **Redirect is used by a router to signal the reroute of a packet to a better router**

ROUTING

RIPNG (RFC 2080)

RST-1305
9719_05_2004_c2

Enhanced Routing Protocol Support
RIPng Overview

Cisco.com

- RIPng for IPv6, RFC 2080
- Same as IPv4:
  - Distance-vector, radius of 15 hops, split-horizon, etc.
  - Based on RIPv2
- Updated features for IPv6
  - IPv6 prefix, next-hop IPv6 address
  - Uses the multicast group FF02::9, the all-rip-routers multicast group, as the destination address for RIP updates
  - Uses IPv6 for transport

RST-1305
9719_05_2004_c2

Enhanced Routing Protocol Support RIPng Configuration and Display

Cisco.com

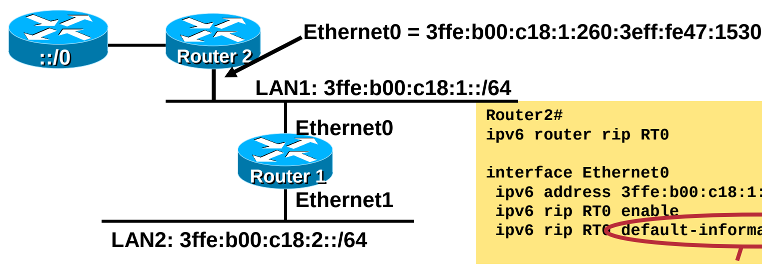


```
Router2#
ipv6 router rip RT0

interface Ethernet0
 ipv6 address 3ffe:b00:c18:1::/64 eui-64
 ipv6 rip RT0 enable
 ipv6 rip RT0 default-information originate
```

```
Router1#
ipv6 router rip RT0

interface Ethernet0
 ipv6 address 3ffe:b00:c18:1::/64 eui-64
 ipv6 rip RT0 enable
 Interface Ethernet1
 ipv6 address 3ffe:b00:c18:2::/64 eui-64
 ipv6 rip RT0 enable
```

```
Router2# debug ipv6 rip
RIPng: Sending multicast update on Ethernet0 for RT0
       src=FE80::260:3eff:fe47:1530
       dst=FF02::9 (Ethernet0)
       sport=521, dport=521, length=32
       command=2, version=1, mbz=0, #rte=1
       tag=0, metric=1, prefix=::/0
```

Multicast All Rip-Routers

Link-local src address

I/IS-IS FOR IPv6

Enhanced Routing Protocol Support
Integrated IS-IS for IPv6 Overview

Cisco.com

- **2 tag/length/values added to introduce IPv6 routing**
- **IPv6 reachability TLV (0xEC)**

  Describes network reachability such as IPv6 routing prefix, metric information and some option bits; the option bits indicates the advertisement of IPv6 prefix from a higher level, redistribution from other routing protocols

  Equivalent to IP internal/external reachability TLV’s described in RFC1195
- **IPv6 interface address TLV (0xE8)**

  Contains 128 bit address

  For Hello PDUs, must contain the link-local address (FE80::/10)

  For LSP, must only contain the non link-local address
- **A new Network Layer Protocol Identifier (NLPID) is defined**

  Allowing IS-IS routers with IPv6 support to advertise IPv6 prefix payload using 0x8E value (IPv4 and OSI uses different values)

Enhanced Routing Protocol Support
I/IS-IS for IPv6-Only Configuration Example

Cisco.com

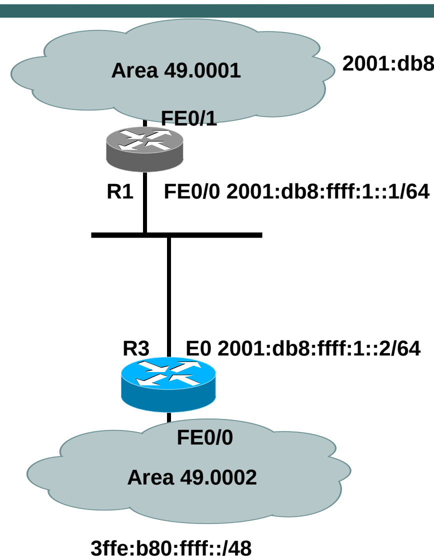


```
Router1#
 interface fastethernet0/0
     ipv6 address 2001:db8:ffff:1::1/64
     ipv6 router isis
     isis circuit-type level-2-only


router isis
      address-family ipv6
      redistribute static
      exit-address-family
      net 49.0001.1921.6801.0001.00
```

Enhanced Routing Protocol Support
Cisco IOS I/IS-IS Dual IP Configuration

Cisco.com

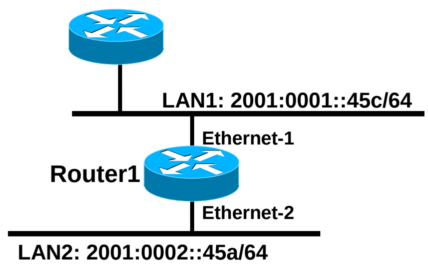

Dual IPv4/IPv6 Configuration Redistributing Both IPv6 Static Routes and IPv4 Static Routes

```
Router1#
 interface ethernet-1
     ip address 10.1.1.1 255.255.255.0
     ipv6 address 2001:0001::45c/64
     ip router isis
     ipv6 router isis

  interface ethernet-2
     ip address 10.2.1.1 255.255.255.0
     ipv6 address 2001:0002::45a/64
     ip router isis
     ipv6 router isis

  router isis
      address-family ipv6
      redistribute static
      exit-address-family
      net 49.0001.1921.6801.0001.00
      redistribute static
```


Enhanced Routing Protocol Support
Cisco IOS I/IS-IS Display

Cisco.com

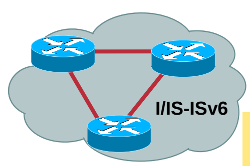


```
brum-45c#sho ipv6 rou is-is
IPv6 Routing Table - 14 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
Timers: Uptime/Expires

I1  2001:45A:1000::/64 [115/20]
     via FE80::210:7BFF:FEC2:ACCC, Ethernet1, 00:10:12/never
I1  2001:72B:2000::/64 [115/10]
     via FE80::210:7BFF:FEC2:ACCC, Ethernet1, 00:05:19/never
I1  2002:49::/64 [115/10]
     via FE80::210:7BFF:FEC2:ACCC, Ethernet1, 00:05:19/never
```

OSPFv3 (RFC 2740)

Enhanced Routing Protocol Support
Similarities with OSPFv2

Cisco.com

- OSPFv3 is OSPF for IPv6 (RFC 2740)
- Based on OSPFv2, with enhancements
- Distributes IPv6 prefixes
- Runs directly over IPv6
- OSPFv3 and v2 can be run concurrently, because each address family has a separate SPF (ships in the night)
- OSPFv3 uses the same basic packet types as OSPFv2 such as hello, Database Description Blocks (DDB), Link State Request (LSR), Link State Update (LSU) and Link State Advertisements (LSA)
- Neighbor discovery and adjacency formation mechanism are identical
- RFC compliant NBMA and point to multipoint topology modes are supported; also supports other modes from Cisco such as point to point and broadcast including the interface
- LSA flooding and aging mechanisms are identical

Enhanced Routing Protocol Support Differences from OSPFv2

Cisco.com

- **OSPF packet type**
- **Ospfv3 will have the same 5 packet type but some fields have been changed**

Packet Type	Description
1	Hello
2	Database Description
3	Link State Request
4	Link State Update
5	Link State Acknowlegdment

- **All OSPFv3 packets have a 16-byte header vs. the 24-byte header in OSPFv2**

OSPFv2 header:

Version	Type	Packet Length	
Router ID			
Area ID			
Checksum		Autotype	
Authentication			
Authentication			

OSPFv3 header:

Version	Type	Packet Length	
Router ID			
Area ID			
Checksum		Instance ID	0

Enhanced Routing Protocol Support Differences from OSPFv2

Cisco.com

- **OSPFv3 protocol processing per-link, not per-subnet**
  - IPv6 connects interfaces to links
  - Multiple IP subnets can be assigned to a single link
  - Two nodes can talk directly over a single even they do not share and common subnet
  - The term “network” and “subnet” is being replaced with “link”
  - An OSPF interface now connects to a link instead of a subnet
- **Multiple OSPFv3 protocol instances can now run over a single link**
  - This allows for separate ASes, each running OSPF, to use a common link. Single link could belong to multiple areas
  - Instance ID is a new field that is used to have multiple OSPFv3 protocol instance per link
  - In order to have 2 instances talk to each other they need to have the same instance ID; by default it is 0 and for any additional instance it is increased

Enhanced Routing Protocol Support Differences from OSPFv2

Cisco.com

- **Uses link local addresses**
  - To identify the OSPFv3 adjacency neighbors
- **Two new LSA types**
  - Link-LSA (LSA Type 0x2008)
  - There is one Link-LSA per link; this LSA advertises the router’s link-local address, list of all IPv6 prefixes and options associated with the link to all other routers attached to the link
  - Intra-Area-Prefix-LSA (LSA Type 0x2009)
  - Carries all IPv6 prefix information that in IPv4 is included in router LSAs and network LSAs
- **Two LSAs are renamed**
  - Type-3 summary-LSAs, renamed to “Inter-Area-Prefix-LSAs”
  - Type-4 summary LSAs, renamed to “Inter-Area-Router-LSAs”

RST-1305
9719_05_2004_c2

Enhanced Routing Protocol Support Differences from OSPFv2

Cisco.com

- **Multicast addresses**
  - FF02::5—Represents all SPF routers on the link local scope, equivalent to 224.0.0.5 in OSPFv2
  - FF02::6—Represents all DR routers on the link local scope, equivalent to 224.0.0.6 in OSPFv2
- **Removal of address semantics**
  - IPv6 addresses are no longer present in OSPF packet header (part of payload information)
  - Router LSA, network LSA do not carry IPv6 addresses
  - Router ID, Area ID and Link State ID remains at 32 bits
  - DR and BDR are now identified by their Router ID and no longer by their IP address
- **Security**
  - OSPFv3 uses IPv6 AH and ESP extension headers instead of variety of mechanisms defined in OSPFv2

RST-1305
9719_05_2004_c2

LSA Types

Cisco.com

	LSA Function Code	LSA Type
Router-LSA	1	0x2001
Network-LSA	2	0x2002
Inter-Area-Prefix-LSA	3	0x2003
Inter-Area-Router-LSA	4	0x2004
AS-External-LSA	5	0x2005
Group-Membership-LSA	6	0x2006
Type 7-LSA	7	0x2007
Link-LSA NEW	8	0x2008
Intra-Area-Prefix-LSA	9	0x2009

Enhanced Routing Protocol Support
OSPFv3 Configuration Example

Cisco.com

```
Router1#
interface POS1/1
 ipv6 address 2001:db8:FFFF:1::1/64
 ipv6 enable
 ipv6 ospf 100 area 0

interface POS2/0
 ipv6 address 3FFE:B00:FFFF:1::2/64
 ipv6 enable
 ipv6 ospf 100 area 1

 ipv6 router ospf 100
   router-id 10.1.1.3

Router2#
interface POS3/0
 ipv6 address 3FFE:B00:FFFF:1::1/64
 ipv6 enable
 ipv6 ospf 100 area 1

ipv6 router ospf 100
   router-id 10.1.1.4
```

Do It Again…

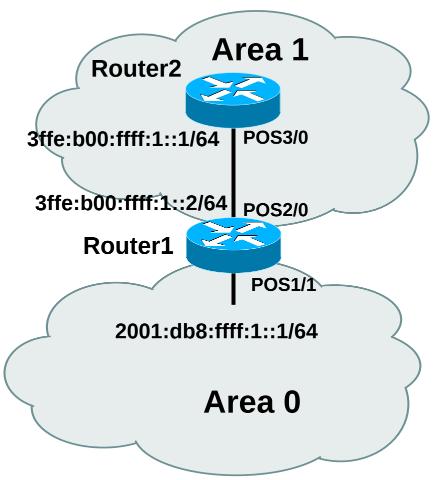

BGP-4 EXTENSIONS FOR IPv6 (RFC 2545)

BGP-4 Extensions for IPv6

Cisco.com

- **BGP-4 carries only 3 pieces of information which is truly IPv4 specific:**
  - NLRI in the UPDATE message contains an IPv4 prefix
  - NEXT_HOP path attribute in the UPDATE message contains an IPv4 address
  - BGP Identifier is in the OPEN message and AGGREGATOR attribute
- **To make BGP-4 available for other network layer protocols, RFC 2858 (obsoletes RFC 2283) defines multi-protocol extensions for BGP-4**
  - Enables BGP-4 to carry information of other protocols e.g MPLS,IPv6
  - New BGP-4 optional and non-transitive attributes:
    - MP_REACH_NLRI
    - MP_UNREACH_NLRI
  - Protocol independent NEXT_HOP attribute
  - Protocol independent NLRI attribute

BGP-4 Extensions for IPv6

Cisco.com

- New optional and non-transitive BGP attributes:
  - MP_REACH_NLRI (Attribute code: 14)
    - "Carry the set of reachable destinations together with the next-hop information to be used for forwarding to these destinations" (RFC2858)
  - MP_UNREACH_NLRI (Attribute code: 15)
    - Carry the set of unreachable destinations
- Attribute 14 and 15 contains one or more triples:
  - Address Family Information (AFI)
  - Next-hop information (must be of the same address family)
  - NLRI

BGP-4 Extensions for IPv6

Cisco.com

- Address Family Information (AFI) for IPv6
  - AFI = 2 (RFC 1700)
  - Sub-AFI = 1 Unicast
  - Sub-AFI = 2 (Multicast for RPF check)
  - Sub-AFI = 3 for both Unicast and Multicast
  - Sub-AFI = 4 Label
  - Sub-AFI = 128 VPN

BGP-4 Extensions for IPv6

Cisco.com

- Next-hop contains a global IPv6 address or potentially a link local (for iBGP update this has to be change to global IPv6 address with route-map)
- The value of the length of the next hop field on MP_REACH_NLRI attribute is set to 16 when only global is present and is set to 32 if link local is present as well
- Link local address as a next-hop is only set if the BGP peer shares the subnet with both routers (advertising and advertised)

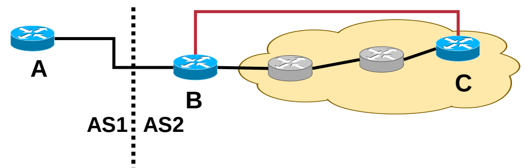

RST-1305
9719_05_2004_c2

BGP-4 Extensions for IPv6

Cisco.com

- TCP Interaction

  BGP-4 runs on top of TCP

  This connection could be setup either over IPv4 or IPv6

- Router ID

  When no IPv4 is configured, an explicit BGP router-id needs to be configured

  This is needed as a BGP Identifier, this is used as a tie breaker, and is send within the OPEN message

RST-1305
9719_05_2004_c2

Presentation_ID.scr

BGP-4 Configurations for IPv6 Non-Link Local Peering

Cisco.com

Router A

router bgp 1

no bgp default ipv4 unicast

bgp router-id 1.1.1.1

neighbor 3ffe:b00:ffff:2::2
  remote-as 2

address-family ipv6

neighbor 3ffe:b00:ffff:2::2
  activate

network 2003:3:2::/64

network 2003:3:3::/64

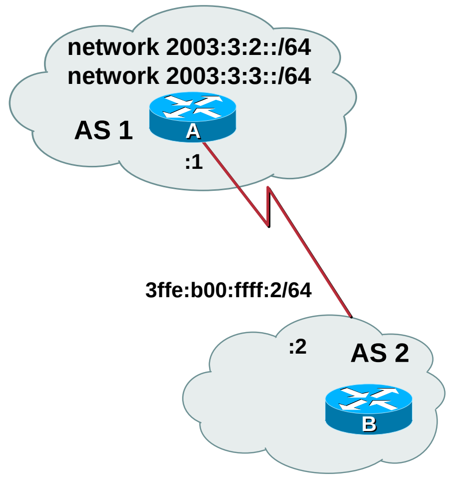

BGP-4 Configurations for IPv6 Link Local Peering

Cisco.com

Router A

Interface e2
ipv6 address 2001:db8:ffco:1::1/64

router bgp 1
no bgp default ipv4 unicast
bgp router-id 1.1.1.1
neighbor fe80::260:3eff:c043:1143 remote-as 2
neighbor fe80::260:3eff:c043:1143 update source e0
address-family ipv6
neighbor fe80::260:3eff:c043:1143 activate
neighbor fe80::260:3eff:c043:1143 route-map next-hop out

route-map next-hop
set ipv6 next-hop 2001:db8:ffco:1::1

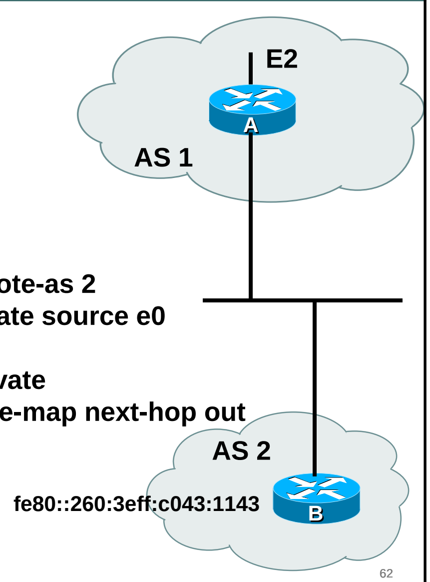

BGP-4 for IPv6 « Show Command »

Cisco.com

- Show bgp ipv6 summary
- Displays summary information regarding the state of the BGP neighbors

```
RouterA# show bgp ipv6 summary
BGP router identifier 1.1.1.1, local AS number 1
BGP table version is 69046, main routing table version 69046
92 network entries and 92 paths using 17756 bytes of memory
826 BGP path attribute entries using 43108 bytes of memory
703 BGP AS-PATH entries using 19328 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
745 BGP filter-list cache entries using 8940 bytes of memory
BGP activity 22978/18661 prefixes, 27166/22626 paths, scan interval 15 secs

Neighbor      V    AS  MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3FFE:B00:FFFF:2::2
              4    2   84194   14725    69044    0    0  3d08h          92
```

Neighbor Information   BGP Messages Activity

DEPLOYMENT

IPv4-IPv6 Transition/Coexistence

Cisco.com

- A wide range of techniques have been identified and implemented, basically falling into three categories:
  - (1) Dual-stack techniques, to allow IPv4 and IPv6 to co-exist in the same devices and networks
  - (2) Tunneling techniques, to avoid order dependencies when upgrading hosts, routers, or regions
  - (3) Translation techniques, to allow IPv6-only devices to communicate with IPv4-only devices
- Expect all of these to be used, in combination

RST-1305
9719_05_2004_c2

Dual Stack Approach

Cisco.com

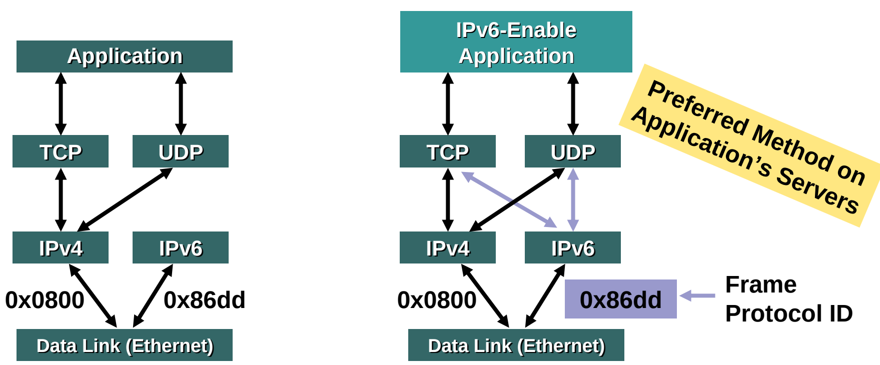

- Dual stack node means:
  - Both IPv4 and IPv6 stacks enabled
  - Applications can talk to both
  - Choice of the IP version is based on name lookup and application preference

RST-1305
9719_05_2004_c2

Host Running Dual Stack

Cisco.com

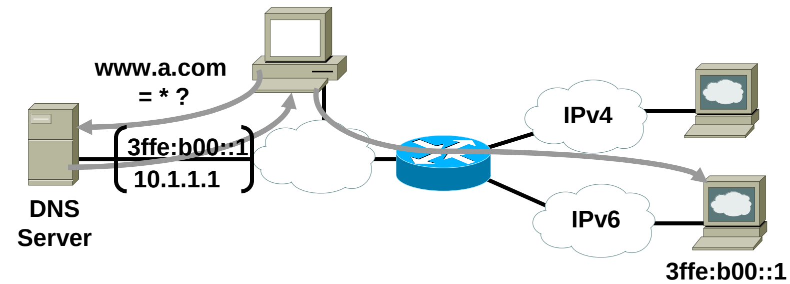

- In a dual stack case, an application that:
  - Is IPv4 and IPv6-enabled
  - Asks the DNS for all types of addresses
  - Chooses one address and, for example, connects to the IPv6 address

Cisco IOS Dual Stack Configuration

Cisco.com

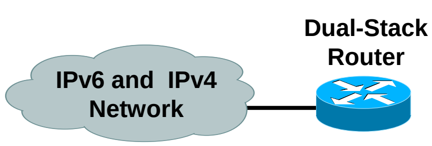


```
router#
ipv6 unicast-routing

interface Ethernet0
 ip address 192.168.99.1 255.255.255.0
 ipv6 address 2001:db8:213:1::/64 eui-64
```

IPv4: 192.168.99.1

IPv6: 2001:db8:213:1::/64 eui-64

- Cisco IOS® is IPv6-enable:
  - If IPv4 and IPv6 are configured on one interface, the router is dual-stacked
  - Telnet, Ping, Traceroute, SSH, DNS client, TFTP,…

TUNNELING

Tunneling

- Many ways to do tunneling
- Some ideas same as before

  GRE, MPLS, IP

- Native IP over data link layers

  ATM PVC, dWDM Lambda, Frame Relay PVC, Serial, Sonet/SDH, Ethernet

- Some new techniques

  Automatic tunnels using IPv4, compatible IPv6 address, 6to4, ISATAP

Manually Configured GRE Tunnel Configuration

Cisco.com

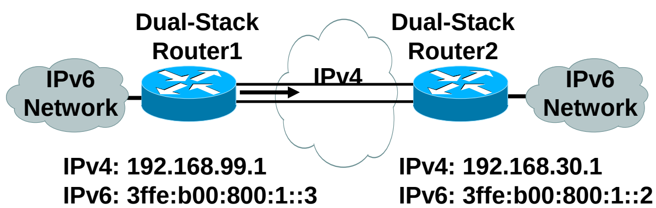


```
router1#

interface Tunnel0
 ipv6 enable
 ipv6 address 3ffe:b00:c18:1::3/128
 tunnel source 192.168.99.1
 tunnel destination 192.168.30.1
 tunnel mode gre ipv6
```

```
router2#

interface Tunnel0
 ipv6 enable
 ipv6 address 3ffe:b00:c18:1::2/128
 tunnel source 192.168.30.1
 tunnel destination 192.168.99.1
 tunnel mode gre ipv6
```


IPv6 over IPv4 Tunnels

Cisco.com

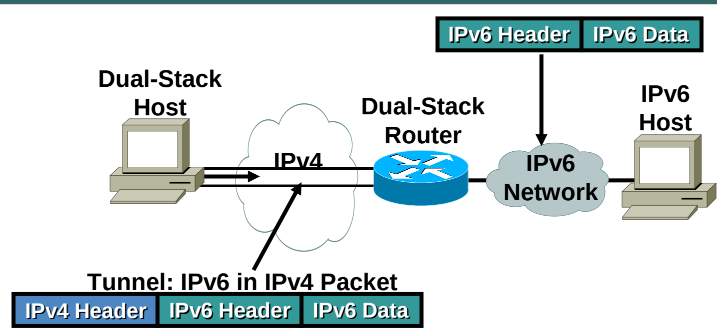

- Tunneling can be used by routers and hosts

Manually Configured
Manual Tunnel Configuration

Cisco.com

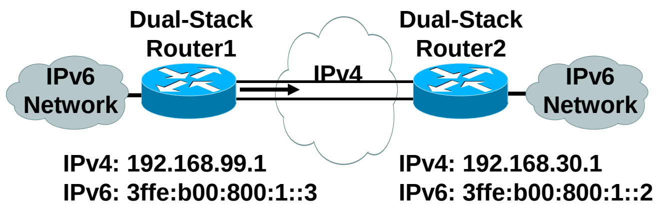


```
router1#

interface Tunnel0
 ipv6 enable
 ipv6 address 3ffe:b00:c18:1::3/127
 tunnel source 192.168.99.1
 tunnel destination 192.168.30.1
 tunnel mode ipv6ip
```

```
router2#

interface Tunnel0
 ipv6 enable
 ipv6 address 3ffe:b00:c18:1::2/127
 tunnel source 192.168.30.1
 tunnel destination 192.168.99.1
 tunnel mode ipv6ip
```


Automatic 6to4 Tunnels

Cisco.com

- Allows automatic 6to4 tunnel allows isolated IPv6 domains to connect over an IPv4 network
- Unlike the manual 6to4 the tunnels are not point to point they are multipoint tunnels
- IPv4 network is treated like a virtual NBMA network
- IPv4 is embedded in the IPv6 address is used to find the other end of the tunnel
- Address format is 2002::IPv4 address

Automatic 6to4 Tunnel (RFC 3056)

Cisco.com

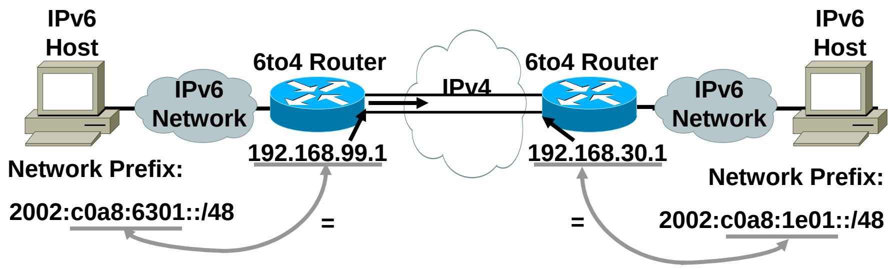

6to4:

- Is an automatic tunnel method
- Gives a prefix to the attached IPv6 network

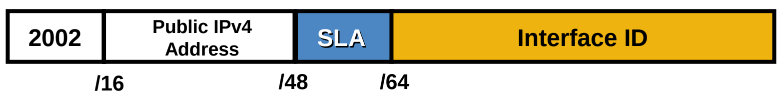

RST-1305
9719_05_2004_c2

Automatic 6to4 Tunnel (RFC 3056)

Cisco.com

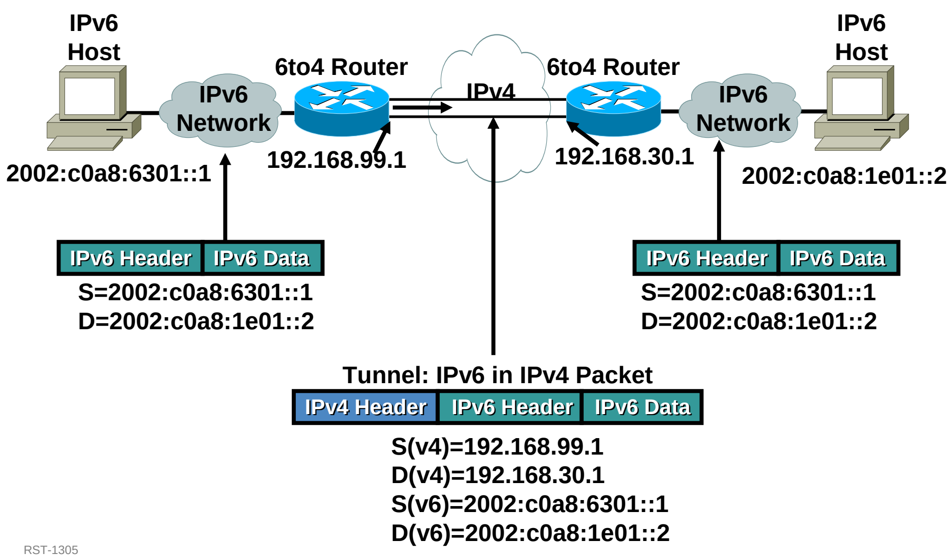

RST-1305
9719_05_2004_c2

Automatic 6to4 Configuration

Cisco.com

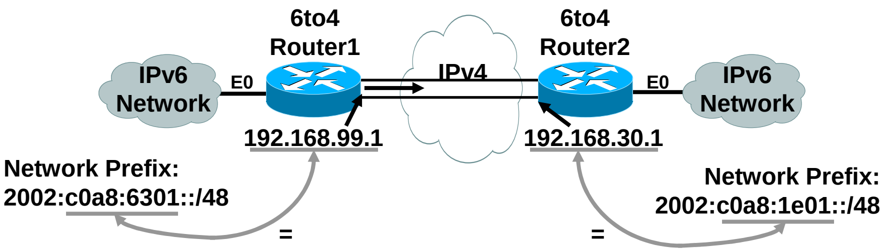


```
router1#
interface Ethernet0
 ipv6 address 2002:c0a8:6301:1::/64 eui-64
Interface Ethernet1
  ip address 192.168.99.1 255.255.0.0
interface Tunnel0
 ipv6 unnumbered Ethernet0
 tunnel source Ethernet1
 tunnel mode ipv6ip 6to4

ipv6 route 2002::/16 Tunnel0
```

```
router2#
interface Ethernet0
 ipv6 address 2002:c0a8:1e01:1::/64 eui-64
Interface Ethernet1
 ip address 192.168.30.1 255.255.0.0
interface Tunnel0
 ipv6 unnumbered Ethernet0
 tunnel source Ethernet1
 tunnel mode ipv6ip 6to4

ipv6 route 2002::/16 Tunnel0
```


Automatic 6to4 Relay

Cisco.com

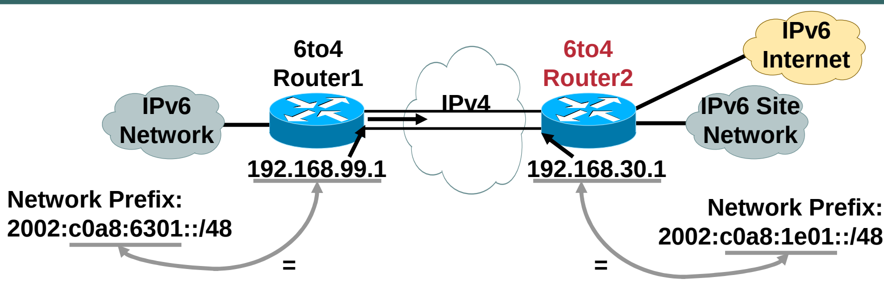

6to4 Relay:

- Is a gateway to the rest of the IPv6 Internet
- Is a default router

Automatic 6to4 Relay Configuration

Cisco.com

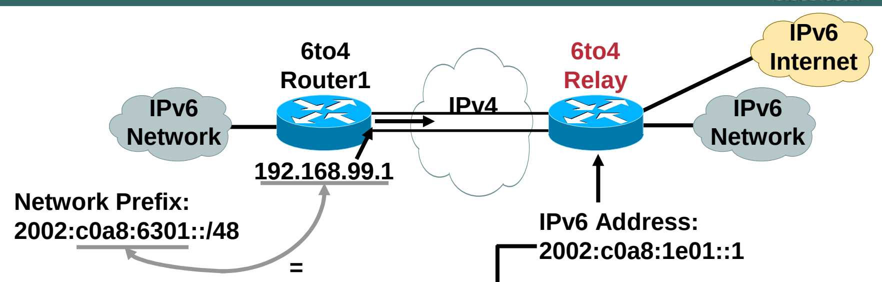


```
router1#
interface Ethernet0
 ipv6 address 2002:c0a8:6301:1::/64 eui-64
Interface Ethernet1
  ip address 192.168.99.1 255.255.0.0
interface Tunnel0
 no ip address
 ipv6 unnumbered Ethernet0
 tunnel source Ethernet1
 tunnel mode ipv6ip 6to4

ipv6 route 2002::/16 Tunnel0
ipv6 route ::/0 2002:c0a8:1e01::1
```


Automatic 6to4 Tunnels

Cisco.com

Requirements for 6to4

- Border router must be dual stack with a global IPv4 address
- Interior routing protocol for IPv6 is required
- DNS for IPv6

Intrasite Automatic Tunnel Address Protocol

Cisco.com

- This is for enterprise networks such as corporate and academic networks
- Scalable approach for incremental deployment
- ISATAP makes your IPv4 infratructure as transport (NBMA) network

Intrasite Automatic Tunnel Address Protocol

Cisco.com

- To deploy a router is identified that carries ISATAP services
- ISATAP routers need to have at least one IPv4 interface and 0 or more IPv6 interface
- DNS entries are created for each of the ISATAP routers IPv4 addresses
- Hosts will automatically discover ISATAP routers and can get access to global IPv6 network
- Host can apply the ISATAP service before all this operation but there interface will only have a link local v6 address until the first router appears

Intrasite Automatic Tunnel Address Protocol

Cisco.com

Use IANA’s OUI 00-00-5E and Encode IPv4 Address as Part of EUI-64

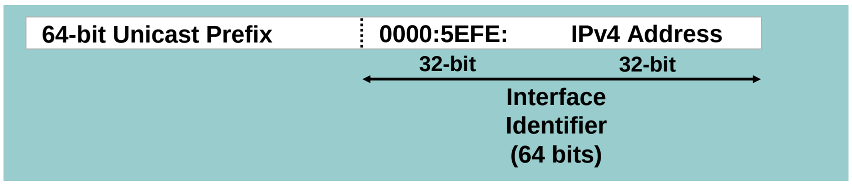

- ISATAP is used to tunnel IPv4 within as administrative domain (a site) to create a virtual IPv6 network over a IPv4 network
- Supported in Windows XP Pro SP1 and others

draft-ietf-ngtrans-isatap-11
draft-ietf-ngtrans-isatap-scenario-01

Automatic Advertisement of ISATAP Prefix

Cisco.com

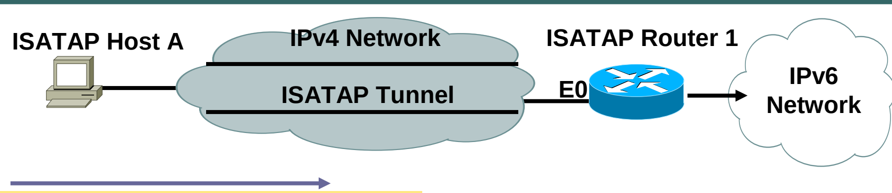

ICMPv6 Type 133 (RS)
IPv4 Source: 206.123.20.100
IPv4 Destination: 206.123.31.200
IPv6 Source: fe80::5efe:ce7b:1464
IPv6 Destination: fe80::5efe:ce7b:1fc8
Send me ISATAP Prefix

ICMPv6 Type 134 (RA)
IPv4 Source: 206.123.31.200
IPv4 Destination: 206.123.20.100
IPv6 Source: fe80::5efe:ce7b:1fc8
IPv6 Destination: fe80::5efe:ce7b:1464
ISATAP Prefix: 3ffe:b00:ffff :2::/64

Automatic Address Assignment of Host and Router

Cisco.com

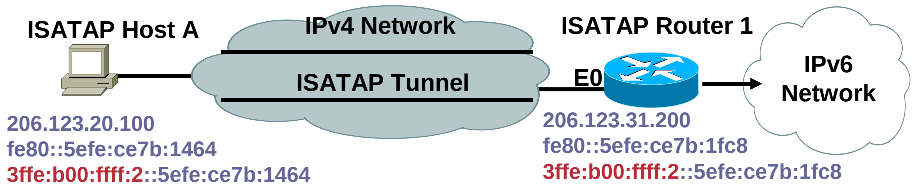

- ISATAP host A receives the ISATAP prefix 3ffe:b00:ffff:2::/64 from ISATAP Router 1
- When ISATAP host A wants to send IPv6 packets to 3ffe:b00:ffff:2::5efe:ce7b:1fc8, ISATAP host A encapsulates IPv6 packets in IPv4
- The IPv4 packets of the IPv6 encapsulated packets use IPv4 source and destination address

Automatic Configuring ISATAP

Cisco.com

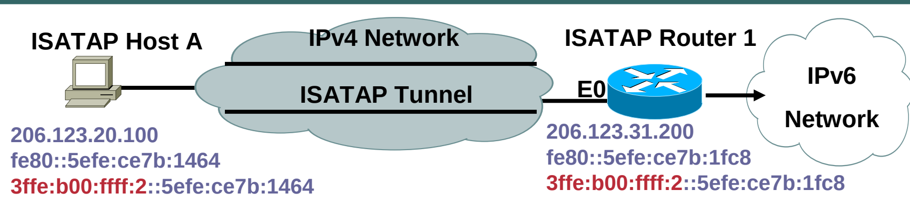


```
ISATAP-router1#
!
interface Ethernet0
 ip address 206.123.31.200
255.255.255.0
!
interface Tunnel0
 ipv6 address 3ffe:b00:ffff:2::/64
eui-64
 no ipv6 nd suppress-ra
 tunnel source Ethernet0
 tunnel mode ipv6ip isatap
```

- The tunnel source command must point to an interface with an IPv4 address configured
- Configure the ISATAP IPv6 address, and prefixes to be advertised just as you would with a native IPv6 interface
- The IPv6 address has to be configured as an EUI-64 address since the last 32 bits in the interface identifier is used as the IPv4 destination address

6PE Routing/Label Distribution

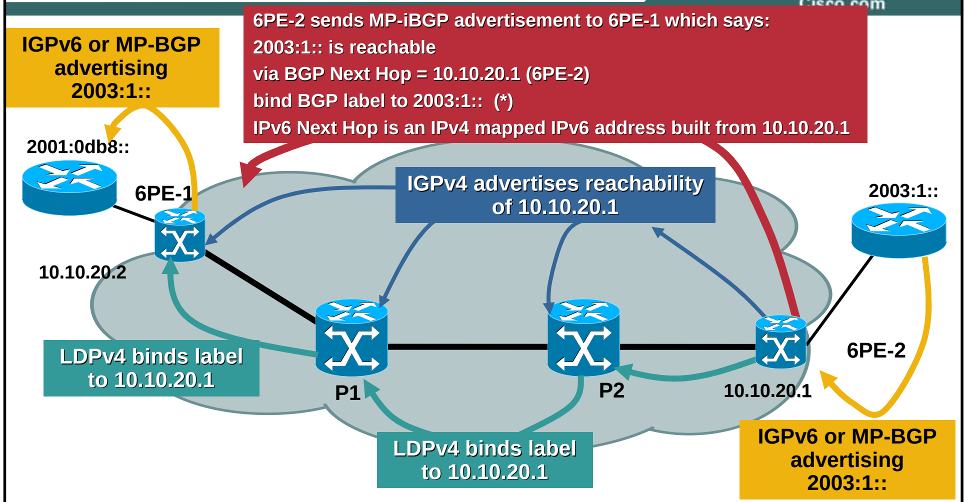

(*) The 2nd Label Allows Operations with Penultimate Hop Popping (PHP)

6PE Configuration

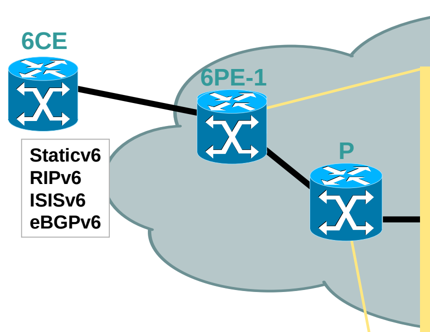


```
ipv6 cef
mpls label protocol ldp
mpls ipv6 source-interface Loopback0
mpls ldp router-id loopback0
!
interface Loopback0
 ip address 10.10.20.2 255.255.255.255
 ipv6 address 2003::/64 eui-64
!
router bgp 100
 no synchronization
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 10.10.20.1 remote-as 100
 neighbor 10.10.20.1 update-source Loopback0
 !
 address-family ipv6
 neighbor 10.10.20.1 activate
 neighbor 10.10.20.1 send-label
 redistribute connected
 redistribute rip ripv6CE1
exit-address-family
!
```

```
ip cef
mpls label protocol ldp
tag-switching tdp router-id loopback0
!
interface Serial2/0
 ip address 10.10.10.2 255.255.255.252
 ip router isis
 tag-switching ip
!
```

Show bgp ipv6 <ipv6-prefix>

Cisco.com

```
6PE-1> show bgp ipv6 2003:1:1:30::/64

BGP routing table entry for 2003:1:1:30::/64, version 2
Paths: (1 available, best #1, table Global-IPv6-Table)
  Not advertised to any peer
  Local
    ::FFFF:10.10.20.1 (metric 10) from 10.10.20.1 (192.168.254.1)
      Origin incomplete, metric 0, localpref 100, valid,
      internal, best
```


Show bgp ipv6 neighbor

Cisco.com

```
6PE-1> show bgp ipv6 neighbors 10.10.20.1

BGP neighbor is 10.10.20.1,  remote AS 100, internal link
  BGP version 4, remote router ID 192.168.254.1
  BGP state = Established, up for 00:04:07
  Last read 00:00:07, hold time is 180,
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv6 Unicast: advertised and received
    ipv6 MPLS Label capability: advertised and received

 For address family: IPv6 Unicast
  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  Route refresh request: received 0, sent 0
  Sending Prefix & Label
  2 accepted prefixes consume 144 bytes
  Prefix advertised 1, suppressed 0, withdrawn 0
  Number of NLRIs in the update sent: max 1, min 0
```

IPv6: Conclusion

Cisco.com

Moving IPv6 to Production?

- **Core IPv6 specifications are well-tested and stable**
  - Some of the advanced features of IPv6 still need specification, implementation, and deployment work
- **Application, middleware and scalable deployment scenario are IPv6 focus and challenge**
- **Plan for IPv6 integration and IPv4-IPv6 coexistence**
  - Training, applications inventory, and IPv6 deployment planning
- **Cisco is committed to deliver advanced IPv6 capabilities to the Internet industry**
  - IPv6 Solutions, ABC of IPv6, e-Learning/Training, ISD,…
  - See http://www.cisco.com/ipv6

Q AND A

More Information

- **CCO IPv6**
  http://www.cisco.com/ipv6
- **The ABC of IPv6**
  http://www.cisco.com/en/US/products/sw/iosswrel/products_abc_ios_overview.html
- **IPv6 e-Learning [requires CCO username/password]**
  http://www.cisco.com/warp/customer/732/Tech/ipv6/elearning/
- **IPv6 Access Services**
  http://www.cisco.com/warp/public/732/Tech/ipv6/docs/ipv6_access_wp_v2.pdf
- **ICMPv6 Packet Types and Codes TechNote**
  http://www.cisco.com/warp/customer/105/icmpv6codes.html
- **Cisco IOS IPv6 Product Manager**
  pgrosset@cisco.com

Complete Your Online Session Evaluation!

WHAT: Complete an online session evaluation and your name will be entered into a daily drawing

WHY: Win fabulous prizes! Give us your feedback!

WHERE: Go to the Internet stations located throughout the Convention Center

HOW: Winners will be posted on the onsite Networkers Website; four winners per day

RST-1305
9719_05_2004_c2