CISCO SYSTEMS

Session ID:

IPv6 Introduction and Deployment

thomas kramer

tkramer@cisco.com

CISCO SYSTEMS

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

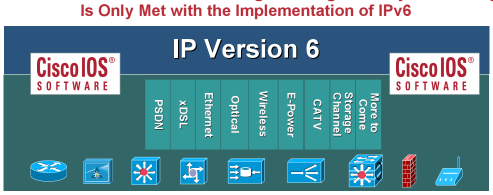

A Need for IPv6?

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!

IP Address Allocation History

Cisco.com

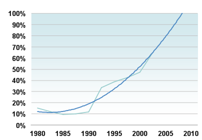

- 1981—IPv4 protocol published
- 1985—1/16 of total space
- 1990—1/8 of total space
- 1995—1/3 of total space
- 2000—1/2 of total space
- 2002.5—2/3 of total space
- This despite increasingly intense conservation efforts
  - PPP/DHCP address sharing
  - NAT (network address translation)
  - CIDR (classless 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)

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

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

CISCO SYSTEMS

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

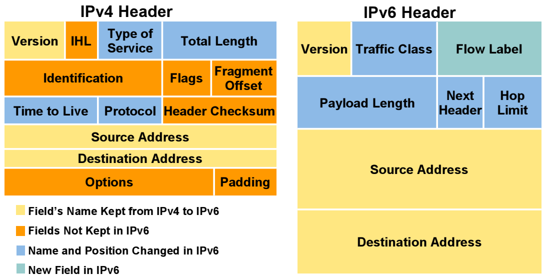

The Chain of Pointers Formed by the Next Header Field

Cisco.com

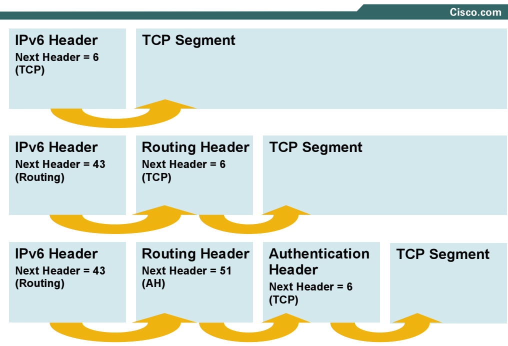

IPv6 Protocol

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 – Flow Label

Cisco.com

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

Flow-Label

T-Class 0-7 time sensitive, 8-15 non-flow traffic

Flow-ID chosen randomly by source, no conflict possible, because a flow is flow-label + source + destination

4 bits	24 bits
TCLASS	Flow Identifier

CISCO SYSTEMS

ADDRESSING

Addressing

Cisco.com

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

Addressing

Three Types of Address

1. Unicast
2. Multicast
3. Anycast

… no more broadcast addresses

Addressing

Representation

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

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

Cisco.com

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

Let's Talk a Little More on Anycast

Cisco.com

RFC 1546:

"… where a host, application, or user wishes to locate a host which supports a particular service but, if several servers support the service, does not particularly care which server is used "

- Anycast: An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the "nearest" one, according to the routing protocols' measure of distance).
- Anycast addresses are taken from the unicast address spaces (of any scope) and are not syntactically distinguishable from unicast addresses.

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

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

IPv6: Addressing Model

Cisco.com

- Addresses are assigned to interfaces
  - Change from IPv4 model:
- Interface "expected" to have multiple addresses
- Addresses have scope
  - Link local
  - Unique local
  - Global
- Addresses have lifetime
  - Valid and preferred lifetime

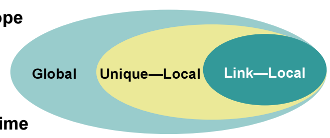

Aggregatable Global Unicast Addresses

Cisco.com

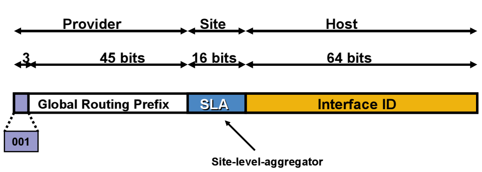

- Aggregatable Global Unicast addresses are:
  - Addresses for generic use of IPv6
  - Structured as a hierarchy to keep the aggregation

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

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

Aggregatable Global Unicast Addresses

Cisco.com

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

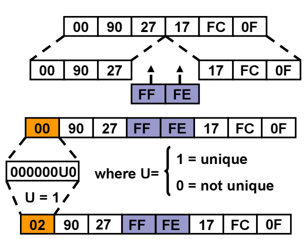

IPv6 Multicast Address

- 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

- 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 (duplicate address detection, details later)

Solicited-Node Multicast Address

- 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

Cisco.com

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

Cisco.com

1 - ICMP Type = 133 (RS)
Src = Link-local Address (FE80::/10)
Dst = All-routers multicast Address (FF02::2)
Query= please send RA

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

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

Neighbor Solicitation

Cisco.com

- 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

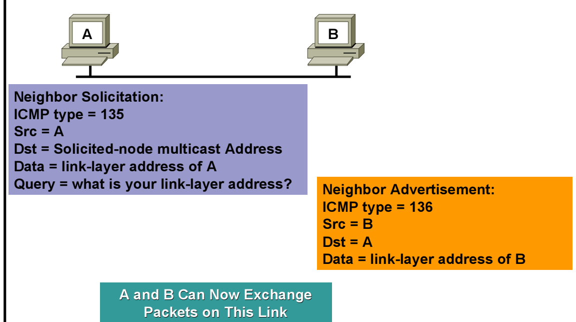

Redirect

Cisco.com

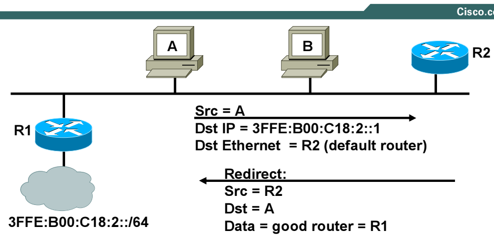

- Redirect is used by a router to signal the reroute of a packet to a better router

Cisco Systems

ROUTING

Cisco Systems

RIPNG (RFC 2080)

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

Enhanced Routing Protocol Support RIPng Configuration and Display

Cisco.com

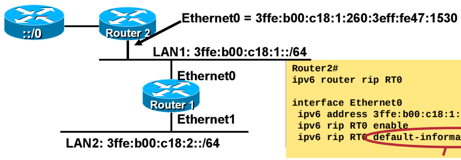


```
Router2#
ipv6 router rip RT0

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

CISCO SYSTEMS

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

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

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

RST-2214
11005_04_2005_c2

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

RST-2214
11005_04_2005_c2

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-2214
11005_04_2005_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

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

© 2005 Cisco Systems, Inc. All rights reserved.

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…

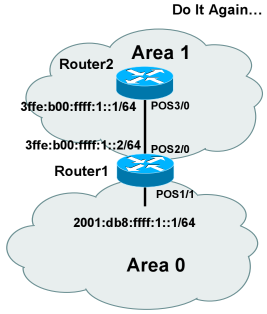

© 2005 Cisco Systems, Inc. All rights reserved.

Cisco Systems

BGP-4 EXTENSIONS FOR IPv6 (RFC 2545)

BGP-4 Extensions for IPv6

- 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

- 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)

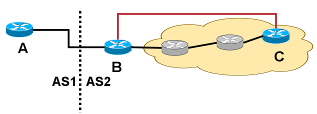

BGP-4 Extensions for IPv6

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

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

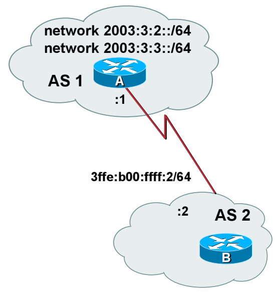

BGP-4 Configurations for IPv6 Link Local Peering

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

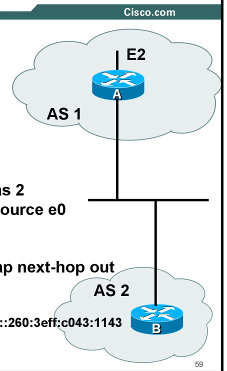

BGP-4 for IPv6 « Show Command »

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

Cisco Systems

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

Dual Stack Approach

Cisco.com

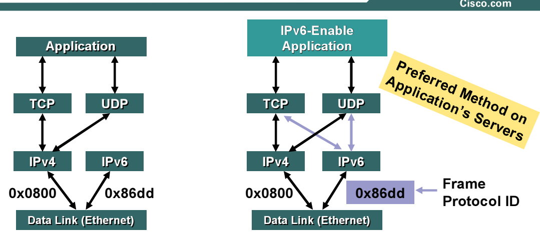

- Dual stack node means:
  - Both IPv4 and IPv6 stacks enabled
  - Applications can talk to both
  - Choice of the IP version is based on name lookup and application preference

Host Running Dual Stack

Cisco.com

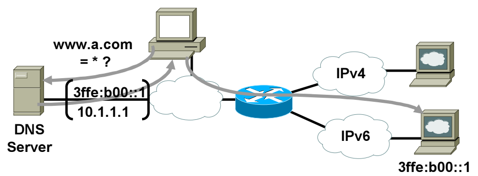

- In a dual stack case, an application that:
  - Is IPv4 and IPv6-enabled
  - Asks the DNS for all types of addresses
  - Chooses one address and, for example, connects to the IPv6 address

Cisco IOS Dual Stack Configuration

Cisco.com

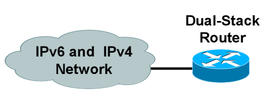


```
router#
ipv6 unicast-routing

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

- 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,…

Cisco Systems

TUNNELING

Tunneling

- Many ways to do tunneling
- Some ideas same as before
  - MPLS, GRE, 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

6PE Routing/Label Distribution

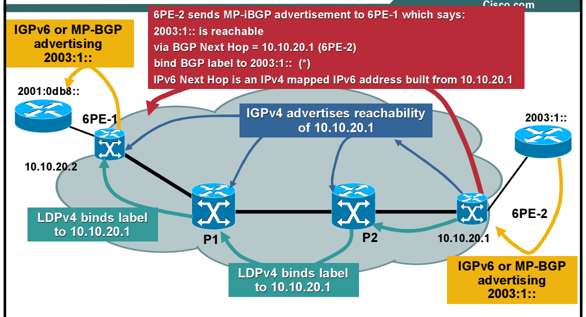

(*) The 2nd Label Allows Operations with Penultimate Hop Popping (PHP)

6PE Configuration

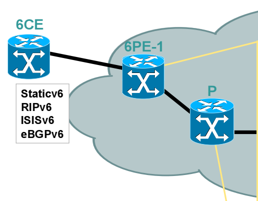


```
ipv6 cef
mpls label protocol ldp
mpls ipv6 source-interface Loopback0
mpls ldp router-id loopback0
!
interface Loopback0
 ip address 10.10.20.2 255.255.255.255
 ipv6 address 2003::/64 eui-64
!
router bgp 100
 no synchronization
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 10.10.20.1 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>


```
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


```
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
```


Manually Configured GRE Tunnel Configuration

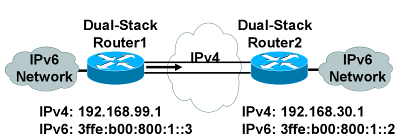


```
router1#

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

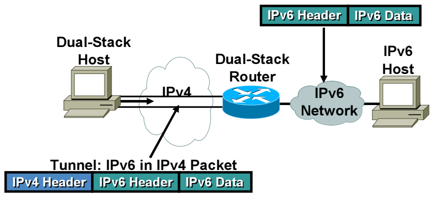

- Tunneling can be used by routers and hosts

Manually Configured
Manual Tunnel Configuration

Cisco.com

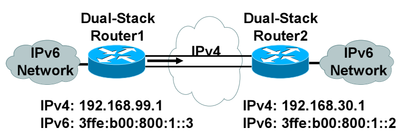


```
router1#

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

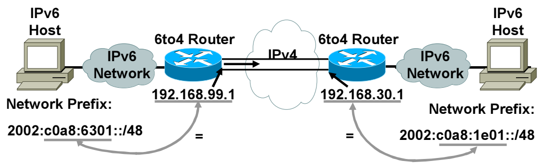

6to4:

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

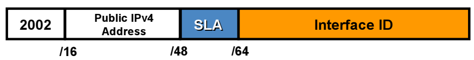

(Note: 2002::/16 is a specially assigned address range for the specific purpose of 6to4.)

Automatic 6to4 Tunnel (RFC 3056)

Cisco.com

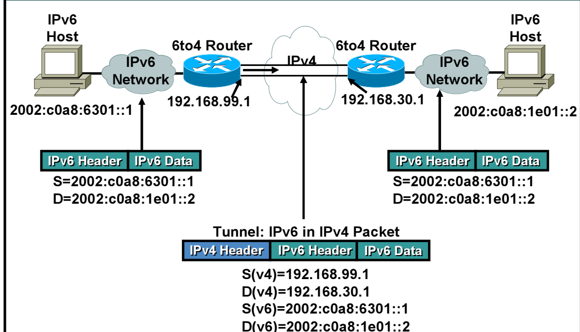

Automatic 6to4 Configuration

Cisco.com

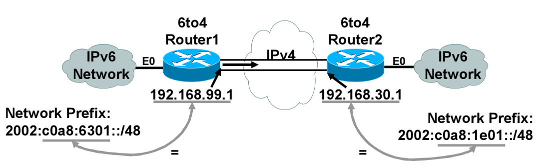


```
router1#
interface Ethernet0
 ipv6 address 2002:c0a8:6301:1::/64 eui-64
Interface Ethernet1
  ip address 192.168.99.1 255.255.0.0
interface Tunnel0
 ipv6 unnumbered Ethernet0
 tunnel source Ethernet1
 tunnel mode ipv6ip 6to4

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

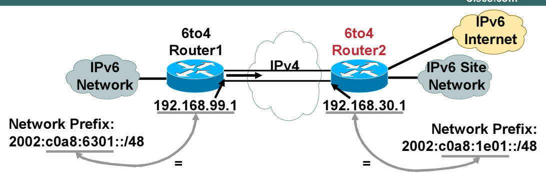

6to4 Relay:

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

Automatic 6to4 Relay Configuration

Cisco.com

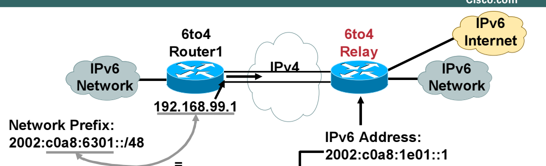


```
router1#
interface Ethernet0
 ipv6 address 2002:c0a8:6301:1::/64 eui-64
Interface Ethernet1
  ip address 192.168.99.1 255.255.0.0
interface Tunnel0
 no ip address
 ipv6 unnumbered Ethernet0
 tunnel source Ethernet1
 tunnel mode ipv6ip 6to4

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

Modified EUI-64 address, that embeds IPv4

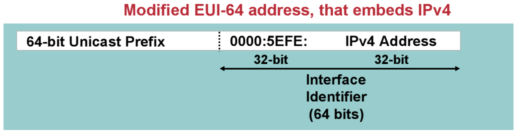

- ISATAP is used to tunnel IPv4 within as administrative domain (a site) to create a virtual IPv6 network over a IPv4 network
- Supported in Windows XP Pro SP1 and others

draft-ietf-ngtrans-isatap-22
draft-ietf-ngtrans-isatap-scenario-01

Automatic Advertisement of ISATAP Prefix

Cisco.com

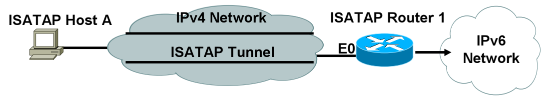

ICMPv6 Type 133 (RS)
IPv4 Source: 206.123.20.100
IPv4 Destination: 206.123.31.200
IPv6 Source: fe80::5efe:ce7b:1464
IPv6 Destination: fe80::5efe:ce7b:1fc8
Send me ISATAP Prefix

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

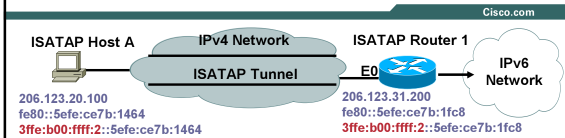

- ISATAP host A receives the ISATAP prefix 3ffe:b00:ffff:2::/64 from ISATAP Router 1
- When ISATAP host A wants to send IPv6 packets to 3ffe:b00:ffff:2::5efe:ce7b:1fc8, ISATAP host A encapsulates IPv6 packets in IPv4
- The IPv4 packets of the IPv6 encapsulated packets use IPv4 source and destination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

Automatic Configuring ISATAP

Cisco.com

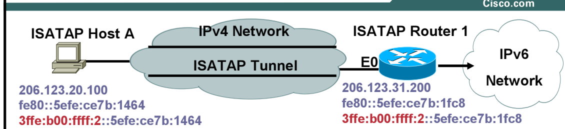


```
ISATAP-router1#
!
interface Ethernet0
 ip address 206.123.31.200
255.255.255.0
!
interface Tunnel0
 ipv6 address 3ffe:b00:ffff:2::/64
eui-64
 no ipv6 nd suppress-ra
 tunnel source Ethernet0
 tunnel mode ipv6ip isatap
```

- 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

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

What was IPv5 ?

Cisco.com

Q AND A

So, why IPv6 and not IPv5

Cisco.com

Internet Stream Protocol (ST, ST2, ST+)

- late 70s
- Accidentally given IPv5

IPv6 was almost called IPv7 because it was thought (before extensive digging through RFCs) that IPv6 had been taken as well!

Cisco.com

Cisco Systems

Cisco Systems

IPV6 MULTICAST TECHNICAL VERSION 2.1

WHY CISCO IOS IPV6 MULTICAST?

Why IPv6 Multicast – End-to-End, Fully Integrated Solutions

- **IPv6 Multicast is deployed as a part of end-to-end solutions:**
  - Target Applications:
    - New - interactive TV, gaming, mobile services, conferencing
    - Not-so-new - application control traffic, software distribution, streaming media and content-delivery
  - IPv6 Stacks and Applications supporting Multicast
  - Network infrastructure

IPV6 VERSUS IPV4 MULTICAST

IPv4 versus IPv6 Multicast

IP Service	IPv4 Solution	IPv6 Solution
Address Range	32-bit, class D	128-bit
Routing	Protocol Independent all IGPs, and BGP4+	Protocol Independent all IGPs and BGP4+ with v6 multicast SAFI
Forwarding	PIM-DM, PIM-SM, PIM-SSM, PIM-bidir	PIM-SM, PIM-SSM, PIM-bidir
Group Management	IGMPv1, v2, v3	MLDv1, v2
Domain Control	Boundary/Border	Scope Identifier
Inter-domain Solutions	MSDP across Independent PIM Domains	Single Rendezvous Point within Globally Shared Domains

ADDRESSES

IPv6 Multicast Addresses (RFC 2373)

Cisco.com

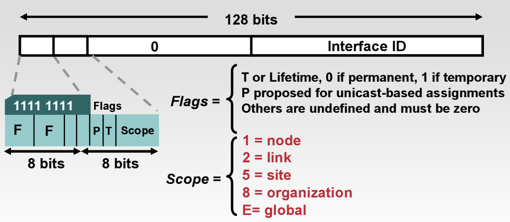

Permanently-Assigned Address Example

Cisco.com

- The "meaning" of a permanently-assigned multicast address is independent of the scope value

  For example, if the "NTP servers group" is assigned a permanent multicast address with a group ID of 101 (hex), then:

  FF01:0:0:0:0:0:0:101 – all NTP servers are on the same node as the sender

  FF02:0:0:0:0:0:0:101 – all NTP servers are on the same link as the sender

  FF05:0:0:0:0:0:0:101 – all NTP servers are at the same site as the sender

  FF0E:0:0:0:0:0:0:101 – all NTP servers are in the internet

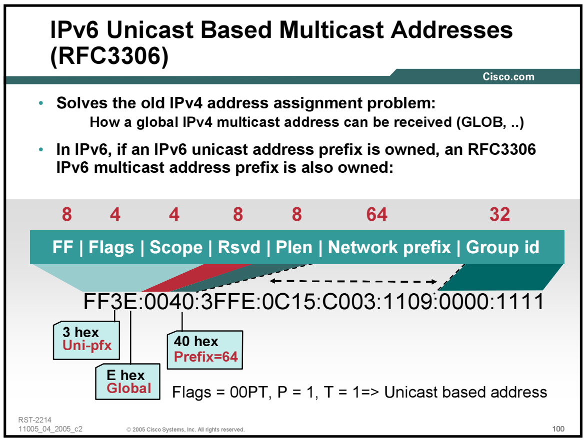
IPv6 Unicast Based Multicast Addresses (RFC3306)
Cisco.com
Solves the old IPv4 address assignment problem:
How a global IPv4 multicast address can be received (GLOB, ..)
In IPv6, if an IPv6 unicast address prefix is owned, an RFC3306 IPv6 multicast address prefix is also owned:
8 4 4 8 8 64 32
FF | Flags | Scope | Rsvd | Plen | Network prefix | Group id
FF3E:0040:3FFE:0C15:C003:1109:0000:1111
3 hex Uni-pfx
E hex Global
40 hex Prefix=64
Flags = 00PT, P = 1, T = 1=> Unicast based address
RST-2214
11005_04_2005_c2
© 2005 Cisco Systems, Inc. All rights reserved.
100

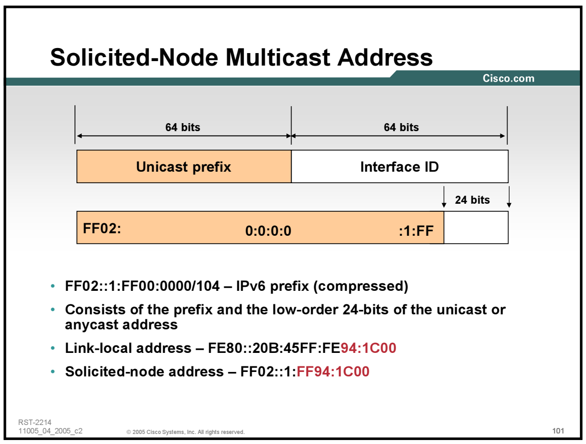
Solicited-Node Multicast Address
Cisco.com
64 bits
64 bits
Unicast prefix
Interface ID
24 bits
FF02:
0:0:0:0
:1:FF
FF02::1:FF00:0000/104 – IPv6 prefix (compressed)
Consists of the prefix and the low-order 24-bits of the unicast or anycast address
Link-local address – FE80::20B:45FF:FE94:1C00
Solicited-node address – FF02::1:FF94:1C00
RST-2214
11005_04_2005_c2
© 2005 Cisco Systems, Inc. All rights reserved.
101

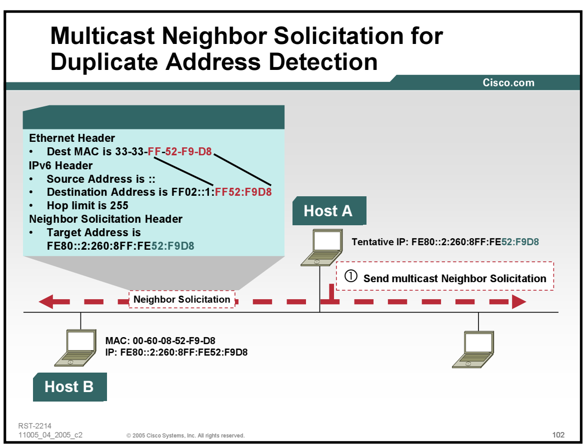
Multicast Neighbor Solicitation for Duplicate Address Detection
Cisco.com
Ethernet Header
• Dest MAC is 33-33-FF-52-F9-D8
IPv6 Header
• Source Address is ::
• Destination Address is FF02::1:FF52:F9D8
• Hop limit is 255
Neighbor Solicitation Header
• Target Address is FE80::2:260:8FF:FE52:F9D8
Host A
Tentative IP: FE80::2:260:8FF:FE52:F9D8
① Send multicast Neighbor Solicitation
Neighbor Solicitation
MAC: 00-60-08-52-F9-D8
IP: FE80::2:260:8FF:FE52:F9D8
Host B
RST-2214
11005_04_2005_c2
© 2005 Cisco Systems, Inc. All rights reserved.
102

Multicast Neighbor Advertisement (Response)

Cisco.com

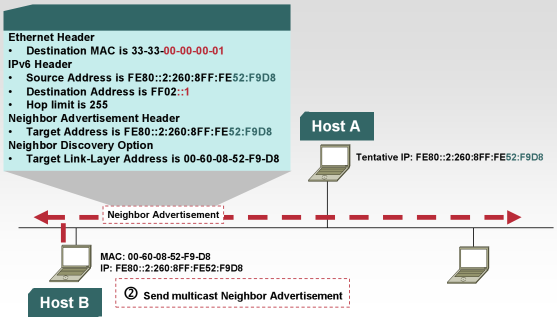

SERVICE MODELS

IPv6 Multicast Service Models

Cisco.com

- Any Source Multicast (ASM)
  - Traditionally called IP Multicast
  - Service description: RFC1112 (no update for IPv6 done yet)
  - MLDv1 RFC2710 or MLDv2 draft-vida-mld-v2-xx.txt
  - PIM-Sparse Mode (PIM-SM) draft-ietf-pim-sm-v2-new-xx.txt
  - Bidirectional PIM (PIM-bidir) draft-ietf-pim-bidir-xx.txt
- Source Specific Multicast (SSM)
  - Service description (IPv4/IPv6): draft-ietf-ssm-overview-xx.txt
  - MLDv2 required
  - PIM-SSM – not a separate protocol, just a subset of PIM-SM
  - Unicast prefix based multicast addresses ff30::/12
  - SSM range is ff3X::/32 (current allocation is from ff3X::/96)

IPv6 Multicast Service Models (Cont.)

Cisco.com

- Which service to use:
  - ASM
  - SSM
- Easily co-exist together (choice is not required)
- Cisco is committed to support both services best and foremost according to customer requirements
- Recommendations:
  - Use SSM for media-broadcast or interdomain applications due to simplicity and protection from denial of service (DoS) attacks
    - Requires moderate amount of application side work
  - Use ASM for legacy, dynamic- or many-source multi-party application, try to limit their use to Intradomain:
    - Start with PIM-SM and consider Bidir-PIM for many-source applications
    - Use PIM-SM with embedded Rendezvous Point for simple and reliable Interdomain ASM

MULTICAST REVERSE PATH FORWARDING SELECTION

Cisco IOS IP Multicast RPF

Cisco.com

- Reverse Path Forwarding (RPF) for PIM
  - Same task in IPv6 as in IPv4
  - Select a route derived from network topology routing protocols that determines where to send PIM joins to
- Cisco IOS RPF selection changes/improvements
  - IPv6 versus IPv4:
    - Longest-match-first instead of Distance-Only lookup (v4)
    - Improved static (m) routes
    - Unchanged support for all routes from Unicast RIB, except for Border Gateway Protocol (BGP) (Intermediate System-to-Intermediate System (ISIS), Open Shortest Path First (OSPF), Enhanced IGPR (EIGRP), Routing Information Protocol (RIP))
    - IPv6 „MBGP“ supported
    - IPv6 Unicast BGP behavior changed
    - No DVMRP (not defined for IPv6)

Cisco IOS IP Multicast RPF (Cont.)

Cisco.com

- IPv4 RPF selection (distance-only)
  - Consider the following tables for RPF information:
    1. IPv4 MBGP RIB
    2. IPv4 static mroutes
    3. IPv4 DVMRP RIB
    4. IPv4 unicast RIB
  - Find route with lowest administrative distance across these tables (within each table a longest-mask-first lookup is done!)
  - If two or more tables have routes with identical distance, take first route according to above table order
- Can select a route with shorter prefix-length over one with a longer prefix
  - Originally devised to allow overriding a whole set of unicast RIB with one (or substantially fewer) static mroute(s)

Cisco IOS IP Multicast RPF (Cont.)

Cisco.com

- IPv6 RPF selection (longest-match-first)
  - Consider the following tables for RPF information:
    1. IPv6 static (m)routes
    2. IPv6 MBGP RIB
    3. IPv6 unicast RIB excluding BGP routes
  - Find route with longest prefix-length (mask) route across these tables
  - If there are routes with equal prefix-length in two or more (three) of these tables, use the one with the lowest administrative distance
  - If two or more tables have routes with identical (longest) prefix-length and equal distance, take the first route according to above table order
- Same algorithm is used in IP Unicast route selection
  - In IP Unicast, results of route selection are merged into the RIB
  - In IPv6 multicast, the above selection is done on demand (today)
  - The result of the selection is a (virtual) Multicast RPF RIB with the same rules applied as in the RIB merging routes from multicast routing protocols

Cisco IOS IP Multicast RPF (Cont.)

Cisco.com

- Improved static (m)routes
  - Cisco IOS IPv4
    - ip route dest mask [nexthop | interface] [distance]
    - ip mroute dest mask [nexthop | interface] [distance]
  - Cisco IOS IPv6
    - ipv6 route dest mask [nexthop | interface]
      [distance] [mdistance | unicast | multicast]
    - Same behavior as ipv4 static routes unless new options are used
    - Support equal-cost multipath mroutes
    - Support unicast only static routes
    - Equal or less than in IPv4 config lines are required

Cisco IOS IP Multicast RPF (Cont.)

Cisco.com

- IPv4 BGP/MBGP
  - Unicast: BGP4, RFC1771
    - BGP4 predates IP multicast with PIM
    - BGP4 routes do not convey whether the announced prefix is IP multicast reachable or not
    - Many users without support for MBGP (older versions of Cisco IOS Software or other vendors) have used and are still using BGP4 to indicate IP multicast reachability (sometimes only for simplicity)
  - Multicast: MBGP, RFC2858, authority and format identifier (AFI)=IPv4, Sub AFI (SAFI)=2
    - SAFI=2 – the route is only usable for IP multicast, but not for IP unicast
    - RFC2858 also defines SAFI=1 (unicast only) and SAFI=3 (unicast and multicast), which are not used/implemented by Cisco IOS Software
  - BGP history requires distance based RPF lookup
    - Prefer an MBGP route with shorter prefix over longer prefix BGP route because BGP route may or may not indicate IP multicast reachability

RST-2214
11005_04_2005_c2 112

Cisco IOS IP Multicast RPF (Cont.)

Cisco.com

- IPv6 BGP/MBGP
  - Unicast: MBGP, RFC2858/2545, AFI=IPv6, SAFI=1
    - SAFI =1 indicates that these prefixes are only usable for IP unicast, but not for IP multicast
    - BGP routes in the IPv6 unicast RIB must be ignored in the IPv6 multicast RPF lookup
    - Cisco IOS Software does not support SAFI=3 (routes reachable for both multicast and unicast), because IETF has removed SAFI=3 from the next version of the MBGP RFC (in Q2'02)
  - Multicast: MBGP, RFC2858/2545, AFI=IPv6, SAFI=2
    - SAFI=2 – the route is only usable for IP multicast, but not for IP unicast.
      This type of announcement is also used in Cisco IOS IPv4 Multicast

RST-2214
11005_04_2005_c2 113

Cisco IOS IP Multicast RPF (7)

Cisco.com

- Support of non-MBGP router/deployments products (not Cisco IOS Software based)
- May not support MBGP SAFI=2 for IP multicast or users even running Cisco IOS Software
- May chose not to enable MBGP SAFI=2
  - Not recommended for transit / most multi-homed stub domains because they must always run MBGP SAFI=2, otherwise they can not transit or learn incongruent routes between unicast / multicast
  - Single homed leaf domains (enterprises) with complete congruent unicast/multicast routing can run just one instance of BGP internally, if necessary:
    - Use BGP translate-update on edge peeering to declare all BGP (IPv4) or MBGP AFI=IPv6,SAFI=1 routes to actually indicate SAFI=1 and SAFI=2
    - If enterprises routers are Cisco IOS Software based, use (hidden) command to enable Cisco IOS IPv6 multicast RPF lookup for consideration of MBGP AFI=IPv6 and SAFI=1 routes

RST-2214
11005_04_2005_c2 114

Cisco IOS IP Multicast RPF (Cont.)

Cisco.com

- Evaluation of IPv4 Multicast RPF selection:
  - IPv4 Multicast RPF selection was developed and optimized in the face of the IPv4 IP Multicast introduction
  - No ubiquitous MBGP for a long time
  - Only partial deployment of IP multicast
  - These deployment necessities resulted in a set of architectural shortcomings:
    - Unicast/Multicast routing selection algorithm dissimilarities increase the management effort to operate IP multicast and reduce the chance of infrastructure sharing (software, operational experience, ..) resulting in lower availability of IP multicast
    - Preference on distance made mixing of different routing protocols for routes of the same network prefixes is difficult and results in higher likelihood of RPF-loops due to inconsistent distance configuration or unwanted ignorance of longer prefix-length routes
    - Reliance on legacy BGP4 in the Interdomain space makes IP multicast reachability undeterminable for such prefixes

Cisco IOS IP Multicast RPF (Cont.)

Cisco.com

- Evaluation of IPv6 Multicast RPF selection:
  - IPv6 Multicast RPF selection was developed and optimized in the face of changed BGP in IPv6, expectance of hopefully ubiquitous MBGP availability, and more strategic / full network deployment of IPv6 multicast in IPv6 networks:
    - Avoid pitfalls of the Distance-only RPF selection
    - In most strategic (not tactically "hacked") deployment cases today, Distance-only and Longest-match-first RPF do not result in different RPFs
    - Routes between different routing protocols need to be overlapping and have different prefix lengths for results to be different
    - Simplify RPF selection by being as close as possible to IP unicast routing
    - Be prepared for potential multi-topology routing protocol implementations (MT-ISIS, etc..).
- Note: Cisco IOS IPv4 multicast also supports longest-match-first RPF selection via a (hidden) command
  - It was lately quite often recommended for complex newer deployments, but due to the large installed user base in IPv4 it will take some more planning to make this the default behavior in IPv4 too

MULTICAST LISTENER DISCOVERY

Multicast Listener Discover

Cisco.com

- Multicast Listener Discovery (MLD) is equivalent to IGMP in IPv4
- MLD messages are transported over ICMPv6
- MLD uses link local source addresses
- MLD packets use “Router Alert” option in IPv6 header (RFC2711)
- Version number confusion:
  - MLDv1 (RFC2710) like IGMPv2 (RFC2236)
  - MLDv2 (draft-vida-mld-v2-07) like IGMPv3 (RFC3376)

Multicast Listener Discover

Cisco.com

- Service Model requirements:
  - ASM – MLDv1 sufficient
  - SSM – Requires MLDv2
- Cisco IOS Software only provides MLDv2 router side
  - Fully backward compatible with MLDv1 on hosts
- SSM transition methods
  - Cisco IOS IPv4 multicast has 3 transition methods for IGMPv3:
    - IGMPv3 lite, URD, SSM-Mapping
    - Methods are equally applicable to IPv6 (MLDv2), but not currently planned
- MLD snooping
  - draft-ietf-magma-snoop-xx.txt
  - Under development for Cisco 7600 Series Router
- Cisco Group Management Protocol for v6 is a potential
  - Currently not considered

MLDv1 - Joining a Group (REPORT)

Cisco.com

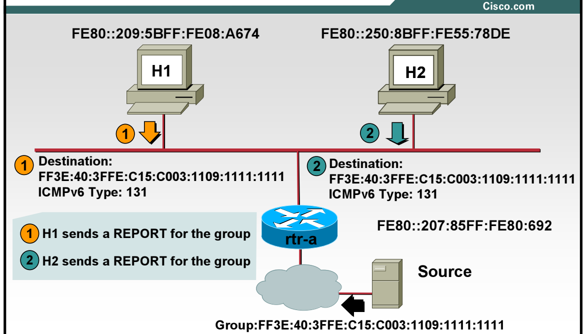

MLDv1 – Host Management (Group-Specific Query)

Cisco.com

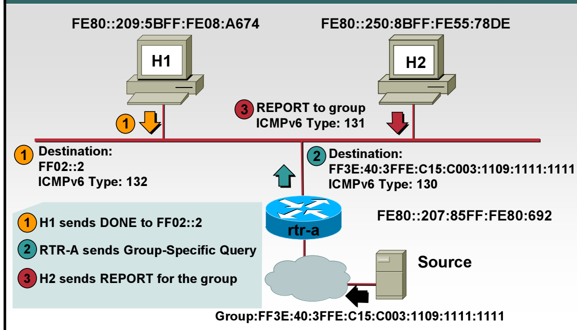

Other MLD Details

Cisco.com

- Leave/DONE
  - Last host leaves
  - Sends DONE (Type 132)
  - Router will respond with Group-Specific Query (Type 130)
  - Router will use the Last member query response interval (Default=1 sec) for each query
  - Query is sent twice, and if no reports occur then entry is removed (2 seconds)
- General Query (Type 130)
  - Sent to learn about listeners on the attached link
  - Sets the Multicast Address Field to zero
  - Sent every 125 seconds (configurable)
- MLDv2 (Type 143) problems in IETF
  - Cisco IOS Software will provide configuration (CSCed45941)

MULTICAST DOMAINS

Multicast Domains – Overview

Cisco.com

- Definitions:
  - A PIM domain is topology served by common Rendezvous Point for all sources and receivers of same group
  - A routing domain is consistent with autonomous system
- It is necessary to constrain the PIM messages, rp-mappings, and data for groups within the PIM domain:
  - In IPv4 multicast, boundary/BSR border is used
  - In IPv6, scopes and zones are used

RST-2214
11005_04_2005_c2

Multicast Domains – Service Models and Rendezvous Point Operations

Cisco.com

- SSM:
  - No Rendezvous Point or shared tree procedures (Shortest Path Tree (SPT) only)
  - MLDv2 (IPv6) required
  - FF3x::/96
- ASM:
  - Bidir-PIM or PIM-SM needs Rendezvous Point
  - Cisco IOS IPv6 multicast does not provide PIM-DM or DVMRP
- Two main elements in Rendezvous Point operations:
  - Group-to-RP mapping
    - All routers in PIM domain need to know the Rendezvous Point for groups that are PIM-SM or Bidir-PIM
    - Group-to-RP mapping also implicitly tells router whether group is PIM-SM, Bidir-PIM, or does NOT have Rendezvous Point at all
  - Redundancy: Rendezvous Point is a single point of failure

RST-2214
11005_04_2005_c2

Multicast Domains – Three Types of Rendezvous Point Operations

Cisco.com

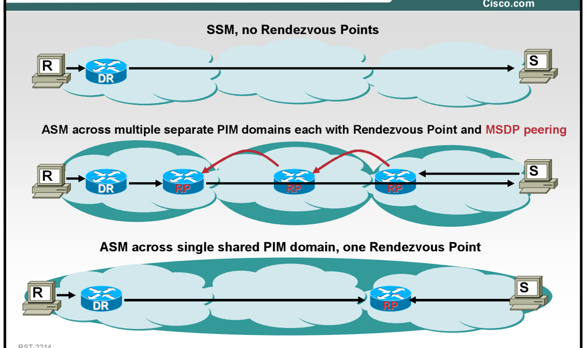

RST-2214
11005_04_2005_c2

11005_04_2005_c2.scr

SOURCE SPECIFIC MULTICAST

Source Specific Multicast

Cisco.com

- NO configuration required other than enabling *ipv6 multicast-routing*
- SSM group ranges are automatically defined

  FF3x::/96
- Very few applications support for MLDv2 yet

```
wan-agg-left#show ipv6 pim range-list
config SSM Exp: never Learnt from : ::
  FF33::/32 Up: 1d00h
  FF34::/32 Up: 1d00h
  FF35::/32 Up: 1d00h
  FF36::/32 Up: 1d00h
  FF37::/32 Up: 1d00h
  FF38::/32 Up: 1d00h
  FF39::/32 Up: 1d00h
  FF3A::/32 Up: 1d00h
  FF3B::/32 Up: 1d00h
  FF3C::/32 Up: 1d00h
  FF3D::/32 Up: 1d00h
  FF3E::/32 Up: 1d00h
  FF3F::/32 Up: 1d00h
```


SSM Mroute

Cisco.com

```
1760-branch#show ipv6 mroute
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host
        Report,
       P - Pruned, R - Rendezvous Point-bit set, F - Register flag, T -
        SPT-bit set,
       J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, State

(3FFE:C15:C003:1109::2, FF3E::DEAD), 00:03:28/never, flags: sTI
  Incoming interface: Serial1.1
  RPF nbr: FE80::210:7FF:FEDD:40
  Outgoing interface list:
    FastEthernet0, Forward, 00:03:28/never
```

RENDEZVOUS POINT OPERATIONS

Rendezvous Point Operations – Alternatives

Cisco.com

- Static Rendezvous Point
  - For PIM-SM and Bidir-PIM
  - Provides Group-to-RP mapping, but no Rendezvous Point Redundancy
- Boot Strap Router (BSR)
  - Provides Group-to-RP mapping and Rendezvous Point Redundancy
- Embedded-Rendezvous Point
  - Group-to-RP mapping only, no Rendezvous Point Redundancy
  - PIM-SM only (as of today), no Bidir-PIM
- Rendezvous Point Redundancy options for static/embedded-Rendezvous Point
  - MSDP mesh-group, PIM/Anycast ?, Prefixlength/Anycast
  - Can also be combined with BSR for faster convergence
- AutoRP
  - No option yet – IPv4 only

IPv6 Multicast Static Rendezvous Point

Cisco.com

- Easier than before, as PIM is auto-enabled on every interface

```
ipv6 multicast-routing
!
interface Loopback0
 description IPV6 IPmc RP
 no ip address
 ipv6 address 3FFE:C15:C003:110A::1/64
!
ipv6 pim rp-address 3FFE:C15:C003:110A::1/64
```

```
ipv6 multicast-routing
!
ipv6 pim rp-address 3FFE:C15:C003:110A::1/64
```

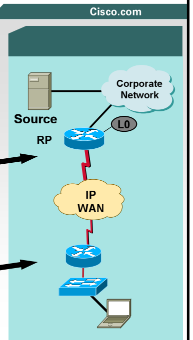

Bidirectional PIM (Bidir)

Cisco.com

- The same Many-to-Many model as before
- Configure Bidir Rendezvous Point and range via the usual ip pim rp-address syntax with the optional bidir keyword

```
ipv6 pim rp-address 3FFE:C15:C003:110A::1 bidir

2691-extra#show ipv6 pim range | include BD

Static BD RP: 3FFE:C15:C003:110A::1 Exp: never Learnt from : ::
```


IPv6 Multicast PIM BSR - Configuration

Cisco.com

```
wan-agg-right#sh run | incl ipv6 pim bsr

ipv6 pim bsr candidate-bsr 3FFE:C15:C003:1116::2
ipv6 pim bsr candidate-rp 3FFE:C15:C003:1116::2
```

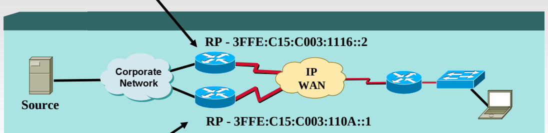


```
wan-agg-left#sh run | incl ipv6 pim bsr

ipv6 pim bsr candidate-bsr 3FFE:C15:C003:110A::1
ipv6 pim bsr candidate-rp 3FFE:C15:C003:110A::1
```


IPv6 Multicast PIM BSR – Election

Cisco.com

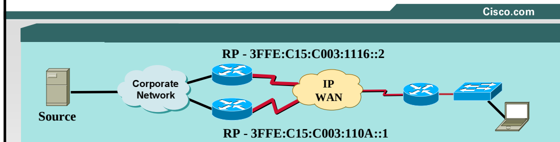


```
wan-agg-left#show ipv6 pim bsr election
PIMv2 BSR information

BSR Election Information
 Scope Range List: ff00::/8
    BSR Address: 3FFE:C15:C003:1116::2
    Uptime: 2d21h, BSR Priority: 0, Hash mask length: 126
    RPF: FE80::201:42FF:FE2D:9580,Serial2/1/0.2
    BS Timer: 00:01:44
 This system is candidate BSR
     Candidate BSR address: 3FFE:C15:C003:110A::1, priority: 0, hash mask
length:126
```


IPv6 Multicast PIM BSR – Non Rendezvous Points

Cisco.com

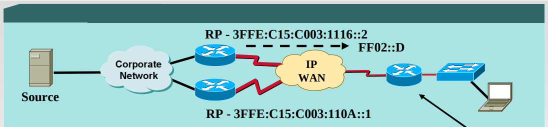


```
IPv6 BSR: Received BSR message from FE80::230:F2FF:FE15:9C00 for
3FFE:C15:C003:1116::2, priority 0 hash mask length 126
IPv6 BSR: Skipping interface FastEthernet0/0, no PIM neighbors found
IPv6 BSR: Skipping interface Serial0/2, incoming interface
IPv6 BSR: Received Group range FF00::/8, RP count 2 Fragment RP count2
IPv6 BSR: RP 3FFE:C15:C003:110A::1, Holdtime 150, Priority 192
IPv6 BSR: RP 3FFE:C15:C003:1116::2, Holdtime 150, Priority 192
```

```
BSR Election Information
  Scope Range List: ff00::/8
     BSR Address: 3FFE:C15:C003:1116::2
     Uptime: 01:51:17, BSR Priority: 0, Hash mask length: 126
     RPF: FE80::230:F2FF:FE15:9C00,Serial0/2
     BS Timer: 00:01:13
```


Embedded – Rendezvous Point Addressing Overview

Cisco.com

- draft-savola-mboned-mcast-rpaddr-00.txt
- Relies on a subset of RFC3306 IPv6 unicast-prefix multicast group addresses with special encoding rules:
  - Group address carries the Rendezvous Point address for the group!
  - From now on, for each Unicast prefix owned, 16 Rendezvous Points for each of the 16 Multicast Scopes (256 total) with 2^32 multicast groups assigned to each Rendezvous Point (2^40 total) are also owned

8	4	4	8	8	64	32
FF	Flags	Scope	Rsvd	Plen	Network prefix	Group id

New Address format defined :

Flags = 0RPT, R = 1, P = 1, T = 1=> Rendezvous Point address embedded

Example Group: FF76:0130:1234:5678:9ab0::01020304

Embedded Rendezvous Point: 1234:5678:9ab0::1

Embedded – Rendezvous Point Address Example

Cisco.com

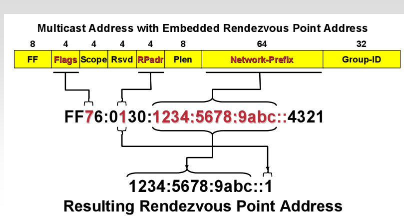

- Rendezvous Point address = network prefix + RPadr
- Sixteen Rendezvous Point addresses per network prefix

Embedded – Rendezvous Point Addressing Benefit

Cisco.com

- PIM-SM protocol operations with embedded-Rendezvous Point:
  - No change in actual PIM-SM protocol operations
  - Embedded-Rendezvous Point can simply be considered as an automatic replacement to static Rendezvous Point configuration
  - Can replace BSR for Rendezvous Point learning
  - Work simplicity due to the large address space of IPv6
    - No equivalent possible in IPv4
  - Intradomain transition into embedded-Rendezvous Point is easy:
    - Non-supporting routers simply need to be configured statically or via BSR for the embedded-Rendezvous Points

Embedded – Rendezvous Point Limitations

Cisco.com

- Embedded-Rendezvous Point is a method to learn ONE Rendezvous Point address for a multicast group:
  - Can not replace Rendezvous Point-redundancy, which is possible with BSR or MSDP/anycast-Rendezvous Point
  - Any Rendezvous Point redundancy solution that ought to work for an embedded Rendezvous Point must be some kind of anycast-Rendezvous Point solution because the embedded Rendezvous Point address is fixed through the mechanism
    - If MSDP was available for IPv6, MSDP/anycast-Rendezvous Point could be used together with embedded Rendezvous Point
- Embedded-Rendezvous Point does not yet support Bidir-PIM
  - Simply extending the mapping function to define Bidir-PIM Rendezvous Points is not sufficient:
    - In Bidir-PIM, routers carry per-Rendezvous Point state (DF per interface) prior to any data packet arriving
    - This would need to be changed in Bidir-PIM, if embedded-Rendezvous Point was to be supported

Embedded – Rendezvous Point Configuration Example

Cisco.com

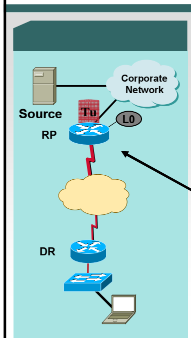

- Rendezvous Point used as an Embedded-Rendezvous Point needs to be configured with address/group range
- All other non-Rendezvous Point routers do not require any special configuration

```
ipv6 pim rp-address 3FFE:C15:C003:111D::1 ERP
!
ipv6 access-list ERP
permit ipv6 any FF7E:140:3FFE:C15:C003:111D::/96
```


Embedded Rendezvous Point – Does it work?

Cisco.com

```
branch#show ipv6 pim group

FF7E:140:3FFE:C15:C003:111D::/96*
    RP      : 3FFE:C15:C003:111D::1
    Protocol: SM
    Client  : Embedded
    Groups  : 1
    Info    : RPF: Se0/0.1,FE80::210:7FF:FEDD:40
```

```
branch#show ipv6 mroute active

Active IPv6 Multicast Sources - sending >= 4 kbps
Group: FF75:140:3FFE:C15:C003:111D:0:1112
  Source: 3FFE:C15:C003:1109::2
    Rate: 21 pps/122 kbps(1sec), 124 kbps(last 100 sec)
```

```
branch#show ipv6 pim range | include Embedded

Embedded SM RP: 3FFE:C15:C003:111D::1 Exp: never Learnt from : ::

  FF7E:140:3FFE:C15:C003:111D::/96 Up: 00:00:24
```

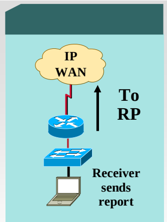

Embedded Rendezvous Point – More Details

Cisco.com

- Embedded-Rendezvous Point allows the control of the multicast groups and Rendezvous Points to be handled by the applications group
- Embedded Rendezvous Point router MUST have configure "ipv6 rp-address" for its own Rendezvous Point address to enable it to be Rendezvous Point with Embedded Rendezvous Point – for security reasons
  - Consider the issue when the WRONG Rendezvous Point is defined within the group address, and a lonely Cisco 800 Series Router on a 128k line becomes Rendezvous Point for hundreds of high-rate video streams (worse yet, IPpv6 pim spt-threshold infinity is used)
- Use the no ipv6 pim rp embedded command to disable Embedded-Rendezvous Point learning
- Embedded Rendezvous Point can be used in interdomain without any additional work:
  - Requires for all routers between sources and receivers in an application (potentially worldwide) to support the embedded-Rendezvous Point mechanism

Embedded Rendezvous Point – Interdomain Concerns and Answers

Cisco.com

- In current IPv4 Multicast:
  - Multicast group is served by typically one set of Rendezvous Points per domain and these Rendezvous Points are interconnected by MSDP
- With IPv6 embedded-Rendezvous Point:
  - There is just one set of anycast-Rendezvous Points globally for a group
- Scalability … flat virtual topology
  - Similar to SSM (with just one added Rendezvous Point – simple != scalable?)
  - No MSDP scalability / reliability / administration concerns
  - Almost arbitrary number of Rendezvous Points that can be used:
  - Each Rendezvous Point may need to serve only very few groups
- Third party Rendezvous Point dependency:
  - Yes – for totally anarchic applications that must not have a single identifiable point of origin
  - No – the majority of IP multicast applications will (for example: web applications) have an identifiable owner. This owner must take care of using an appropriate Rendezvous Point under his control

Rendezvous Point Redundancy – Overview

Cisco.com

- ASM always requires a Rendezvous Point, whether it is PIM-SM or Bidir-PIM
  - PIM-DM would be the exception to this rule for ASM
- Rendezvous Point is a single point of failure, and redundancy is a basic operational requirement
  - Today, BSR is the only available Rendezvous Point redundancy solution for IPv6:
    - Static-Rendezvous Point configuration by itself can not provide redundancy
    - MSDP (for anycast Rendezvous Point redundancy) is not defined for IPv6
    - BSR / AutoRP in IPv4 are considered to be inferior solutions to anycast:
      - Worse convergence times
      - Active protocol operations required in all routers
      - BSR has a set of limitation, but further protocol work does not seem to happen in the IETF
- An anycast-Rendezvous Point solution for IPv6 could solve the issues at hand if combined together with embedded-Rendezvous Point

Rendezvous Point Redundancy – Potential Anycast Rendezvous Point Alternatives

Cisco.com

- MSDPv6
  - Perfectly well suited to do support for anycast-Rendezvous Point (one mesh-group)
  - Complex protocol – only a small subset of functions of MSDP is really required for anycast Rendezvous Point function
  - MSDP was determined to be a temporary solution due to its intrinsic (not anycast-Rendezvous Point related) restrictions
  - Reviving MSDP for IPv6 is considered counter productive
- Draft-ietf-farinacci-pim-anycast-rp-00.txt
  - Most simple protocol is doing exactly what MSDP needs to do in one mesh-group:
    - PIM-SM register messages are unicast forwarded between the redundant Rendezvous Points
    - Almost no operational differences to MSDP for anycast-Rendezvous Point
- Prefixlength/Anycast-Rendezvous Point (Cisco internal idea)
  - Solution without any new protocol (similar to embedded-Rendezvous Point) – aka: most simple solution ?
  - Could support PIM-SM and Bidir-PIM, IPv4 and IPv6

Rendezvous Point Redundancy – Prefixlength/Anycast-Rendezvous Point

Cisco.com

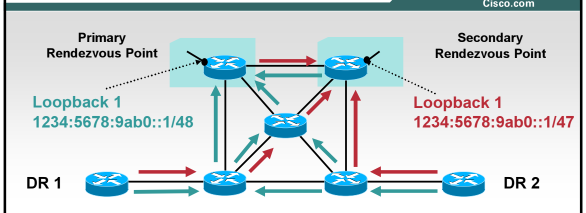

- NEW: Designate a primary and a secondary (tertiary and etc are also possible) Rendezvous Point for the anycast group
- NEW: Configure Primary Rendezvous Point with longest subnet mask on the loopback (secondary has longer mask)
- OLD: Distribute loopback interfaces routes into Interior Gateway Protocol (IGP)

Rendezvous Point Redundancy – Anycast-RP with Prefixlength Arbitration

Cisco.com

- Result:
  - All routers will converge on the primary Rendezvous Point, if it is available – because the longer mask route always wins
  - If the primary Rendezvous Point fails, failover is as fast as with the known MSDP/anycast-Rendezvous Point
    - Depends only on the convergence speed of the IGP
  - Because only one Rendezvous Point of the anycast group is active at any time, MSDP (or equivalent) is not needed
- Major difference:
  - No load-sharing between Rendezvous Points
    - Load-sharing is not necessary in IPv4
    - Load –sharing comes for free with the MSDP/anycast-Rendezvous Point redundancy
  - Scalability behavior is also different from MSDP
- No new protocol, but requires a few Cisco IOS Software code fixes on Rendezvous Point / DR before it can be used correctly

TUNNELING

A Few Notes On Tunnels

Cisco.com

- PIM uses tunnels when Rendezvous Points/Sources are known
- Source registering (on first-hop router)
  - Uses virtual tunnel interface (appear in OIL for (S,G))
  - Created automatically on first-hop router when Rendezvous Point is known
  - Cisco IOS Software keeps tunnel as long as Rendezvous Point is known
  - Unidirectional (transmit only) tunnels
  - PIM Register-Stop messages are sent directly from Rendezvous Point to registering router (not through tunnel!)

PIM Tunnels (DR-to-RP)

Cisco.com

```
branch#show ipv6 pim tunnel
Tunnel1*
  Type  : PIM Encap
  RP    : 3FFE:C15:C003:1116::2
  Source: 3FFE:C15:C003:111E::2
```

```
branch#show interface tunnel 1
Tunnel1 is up, line protocol is up
  Hardware is Tunnel
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 3FFE:C15:C003:111E::2 (Serial0/2),
destination 3FFE:C15:C003:1116::2
  Tunnel protocol/transport PIM/IPv6, key disabled,
sequencing disabled
  Checksumming of packets disabled
  Tunnel is transmit only
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
… output truncated…
```

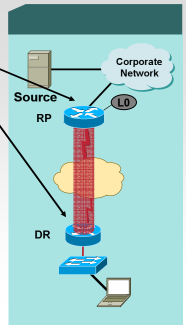

Protocol Independent Multicast Tunnels (Rendezvous Point)

Cisco.com

- Source registering on Rendezvous Point → 2 virtual tunnels are created

  First one transmits only registering sources locally connected to the Rendezvous Point

  Second one receives only decapsulating incoming registers from remote designated routers

- No one-to-one relationship between virtual tunnels on designated routers and Rendezvous Point

Protocol Independent Multicast Tunnels (RP-for-Source)

Cisco.com

```
RP-router#show ipv6 pim tunnel
Tunnel0*
  Type  : PIM Encap
  RP    : 3FFE:C15:C003:1116::2
  Source: 3FFE:C15:C003:1116::2
Tunnel1*
  Type  : PIM Decap
  RP    : 3FFE:C15:C003:1116::2
  Source:  -
```

```
RP-router#show interface tunnel 1
Tunnel1 is up, line protocol is up
  Hardware is Tunnel
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 3FFE:C15:C003:1116::2
(FastEthernet0/0), destination 3FFE:C15:C003:1116::2
  Tunnel protocol/transport PIM/IPv6, key disabled,
sequencing disabled
  Checksumming of packets disabled
  Tunnel is receive only

… output truncated…
```

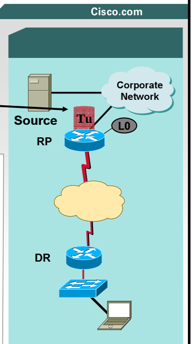

Tunneling Version 6 Multicast

Cisco.com

- v6 in v4
  - Most Widely Used
  - Tunnel mode ipv6ip
  - Intermediate System to Intermediate System (IS-IS) cannot traverse
  - v6 in v4 GRE
  - tunnel mode gre ip
  - ISATAP does not support IPv6 Multicast yet
- v6 in v6
  - Tunnel mode ipv6
  - v6 in v6 GRE
  - Tunnel mode gre ipv6

SCOPING

IPv6 Scoping Support

Cisco.com

- Scopes: draft-ietf-ipngwg-addr-arch-v3-11.txt
- Example scopes:
  - link-local (2)
  - site-local (5)
  - global (E or 14)
- Zone is a connected region of a given scope
- Initial implementation similar to v4 boundaries:
  - Can configure interface with zone and scope
  - ipv6 zone <zoneid> scope <2-15>
  - CAUTION: this is still in the working stage
- PIM messages and data traffic within that scope are ignored on that interface
- Initially a zone can only contain one interface

IPv6 ASM Solutions Summary

Cisco.com

- The solution elements (BSR, static-Rendezvous Point, embedded-Rendezvous Point, prefixlength/Anycast-Rendezvous Point, or MSDP/replacement) are not independent of each other, but can form a potential framework:
  - Due to intradomain interoperability reasons BSR is considered to be primarily important
  - ASM single-Rendezvous Point with embedded Rendezvous Point and future prefixlength/Anycast-Rendezvous Point is considered to be the best approach to reduce complexity of interdomain PIM-SM
  - Reduced complexity is considered to be an important factor to reduce TCO and improve serviceability of IPv6 multicast
  - If customers point towards MSDP / other solutions, then they will also be considered

Ease of Deployment (non-Rendezvous Point)

Cisco.com

IPv4 Multicast w/Auto-RP

```
ip multicast-routing
!
interface <all-interfaces>
ip pim sparse-dense-mode
!
ip pim rp-address 2.2.2.2 NO-DM
!
ip access-list standard NO-DM
 deny   224.0.1.39
 deny   224.0.1.40
 permit any
```

This is problematic

IPv6 Multicast w/SSM or E-RP

```
ipv6 multicast-routing
```


APPENDIX

Support Outside of Cisco - Operating Systems

Cisco.com

Vendor	IPv6 Support	Versions	More Info
Microsoft	Yes	W2K – Preview XP (SP1) and 2003 CE .NET (Pocket PC 4.1)	http://www.microsoft.com/ipv6
Sun	YES	Solaris 8 and 9	http://wwws.sun.com/software/solaris/ipv6/
IBM	YES	z/OS Rel. 1.4 AIX 4.3 - > OS/390 V2R6 eNCS	http://www-3.ibm.com/software/os/zseries/ipv6/
BSD	YES	FreeBSD 4.0 - > OpenBSD 2.7 - > NetBSD 1.5 - > BSD/OS 4.2 - >	http://www.kame.net/
Linux	YES	RH 6.2 - > Mandrake 8.0 - > SuSE 7.1 - > Debian 2.2 - >	http://www.bieringer.de/linux/IPv6/status/IPv6+Linux-status-distributions.html
HP/Compaq	YES	HP-UX 11i Tru64 UNIX V5.1 OpenVMS V5.1	http://h18000.www1.hp.com/ipv6/next_gen.html
Novell	YES	Netware 6.1	http://www.novell.com/documentation/lg/nw65/index.html?page=/documentation/lg/nw65/readme/data/ajzlp6r.html
Apple	YES	MAC OS X 10.2 - >	http://developer.apple.com/macosx/

RST-2214
11005_04_2005_c2 160

Generic Multicast Group Addresses

Cisco.com

```
|   8    |  4 |  4 |                   112 bits                    |
+------ -+----+----+-----------------------------------------------+
|11111111|flgs|scop|                   group ID                    |
+--------+----+----+-----------------------------------------------+

flgs
                              +-+-+-+-+
is a set of 4 flags:          |0|0|0|T|
                              +-+-+-+-+
```

```
scop
0  reserved
1  interface-local scope
2  link-local scope
3  Reserved
4  admin-local scope
5  site-local scope
6  (unassigned)
7  (unassigned)
8  organization-local scope
9  (unassigned)
A  (unassigned)
B  (unassigned)
C  (unassigned)
D  (unassigned)
E  global scope
F  reserved
```

- RFC3513
- FF00::/8
- T=0 → permanent IANA groups
- T=1 → FF10::/12 → user groups
- Scope 1-2 → local link only
- Scope 4 and higher → across router interfaces and links

RST-2214
11005_04_2005_c2 161

Unicast Prefix-Based Addresses

Cisco.com

```
|   8    |  4 |  4 |   8    |   8    |       64       |    32    |
+--------+----+----+--------+--------+----------------+----------+
|11111111|flgs|scop|reserved|  plen  | network prefix | group ID |
+--------+----+----+--------+--------+----------------+----------+
```

- RFC3306
- Like GLOP for IPv4 (233/8 + ASN = 256 group addresses)
- 2**32 IPv6 multicast addresses
- FF30::/12
- P=1, T=1 (must)
- plen is #significant bits in network prefix
- Example: unicast prefix 3FFE:FFFF:1::/48 → unicast prefix-based multicast prefix FF3x:0030:3FFE:FFFF:1::/96 (“30” is 48 bit length)

```
flgs
+-+-+-+-+
|0|0|P|T|
+-+-+-+-+
```

RST-2214
11005_04_2005_c2 162

SSM Addresses

Cisco.com

- Special case of unicast prefix-based addresses
- P=T=1, plen=0, network prefix=0
- FF3x::/96

Embedded Rendezvous Point Addresses

Cisco.com

```
|   8    |  4 |  4 |  4 |  4 |    8    |        64        |     32    |
+--------+----+----+----+----+--------+----------------+----------+
|11111111|flgs|scop|rsvd|RPad|  plen  | network prefix | group ID |
+--------+----+----+----+----+--------+----------------+----------+
```

- Special case of unicast prefix-based addresses
- draft-savola-mboned-mcast-rpaddr
- R=P=T=1 (must)
- Rpad is Rendezvous Point address for unicast prefix-based multicast address
- plen is #significant bits in network prefix
- FF70::/12

```
flgs
+-+-+-+-+
|0|R|P|T|
+-+-+-+-+
```


MAC Address Mapping

Cisco.com

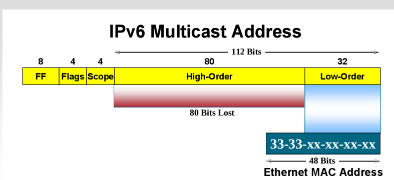

- RFC2464
- Example: FF05:1::5 → 33:33:0:0:0:5
- More than one IPv6 multicast address will map to the same MAC address (80 bit is lost)

Unicast Routing

Cisco.com

```
7500>show ipv6 rpf 2001:abcd:0:B::1
RPF information for 2001:abcd:0:B::1
  RPF interface: POS4/1/0
  RPF neighbor: FE80::208:E2FF:FE3C:300
  RPF route/mask: 2001:abcd:0:B::/64
  RPF type: MBGP
  RPF recursion count: 0
  Metric preference: 20
  Metric: 20
```

- sh ipv6 rpf

PIM Tree Building Procedures

Cisco.com

- RFC2362 and draft-ietf-pim-sm-v2-new-07
- IPv6 multicast-routing → enables PIMv6 on all IPv6 interfaces
- Disable PIM on interface: no ipv6 pim
- Note that there is difference with PIMv4!
- Only sparse mode

PIM Tree Building Procedures

Cisco.com

```
12000-2>show ipv6 pim interface

Interface          PIM  Nbr   Hello  DR
                        Count Intvl  Prior

Ethernet0          on   0     30     1
    Address: FE80::208:E2FF:FE3C:3FF
    DR     : this system
POS1/0             on   1     30     1
    Address: FE80::208:E2FF:FE3C:300
    DR     : this system
```

- sh ipv6 pim interface

PIM Tree Building Procedures

Cisco.com

- PIM neighbor resolution
  - PIM hellos and PIM neighbor discovery use link local addresses
  - PIM hellos advertise all IPv6 addresses on PIM interface to match with RPF neighbor IPv6 address

PIM Tree Building Procedures

Cisco.com

```
7200-2>show ipv6 pim neighbor detail

Neighbor Address(es)       Interface        Uptime    Expires DR pri Bidir

FE80::C0A8:F01             Tunnel10           06:50:56  00:01:30 1      B
2002::1
```

- sh ipv6 pim neighbor detail

PIM Tree Building Procedures

Cisco.com

- Restricting PIM registers at Rendezvous Point
  - ipv6 pim accept-register list <acl_name>
  - ipv6 pim accept-register route-map <rmap_name>

PIM Tree Building Procedures

Cisco.com

- Distributing Rendezvous Point information
  - Static
    - ipv6 pim rp-address <ipv6_address> <acl_name>
  - Boot-Strap Router (BSR)
    - Currently Cisco IOS Software only forwards and does not interpret or learn Rendezvous Point information
    - Full support under development
  - Embedded Rendezvous Point
    - Control from network administrator to end user
    - Make sure that physical interfaces of any router can never become Rendezvous Point address (make 65th – 124th bit non-zero) → protects against possible DoS attacks
    - Rendezvous Point address should be loopback

RST-2214
11005_04_2005_c2 172

PIM Tree Building Procedures

Cisco.com

- PIM topology information
  - Immediate OIL → Tag Information Base (TIB) (PIM Topology Table)
    - Requires explicit PIM joins or MLD responses before interfaces are added to OIL
    - sh ipv6 mroute
    - sh ipv6 pim topology
  - Inherited OIL → MRIB
    - Inherited from (*,G) Immediate OIL
    - sh ipv6 mrib route

RST-2214
11005_04_2005_c2 173

Multicast Data Forwarding

Cisco.com

- From RIB, TIB → MRIB, MFIB
- sh ipv6 mfib (enabled by default with ipv6 multicast-routing)
- no ipv6 mfib (software based hardware only; router still runs PIM)
- sh ipv6 mfib interface

RST-2214
11005_04_2005_c2 174

Multicast Data Forwarding

Cisco.com

- Software based hardware
  - Cisco 7200 Series Router:
    - centralized
  - Cisco 7500 Series Router:
    - ip cef distributed / ipv6 cef distributed

Multicast Data Forwarding

Cisco.com

```
7200-2>show ipv6 mfib ff05:1::5 verbose

IP Multicast Forwarding Information Base
Entry Flags: C - Directly Connected, S - Signal, IA - Inherit A flag,
             AR - Activity Required, D - Drop
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts: Total/RPF failed/Other drops
Interface Flags: A - Accept, F - Forward, NS - Negate Signalling
             IC - Internal Copy, NP - Not platform switched
             SP - Signal Present
Interface Counts: FS Pkt Count/PS Pkt Count
(*,FF05:1::5) Flags: C
   Forwarding: 0/0/0/0, Other: 4246/4246/0
   Tunnel10 Flags: A NS
   GigabitEthernet0/0 Flags: F NS
     Pkts: 0/0 (process switching)
   Loopback1 Flags: F IC NS
     Pkts: 0/0 (process switching)
(2001:abcd:0:D::64,FF05:1::5) Flags:
   Forwarding: 141504/100/186/145, Other: 0/0/0
   Tunnel10 Flags: A
   GigabitEthernet0/0 Flags: F NS
     Pkts: 0/94682 (process switching)
   Loopback1 Flags: F IC NS
     Pkts: 0/141504 (process switching)
```

- A (rpf interface), F (outgoing forwarding interface), IC (mld join-group) flags

Multicast Data Forwarding

Cisco.com

- Hardware based hardware (GSR)
  - Rendezvous Point forwarding (GRP) – rate limited to 1000 pps
  - Line Cards forwarding (Line Card CPU)
  - Hardware forwarding (Line Card hardware)

Multicast Data Forwarding

Cisco.com

- 12000-1>show ipv6 mfib ff05:1::5
- IP Multicast Forwarding Information Base
- Entry Flags:
  - C - Directly Connected, S – Signal, IA - Inherit A flag, AR - Activity Required, D - Drop
- Forwarding Counts:
  - Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
- Other counts:
  - Total/RPF failed/Other drops
- Interface Flags:
  - A - Accept, F - Forward, NS - Negate Signalling, IC - Internal Copy, NP - Not hardware switched, SP - Signal Present
- Interface Counts:
  - Distributed FS Pkt Count/FS Pkt Count/PS Pkt Count

Multicast Data Forwarding (Cont.)

Cisco.com

```
(*,FF05:1::5) Flags: C
  RP Forwarding: 0/0/0/0, Other: 0/0/0
  LC Forwarding: 0/0/0/0, Other: 0/0/0
  POS1/0 Flags: A
  Tunnel11 Flags: F NS
    Pkts: 0/0/0
  (2001:abcd:0:D::64,FF05:1::5) Flags:
  RP Forwarding: 7816536/29/186/42, Other: 0/0/0
  LC Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwd:   0/0/0/0, Other: 744/744/0
  POS1/0 Flags: A F
    Pkts: 0/0/0
  Tunnel11 Flags: F NS
    Pkts: 0/0/7816612
```

- Rendezvous Point, Line Card, Hardware forwarding

User to Network Signalling

Cisco.com

- MLDv2 enabled by default with ipv6 multicast-routing
- No ipv6 mld router (interface configuration) will disable MLD
  - Only router side (e.g. querying, incoming MLD host messages) will still be processed
- No ipv6 pim (interface configuration) will disable MFIB forwarding
  - MLD message processing will still be processed
- Sh ipv6 mld interface
- Restricting access:
  - ipv6 mld access-group <acl_name>

Enabling and Disabling Multicast Components

Cisco.com

- Enable IPv6 multicast globally → ipv6 multicast-routing
  - Enables PIM
  - Enables MFIB
  - Enables MLD processing
- On RSP/VIP → ipv6 cef distributed

Enabling and Disabling Multicast Components (Cont.)

Cisco.com

- Disable PIM on a per interface basis → no ipv6 pim
  - Disables PIM on interface
  - Disables MFIB on interface
  - MLD stays operational, but interface is not in OIL, and no data is forwarded
- Status → sh ipv6 pim interface

Enabling and Disabling Multicast Components (Cont.)

Cisco.com

- Disable MFIB globally → no ipv6 mfib
  - Disables all multicast data forwarding
  - PIM and MLD processing still allowed
- Disable MFIB on a per interface basis → no ipv6 mfib-cef
  - Disable multicast data forwarding in interrupt mode → traffic process switched
  - PIM and MLD processing still allowed
- Status → sh ipv6 mfib interface
  - Caveat: currently does not check MFIB status on interface
  - MFIB interface is down when line protocol on interface is down
  - Disable MFIB on interface due to “no ipv6 pim” is not reflected

Enabling and Disabling Multicast Components (Cont.)

Cisco.com

- Disable MLD on a per interface basis → no ipv6 mld router
  - Only router side (e.g. querying, incoming MLD host messages) will still be processed
- Status → sh ipv6 mld interface

RST-2214
11005_04_2005_c2

Deployment Scenarios

Cisco.com

- Single Domain Multicast
  - Congruent-native topology
  - Non congruent native topology
  - Non congruent tunnelled topology (CESNET!)
- Source Specific Multicast (SSM)
- Interdomain Multicast

RST-2214
11005_04_2005_c2

Deployment Scenarios (Cont.)

Cisco.com

- Single Domain Multicast – Congruent-native topology
  - Whole network is fully ipv6 unicast + multicast enabled
  - Normal unicast routing will take care of RPF check
  - Enable globally ipv6 multicast routing
  - Decide on Rendezvous Point distribution method:
    - Static
    - SSM – FF3x::/96
    - Embedded Rendezvous Point – FF70::/12
- Single Domain Multicast – Non congruent native topology
  - Separate unicast table for RPF check
  - Only static multicast or Multiprotocol-Border Gateway Protocol (MP-BGP) (SAFI=2) is possible

RST-2214
11005_04_2005_c2

Deployment Scenarios (Cont.)

Cisco.com

- Single Domain Multicast – Non congruent tunnelled topology
  - No support for IPv6 in some parts of the network
  - No support for IPv6 multicast in some parts of the network

RST-2214
11005_04_2005_c2

Deployment Scenarios (Cont.)

Cisco.com

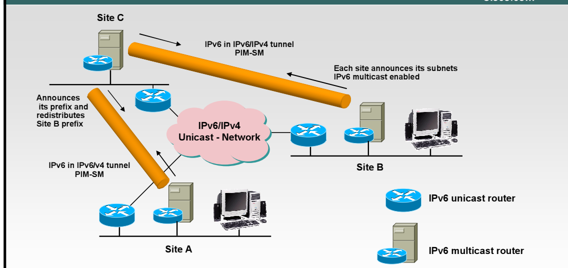

- Cloud can be IPv4 or IPv6

RST-2214
11005_04_2005_c2

Deployment Scenarios (Cont.)

Cisco.com

- Two possibilities:
  - Separate multicast routers are available:
    - Create unicast routing table for RPF check with dynamic routing protocol running over tunnels
    - Should not interfere with unicast routers
  - Same access router runs unicast and multicast:
    - Use static routing or MP-BGP (SAFI=2) for RPF check
    - See non congruent native topology scenario

RST-2214
11005_04_2005_c2

Deployment Scenarios (Cont.)

Cisco.com

- **Source Specific Multicast (SSM)**
  - RFC3569, subset of PIM-SM (RFC2362)
  - No Rendezvous Point or shared tree procedures (Soul Pattinson Telecommunications (SPT) only)
  - IGMPv3 (IPv4) or MLDv2 (IPv6) required
  - FF3x::/96

Deployment Scenarios (Cont.)

Cisco.com

- **Source Specific Multicast**

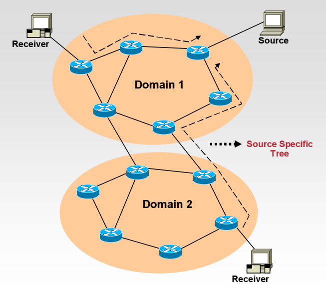

Deployment Scenarios (Cont.)

Cisco.com

- **Interdomain Multicast**
  - No equivalent for IPv4 MSDP (IETF refuses to discuss)
  - Administrative domains forced to share Rendezvous Point unless other mechanisms are used
  - IETF focuses on SSM to replace ASM
  - MP-BGP (SAFI=2)
    - Not strictly required when unicast and multicast topology are the same

Deployment Scenarios (Cont.)

Cisco.com

- **Interdomain Multicast**
  - Any Source Multicast (ASM)
  - Shared Rendezvous Point between PIM domains for particular multicast group range→ technical and administrative issues (no experience)

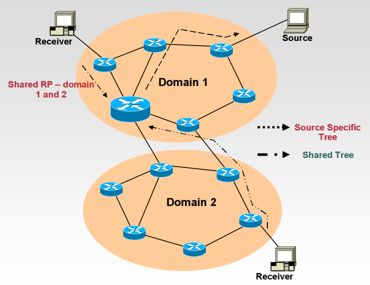

Deployment Scenarios (Cont.)

Cisco.com

- Interdomain Multicast
  - Embedded Rendezvous Point address
    - All network devices need to implement Embedded Rendezvous Point (different for SSM)
  - Shared static Rendezvous Point
    - Each domain have a set of Rendezvous Points to handle group ranges specific for local domain
    - May use BSR or embedded Rendezvous Point distribution
    - Common Rendezvous Point to handle groups with sources and receivers in all domains
    - Use static Rendezvous Point for interdomain shared Rendezvous Point
    - Position shared Rendezvous Point centrally

Deployment Scenarios (Cont.)

Cisco.com

- **Interdomain Multicast**
  - Shared BSR domain
  - All domains share same set of Rendezvous Points for all groups
  - Locally significant Rendezvous Point information only through static configuration

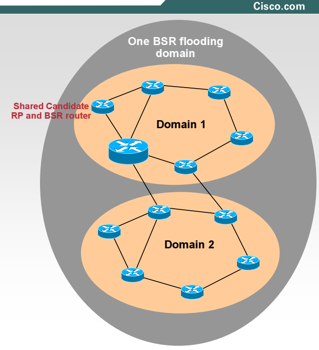

Deployment Scenarios (Cont.)

Cisco.com

- **Interdomain Multicast**
  - Constrained BSR domains
  - Each PIM domain is BSR domain
  - Each domain blocks BSR flooding/traffic to other domains
  - Domains need to agree on common set of Rendezvous Points and group ranges to be flooded identically in each domain
  - Requires certain level of administrative coordination, but decreases size of BSR domains

Deployment Scenarios (Cont.)

Cisco.com

- **Interdomain Multicast**
  - Constrained BSR domains

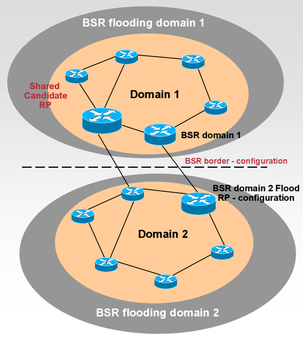

CISCO SYSTEMS

APPENDIX SLIDES FOR REFERENCE ONLY

IPv6 Support Outside of Cisco: Operating Systems

Cisco.com

Vendor	IPv6 Support	Versions	More Info
Microsoft	YES	XP (SP1) and 2003 CE .NET (Pocket PC 4.1)	http://www.microsoft.com/ipv6
Sun	YES	Solaris 8 and 9	http://wwws.sun.com/software/solaris/ipv6/
IBM	YES	z/OS Rel. 1.4 AIX 4.3 - > OS/390 V2R6 eNCS	http://www-3.ibm.com/software/os/zseries/ipv6/
BSD	YES	FreeBSD 4.0 - > OpenBSD 2.7 - > NetBSD 1.5 - > BSD/OS 4.2 - >	http://www.kame.net/
Linux	YES	RH 6.2 - > Mandrake 8.0 - > SuSE 7.1 - > Debian 2.2 - >	http://www.bieringer.de/linux/IPv6/status/IPv6+Linux-status-distributions.html
HP/Compaq	YES	HP-UX 11i Tru64 UNIX V5.1 OpenVMS V5.1	http://h18000.www1.hp.com/ipv6/next_gen.html
Novell	YES	Netware 6.1	http://www.novell.com/documentation/lg/nw65/index.html?page=/documentation/lg/nw65/readme/data/ajzlp6r.html
Apple	YES	MAC OS X 10.2 - >	http://developer.apple.com/macosx/

IPv6 OVER CLIENT VPN: REFERENCE SLIDES FOR NON-WINDOWS PLATFORMS

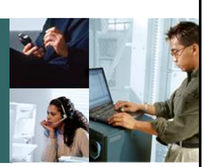

Router Configuration: Configured Tunnels

Cisco.com

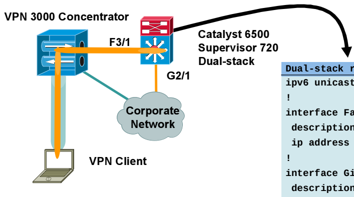


```
Dual-stack router configuration
ipv6 unicast-routing
!
interface FastEthernet3/1
 description TO VPN 3000
 ip address 20.1.1.1 255.255.255.0
!
interface GigabitEthernet2/1
 description TO Campus Network
 ipv6 address 2001:DB8:C003:111C::2/64
!
interface Tunnel1
 description Configured Tunnel for Client1
 no ip address
 ipv6 address 2001:DB8:C003:1123::1/64
 tunnel source FastEthernet3/1
 tunnel destination 10.1.99.103
 tunnel mode ipv6ip
```


Client Configuration (Windows XP): Configured Tunnels

Cisco.com

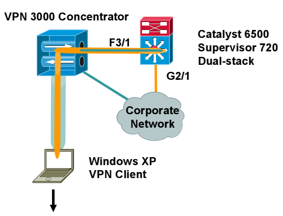

- Create v6v4tunnel
- Add IPv6 address to tunnel interface
- Create a default route (::/0) for the tunnel

VPN IP Router IP

```
netsh interface ipv6>add v6v4tunnel "CISCO" 10.1.99.103 20.1.1.1
Ok.
netsh interface ipv6>add address "CISCO" 2001:DB8:c003:1123::2
Ok.
netsh interface ipv6>add route ::/0 "CISCO"
Ok.
```


Does It Work?

Cisco.com

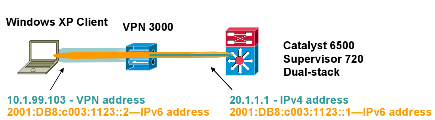


```
Interface 21: CISCO

Addr Type  DAD State  Valid Life   Pref. Life   Address
---------  ---------- ------------ ------------ -----------------------------
Manual     Preferred      infinite     infinite 2001:DB8:c003:1123::2
Link       Preferred      infinite     infinite fe80::a01:6368
```

```
netsh interface ipv6>show neighbors 21
Interface 2: Automatic Tunneling Pseudo-Interface

Internet Address                              Physical Address   Type
--------------------------------------------- -----------------  ----------
2001:DB8:c003:1123::1                             20.1.1.1        Permanent
fe80::1401:0101                                   20.1.1.1        Permanent
```

Client Configuration (Linux): ISATAP Tunnels

Cisco.com

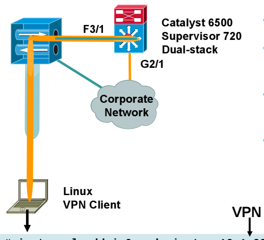

- IPv6-enabled
- Requires Kernel support for ISATAP—USAGI
- Modified iproute package—USAGI
- Must configure ISATAP router—NOT automatic

VPN IP   Router IP

```
# ip tunnel add is0 mode isatap 10.1.99.104 v4any 20.1.1.1 ttl 64
# ip link set is0 up
```

*See notes for full instructions for enabling IPv6 on Linux

Client Configuration (Sun Solaris): Configured Tunnels With 3002 Client

Cisco.com

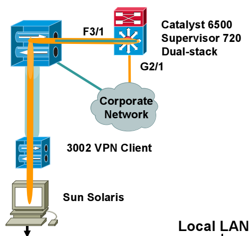

- IPv6-enabled
- Example of Solaris behind a 3002 VPN Client
- Basic configured tunnel—manual commands given
- Can maintain configuration permanently using /etc/hostname6.ip.tunN (where N is 0, 1, 2, and so on)

Local LAN IP   Router IP

```
# ifconfig ip.tun0 inet6 plumb
# ifconfig ip.tun0 inet6 tsrc 192.168.0.1 tdst 20.1.1.1 up
# ifconfig ip.tun0 inet6 addif 2001:DB8:c003:1123::2/64 2001:DB8:c003:1123::1 up
Created new logical interface ip.tun0:2
```

*See notes for full instructions for enabling IPv6 on Solaris

Client Configuration (Mac): Configured Tunnels With 3002 HW Client

Cisco.com

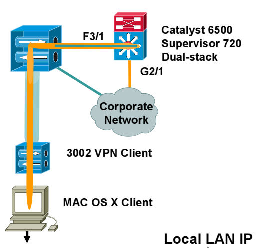

- IPv6-enabled
- Have permissions (root user)
- Example of Mac behind a 3002 VPN Client

Local LAN IP   Router IP

```
# ifconfig gif0 tunnel create
# ifconfig gif0 tunnel 192.168.0.1 20.1.1.1
# ifconfig gif0 inet6 alias 2001:DB8:c003:1123::2
# route add -inet6 default -interface gif0
```

OPERATING SYSTEM CONFIGURATION REFERENCE

Client Configuration (Windows XP): Dual-Stack

Cisco.com

- Required
  - Microsoft Windows XP (SP1 or higher)
  - Microsoft Windows Server 2003
- IPv6 must be installed
  - C:\>ipv6 install
- Have network (routers/switches) configured for IPv6
  - Stateless autoconfiguration and/or DHCPv6


```
C:\>ipconfig

Windows IP Configuration

Ethernet adapter Local Area Connection 1:
        Connection-specific DNS Suffix  . :
        IP Address. . . . . . . . . . . . : 10.1.1.100
        Subnet Mask . . . . . . . . . . . : 255.255.255.0
        IP Address. . . . . . . . . . . . : 2001:DB8:C003:1122:203:ffff:fe81:d6da
        IP Address. . . . . . . . . . . . : fe80::203:ffff:fe81:d6da%4
        Default Gateway . . . . . . . . . : 10.1.1.1
                                            fe80::201:42ff:fe2d:9580
```


LINUX

What Is Required

Cisco.com

- Red Hat 6.2 and higher
  - RH 8, 9, WS, and ES preferred
  - Fedora project builds
- Mandrake 8.0 and higher
- SuSE 7.1 and higher
- Debian 2.2 and higher
- ISATAP support requires
  - Requires Kernel support for ISATAP—USAGI
  - Modified iproute package—USAGI
  - http://www.linux-ipv6.org/

Client Configuration (Linux): Dual-Stack

Cisco.com

- ENABLE IPv6 support on Linux
  - Edit—/etc/sysconfig/network
  - Add entry—NETWORKING_IPV6=yes
  - Restart networking or reboot

```
# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:40:F4:6C:C8:AF
          inet addr:10.1.1.100  Bcast:10.1.1.255  Mask:255.255.255.0
          inet6 addr: 2001:DB8:C003:1122:240:f4ff:fe6c:c8af/64 Scope:Global
          inet6 addr: fe80::240:f4ff:fe6c:c8af/10 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:289223 errors:0 dropped:0 overruns:0 frame:0
          TX packets:13452 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          RX bytes:53425777 (50.9 Mb)  TX bytes:3381080 (3.2 Mb)
          Interrupt:5 Base address:0xf000
```


MICROSOFT

Client Configuration (Linux): ISATAP Tunnels

Cisco.com

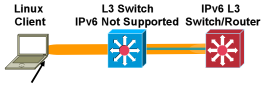

10.1.1.100—Client IPv4 address
2001:DB8:C003:111f:0:5efe:10.1.1.100—IPv6 address

- IPv6-enabled
- Requires Kernel support for ISATAP—USAGI
- Modified IProute package—USAGI
- Must configure ISATAP router—NOT automatic

Host IP — Router IP

```
# ip tunnel add is0 mode isatap 10.1.1.100 v4any 30.1.1.1 ttl 64
# ip link set is0 up
```


Client Configuration (Linux): Configured Tunnels

Cisco.com

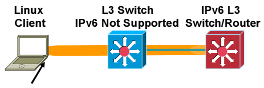

10.1.1.100—Client IPv4 address
2001:DB8:C003:1123::2—IPv6 address

- Create tunnel
- Enable the tunnel interface
- Add IPv6 address to tunnel interface
- Create a default route (::/0) for the tunnel

Router IP — Host IP

```
# ip tunnel add sit1 mode sit remote 30.1.1.1 local 10.1.1.100
# ip link set sit1 up
# ip address add dev sit1 2001:DB8:C003:1123::2/64
# ip route add ::/0 dev sit1
```


Does It Work?

Cisco.com

```
#ip tunnel show sit1
sit1: ipv6/ip  remote 30.1.1.1  local 10.1.1.100  ttl inherit

#route -A inet6 | grep sit1
Kernel IPv6 routing table
Destination                        Next Hop     Flags Metric Ref    Use Iface
2001:DB8:C003:1123::/64            ::           UA    256    10       0 sit1
fe80::/10                          ::           UA    256    6        0 sit1
ff02::9/128                        ff02::9      UAC   0      1        0 sit1
ff00::/8                           ::           UA    256    0        0 sit1
::/0                               ::           U     1024   0        0 sit1

# ip -6 addr show sit1
6: sit1@NONE: <POINTOPOINT,NOARP,UP> mtu 1480 qdisc noqueue
    inet6 fe80::a5e:a64d/128 scope link
    inet6 2001:DB8:C003:1123::2/64 scope global

#ping6 -I sit1 2001:DB8:C003:1123::1
PING 2001:DB8:C003:1123::1 from 2001:DB8:C003:1123::2 sit1:
64 bytes from 2001:DB8:C003:1123::1: icmp_seq=1 ttl=64 time=0.454 ms
64 bytes from 2001:DB8:C003:1123::1: icmp_seq=2 ttl=64 time=0.371 ms
64 bytes from 2001:DB8:C003:1123::1: icmp_seq=3 ttl=64 time=0.392 ms
64 bytes from 2001:DB8:C003:1123::1: icmp_seq=4 ttl=64 time=0.377 ms
```


MAC OS X

Client Configuration (Mac OS X 10.2 +): Dual-Stack via GUI

Cisco.com

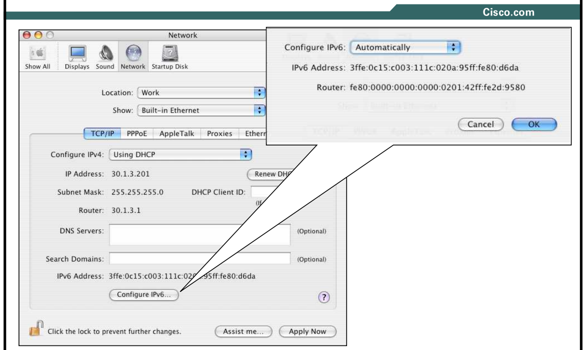

Client Configuration (Mac): Configured Tunnels

Cisco.com

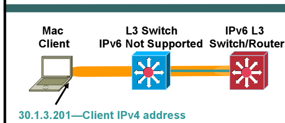

- Create tunnel interface
- Set tunnel end-points
- Add IPv6 address to tunnel
- Set default route

Local LAN IP   Router IP

```
# ifconfig gif0 tunnel create
# ifconfig gif0 tunnel 30.1.3.201 30.1.1.1
# ifconfig gif0 inet6 alias 2001:DB8:C003:1124::2
# route add -inet6 default -interface gif0
```

```
# ifconfig gif0

gif0: flags=8051<UP,POINTOPOINT,RUNNING,MULTICAST> mtu 1280
        tunnel inet 30.1.3.201 --> 30.1.1.1
        inet6 fe80::203:93ff:feee:9f1f prefixlen 64 scopeid 0x2
        inet6 2001:DB8:C003:1124::2 prefixlen 64
```

SUN SOLARIS

Things to Know

- Sun Solaris 8 and 9 will prompt for IPv6 activation during the installation process
  - Say yes and you will be ready for dual-stack with autoconfiguration
- You can also create the /etc/hostname6.<interface> file manually
  - For example if your physical Ethernet adapter is eri0 then you will find a /etc/hostname.eri0 file
  - You can create a /etc/hostname6.eri0 file manually or if you opted to have IPv6 support during installation then the file will already exist

```
#touch /etc/hostname6.eri0
reboot
ifconfig -a
```

and you will see a link local address on the interfaces

Client Configuration (Sun Solaris): Configured Tunnels

- Create tunnel interface
- Create tunnel end-points
- Add IPv6 address to interface
- Can maintain configuration permanently using /etc/hostname6.ip.tunN (where N is 0, 1, 2, and so on)

Local LAN IP   Router IP

```
# ifconfig ip.tun0 inet6 plumb
# ifconfig ip.tun0 inet6 tsrc 10.1.1.100 tdst 30.1.1.1 up
# ifconfig ip.tun0 inet6 addif 2001:DB8:C003:1123::2/64 2001:DB8:C003:1123::1 up
Created new logical interface ip.tun0:2
```

```
ip.tun0: flags=2200851<UP,POINTOPOINT,RUNNING,MULTICAST,NONUD,IPv6> mtu 1480 index 3
        inet tunnel src 10.1.1.100 tunnel dst 30.1.1.1
        tunnel hop limit 60
        inet6 fe80::4065:406a/10 --> fe80::a5e:a644
ip.tun0:1: flags=2200851<UP,POINTOPOINT,RUNNING,MULTICAST,NONUD,IPv6> mtu 1480 index 3
        inet6 2001:DB8:C003:1123::2/64 --> 2001:DB8:C003:1123::1
```

Links/Reference

- Microsoft
  - http://www.microsoft.com/ipv6: Main page
  - http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/windowsserver2003/proddocs/standard/sag_IP_v6topnode.asp: IPv6 documentation
- Linux
  - http://www.deepspace6.net/sections/docs.html: Linux IPv6 documentation
  - http://www.linux-ipv6.org/: Site for the USAGI port
  - http://v6web.litech.org/isatap: ISATAP for Linux site
- Mac
  - http://lists.apple.com/mailman/listinfo/ipv6: Mailing list for IPv6 on Mac
  - http://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/network-ipv6.html: IPv6 documentation that applies to Mac (Darwin) and general FreeBSD ports
- Sun
  - http://docs.sun.com/db/doc/817-0573: Sun Solaris IPv6 documentation
- Multi-OS tunnel configuration
  - http://www.join.uni-muenster.de/Dokumente/Howtos/index_howtos.php?lang=en: Great site for general IPv6 tunneling for most operation systems

Cisco Systems

networkers 2005

poweredbycisco.

DEPLOYING IPv6

Agenda

- **Enterprise Deployment**
  - Campus
  - WAN
  - S2S VPN
  - Remote Access
- Service Provider Deployment
  - Core
  - Access
- IPv6 Services
  - Multicast
  - QoS
  - Security
- Appendix—For Reference Only

Cisco.com

ENTERPRISE DEVELOPMENT

Start Here: Cisco IOS Software Release Specifics for IPv6 Features
http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/ipv6_c/ftipv6s.htm

IPv6 Coexistence in the Enterprise

Cisco.com

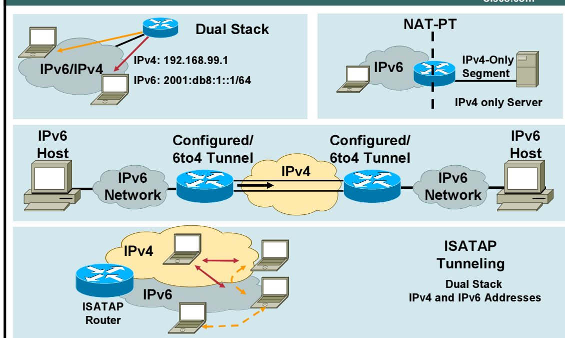

Cisco.com

ENTERPRISE DEPLOYMENT: CAMPUS

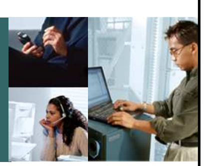

Campus IPv6 Deployment

Cisco.com

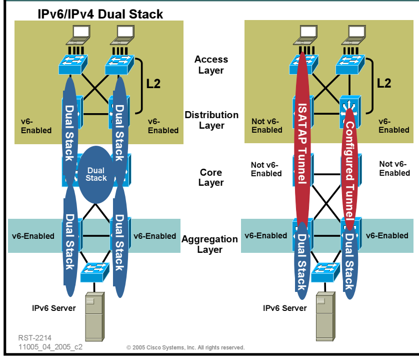

- Dual Stack and Tunnels in use depending on platform support
- Configured between IPv6-enabled L3-switches
- ISATAP between clients and a L3-switch
  - ISATAP – Intra-Site Automatic Tunnel Addressing Protocol

IPv6 on a Campus: Dual-Stack IPv4-IPv6

Cisco.com

- Requires switching/routing platforms to support hardware based forwarding for IPv4 and IPv6
- IPv6 is transparent on L2 switches except for multicast - MLD snooping
  - IPv6 management—Telnet/SSH/HTTP/SNMP
- Requires robust control plane for both IPv4 and IPv6
  - Variety of routing protocols—The same ones in use today with IPv4
- IPv6 multicast, QoS, infrastructure security, etc…
- IPv4 and IPv6 control planes and data planes must not impact each other

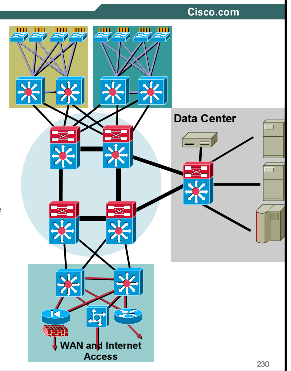

Distribution Layer: Dual Stack

Cisco.com

```
ipv6 unicast-routing
ipv6 multicast-routing
ipv6 cef
!
interface GigabitEthernet1/1
 description To 6k-core-right
 ipv6 address 2001:DB8:C003:1105::1/127
 ipv6 ospf 1 area 0
 ipv6 ospf hello-interval 1
 ipv6 ospf dead-interval 3
 ipv6 cef
!
interface GigabitEthernet1/2
 description To 6k-core-left
 ipv6 address 2001:DB8:C003:1106::3/127
 ipv6 ospf 1 area 0
 ipv6 ospf hello-interval 1
 ipv6 ospf dead-interval 3
 ipv6 cef
```

```
interface Vlan10
 description Data VLAN for Access
 ipv6 address 2001:DB8:C003:1102::1/64
 ipv6 nd prefix 2001:DB8:C003:1102::/64
86400 86400
 ipv6 nd reachable-time 5000
 ipv6 ospf 1 area 1
 ipv6 cef
!
ipv6 router ospf 1
 router-id 10.122.0.25
 log-adjacency-changes
 passive-interface Vlan10
 timers spf 1 1
```

- Optional: lower valid/preferred lifetimes from defaults (2592000/604800)—in seconds to match DHCPv4 lease times
- Optional: lower Neighbor Unreachability Detection (NUD) from 30 seconds (faster failover until HSRP is available)

IPv6 Campus ISATAP Configuration

Cisco.com

- ISATAP connections look like one flat network
- Create DNS "A" record for "ISATAP" = 10.1.1.1
- Use Static Config if DNS use is not desired:
  `C:\>netsh interface ipv6 isatap set router 10.1.1.1`
- Currently ISATAP does not support multicast!!

No Configuration Change on Non-v6 Switches

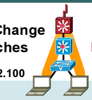


```
ipv6 unicast-routing
ipv6 cef
!
interface Loopback0
 description ISATAP address for Access Layer
 ip address 10.1.1.1 255.255.255.255
!
interface GigabitEthernet2/10
 ipv6 address 2001:DB8:C003:111C::2/64
 ipv6 cef
!
interface Tunnel0
 ipv6 address 2001:DB8:C003:111F::/64 eui-64
 no ipv6 nd suppress-ra
 ipv6 cef
 tunnel source Loopback0
 tunnel mode ipv6ip isatap
```

ISATAP Address Format:

64-bit Unicast Prefix	0000:5EFE:	IPv4 Address
	32-bit	32-bit
	Interface ID	

2001:DB8:C003:111F:0:5EFE:10.1.2.100

Cisco.com

ENTERPRISE DEPLOYMENT: WAN

WAN Deployment

Cisco.com

- Cisco WAN routers support IPv6
- Dual-stack is recommended due to ease of deployment, security advantage and performance
- Support for every media/WAN type you want to use (Frame Relay, leased-line, broadband, MPLS, etc…)

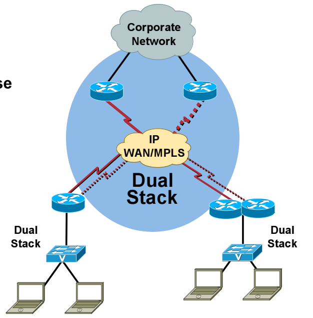

HSRP for IPv6

Cisco.com

- Basically the same as HSRP for IPv4
- Changes occur in Neighbor Advertisement, Router Advertisement, and ICMPv6 redirects
- Virtual MAC derived from HSRP group number and virtual IPv6 Link-local address
- IPv6 Virtual MAC range:
  - 0005.73A0.0000 - 0005.73A0.0FFF (4096 addresses)
- HSRP IPv6 UDP Port Number 2029 (IANA Assigned)
- No HSRP IPv6 secondary address
- No HSRP IPv6 specific debug

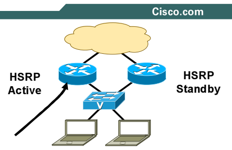


```
interface FastEthernet0/1
 ipv6 address 2001:66:67::2/64
 ipv6 cef
 standby version 2
 standby 1 ipv6 autoconfig
 standby 1 timers msec 250 msec 800
 standby 1 preempt
 standby 1 preempt delay minimum 180
 standby 1 authentication md5 key-string cisco
 standby 1 track FastEthernet0/0
```

Host with GW of Virtual IP

```
#route -A inet6 | grep ::/0 | grep eth2
::/0      fe80::207:85ff:fef3:2f60          UGDA  1024   3         0 eth2
::/0      fe80::205:9bff:febf:5ce0          UGDA  1024   0         0 eth2
::/0      fe80::5:73ff:fea0:1               UGDA  1024   0         0 eth2
```


Cisco.com

OTHER TRANSITION TYPES

Configured Tunnel
Building the Tunnel

Cisco.com

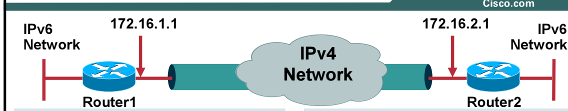


```
ipv6 unicast-routing
ipv6 cef
!
interface Tunnel0
 no ip address
 ipv6 cef
 ipv6 address 2001:DB8:C003:3::1/64
 ipv6 rip v6 enable
 tunnel source FastEthernet0/1
 tunnel destination 172.16.2.1
 tunnel mode ipv6ip
!
interface FastEthernet0/0
 ipv6 address 2001:DB8:C003:1::1/64
 ipv6 cef
!
interface FastEthernet0/1
 ip address 172.16.1.1 255.255.255.252
```

```
ipv6 unicast-routing
ipv6 cef
!
interface Tunnel0
 no ip address
 ipv6 cef
 ipv6 address 2001:DB8:C003:3::2/64
 ipv6 rip v6 enable
 tunnel source FastEthernet0/1
 tunnel destination 172.16.1.1
 tunnel mode ipv6ip
!
interface FastEthernet0/0
 ipv6 address 2001:DB8:C003:2::1/64
 ipv6 cef
!
interface FastEthernet0/1
 ip address 172.16.2.1 255.255.255.252
```

Internet Key Exchange (IKE) Policy
Configured Tunnel (Static Maps)

Cisco.com

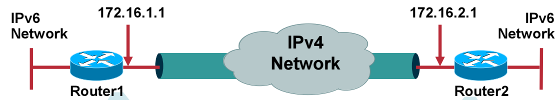


```
crypto isakmp policy 1
   encr 3des
   authentication pre-share
   group 2
!
crypto isakmp key CISCO address
   172.16.2.1
```

```
crypto isakmp policy 1
   encr 3des
   authentication pre-share
   group 2
!
crypto isakmp key CISCO address
   172.16.1.1
```


IPSec Policy
Configured Tunnel (Static Maps)

Cisco.com

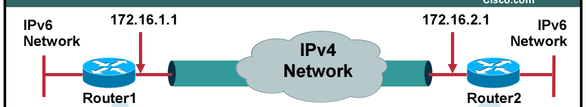


```
crypto ipsec transform-set STRONG esp-
3des esp-sha-hmac
!
crypto map STATIC-MAP local-address
FastEthernet0/1
!
crypto map STATIC-MAP 1 ipsec-isakmp
 set peer 172.16.2.1
 set transform-set STRONG
 set pfs group2
 match address VPN-TO-R2
!
ip access-list extended VPN-TO-R2
 permit 41 host 172.16.1.1 host
172.16.2.1
```

```
crypto ipsec transform-set STRONG esp-
3des esp-sha-hmac
!
crypto map STATIC-MAP local-address
FastEthernet0/1
!
crypto map STATIC-MAP 1 ipsec-isakmp
 set peer 172.16.1.1
 set transform-set STRONG
 set pfs group2
 match address VPN-TO-R1
!
ip access-list extended VPN-TO-R1
 permit 41 host 172.16.2.1 host
172.16.1.1
```


Apply VPN Configuration
Configured Tunnel

Cisco.com

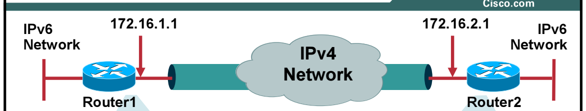


```
interface Tunnel0
 no ip address
 ipv6 address 2001:DB8:C003:3::1/64
 ipv6 cef
 ipv6 mtu 1400
 ipv6 rip V6 enable
 tunnel source FastEthernet0/1
 tunnel destination 172.16.2.1
 tunnel mode ipv6ip
!
interface FastEthernet0/1
 ip address 172.16.1.1 255.255.255.252
 crypto map STATIC-MAP
!
ip route 172.16.2.1 255.255.255.255
172.16.1.2
```

```
interface Tunnel0
 no ip address
 ipv6 address 2001:DB8:C003:3::2/64
 ipv6 cef
 ipv6 mtu 1400
 ipv6 rip V6 enable
 tunnel source FastEthernet0/1
 tunnel destination 172.16.1.1
 tunnel mode ipv6ip
!
interface FastEthernet0/1
 ip address 172.16.2.1 255.255.255.252
 crypto map STATIC-MAP
!
ip route 172.16.1.1 255.255.255.255
172.16.2.2
```


IPv6 IPSec Example (FUTURE-EFT)
IKE/IPSec Policies

Cisco.com

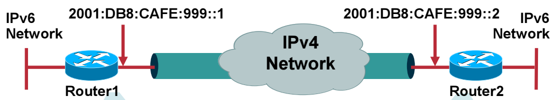

Router1:

```
crypto isakmp policy 1
 authentication pre-share
crypto isakmp key CISCOKEY address ipv6
2001:DB8:CAFE:999::2/128
crypto isakmp keepalive 30 30
!
crypto ipsec transform-set v6STRONG esp-
3des esp-sha-hmac
!
crypto ipsec profile v6PRO
 set transform-set v6STRONG
```

Router2:

```
crypto isakmp policy 1
 authentication pre-share
crypto isakmp key CISCOKEY address ipv6
2001:DB8:CAFE:999::1/128
crypto isakmp keepalive 30 30
!
crypto ipsec transform-set v6STRONG esp-
3des esp-sha-hmac
!
crypto ipsec profile v6PRO
 set transform-set v6STRONG
```

RST-2214
11005_04_2005_c2 241

IPv6 IPSec Example (FUTURE-EFT)
Tunnels

Cisco.com

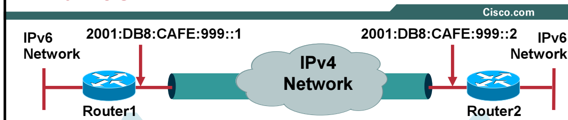

Router1:

```
interface Tunnel0
 ipv6 address 2001:DB8:CAFE:F00D::/64
eui-64
 ipv6 ospf 1 area 0
 tunnel source Serial2/0
 tunnel destination 2001:DB8:CAFE:999::2
 tunnel mode ipsec ipv6
 tunnel protection ipsec profile v6PRO
!
interface Ethernet0/0
 ipv6 address 2001:DB8:CAFE:100::1/64
 ipv6 ospf 1 area 1
!
interface Serial2/0
 ipv6 address 2001:DB8:CAFE:999::1/64
```

Router2:

```
interface Tunnel0
 ipv6 address 2001:DB8:CAFE:F00D::/64
eui-64
 ipv6 ospf 1 area 0
 tunnel source Serial2/0
 tunnel destination 2001:DB8:CAFE:999::1
 tunnel mode ipsec ipv6
 tunnel protection ipsec profile v6PRO
!
interface Ethernet0/0
 ipv6 address 2001:DB8:CAFE:200::1/64
 ipv6 ospf 1 area 2
!
interface Serial2/0
 ipv6 address 2001:DB8:CAFE:999::2/64
```

11005_04_2005_c2 242

IPv6 IPSec Example
Show Output

Cisco.com

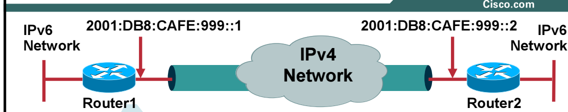


```
Router1#show crypto engine connections active
Crypto Engine Connections
   ID Intfc  Type  Algorithm           Encrypt  Decrypt IP-Address
    3 Tu0    IPsec 3DES+SHA                  0       17 2001:DB8:CAFE:999::1
    4 Tu0    IPsec 3DES+SHA                 16        0 2001:DB8:CAFE:999::1
 1006 Tu0    IKE   SHA+DES                   0        0 2001:DB8:CAFE:999::1

Router1#show crypto sessions
Crypto session current status
Interface: Tunnel0
Session status: UP-ACTIVE
Peer: 2001:DB8:CAFE:999::2 port 500
  IKE SA: local 2001:DB8:CAFE:999::1/500
          remote 2001:DB8:CAFE:999::2/500 Active
  IPSEC FLOW: permit 41 ::/0 ::/0
        Active SAs: 2, origin: crypto map
```

RST-2214
11005_04_2005_c2 243

Configured Tunnels vs. Automatic Tunnels

Cisco.com

	Configured	ISATAP	6to4
Manual Configuration per Client (Router-Side)	YES	NO	NO
Manual Configuration per Client (Client-Side)	YES	NO	NO
IPv6 Multicast Support	YES	NO	NO
Broad Client OS Support	YES	NO	YES
Optimal for Remote Access Clients	NO	YES	YES

*GRE must be used if ISIS is used as the routing protocol

Legacy Services (IPv4 Only)

Cisco.com

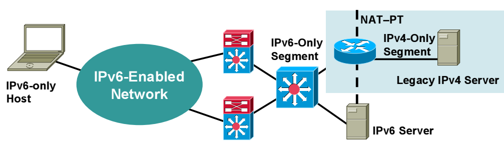

- Many of the non-routing/switching products do not yet support IPv6 (i.e., content switching modules)
- NAT-PT (Network Address Translation–Protocol Translation) as an option to front-end IPv4-only server—NOTE: NAT-PT IS BEING MOVED TO EXPERIMENTAL
- Place NAT-PT box as close to IPv4 only server as possible
- Be VERY aware of performance and manageability issues

Configuring Cisco IOS NAT-PT

Cisco.com

- NAT-PT enables communication between IPv6-only and IPv4-only nodes
- CEF switching in 12.3(14)T

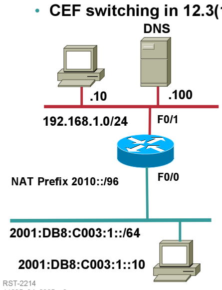


```
interface FastEthernet0/0
   ipv6 address 2001:DB8:C003:1::1/64
   ipv6 cef
   ipv6 nat
!
interface FastEthernet0/1
   ip address 192.168.1.1 255.255.255.0
   ipv6 nat prefix 2010::/96
   ipv6 nat
!
ipv6 nat v4v6 source 192.168.1.100 2010::100
!
ipv6 nat v6v4 source route-map MAP1 pool V4POOL
ipv6 nat v6v4 pool V4POOL 192.168.2.1
192.168.2.10 prefix-length 24
!
route-map MAP1 permit 10
 match interface FastEthernet0/1
```


ENTERPRISE DEPLOYMENT: REMOTE ACCESS

IPv6 for Remote Devices

- Remote nodes can use a VPN client or router to establish connectivity back to enterprise
- Possible over IPv4 today, not possible over IPv6 today (key management is still in progress)
- How could we allow access to IPv6 services at central site or Internet in a secure fashion?
  - Enabling IPv6 traffic inside the Cisco VPN client tunnel
  - Allow remote host to establish a v6-in-v4 tunnel either automatically or manually
    - ISATAP—Intra Site Automatic Tunnel Addressing Protocol
    - Configured—Static configuration for each side of tunnel
  - Same split-tunneling issues exists

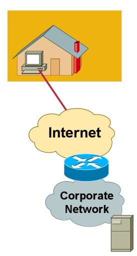

IPv6-in-IPv4 Tunnel Example

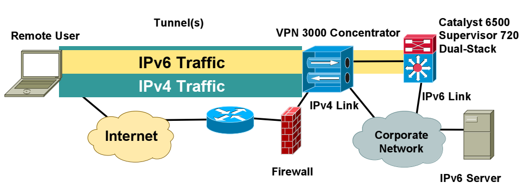

IPsec VPN
IPv6-in-IPv4 Tunnel

Note: The VPN Concentrator could be replaced with a VPN-enabled Cisco IOS Router or PIX®

Considerations

Cisco.com

- Cisco IOS® version supporting IPv6 configured/ ISATAP tunnels
  - Configured—12.3(1)M/12.3(2)T/12.2(14)S and above (12.4M/12.4T)
  - ISATAP—12.3(1)M, 12.3(2)T, 12.2(14)S and above (12.4M/12.4T)
  - Catalyst® 6500 with Sup720—12.2(17a)SX1—HW forwarding
- Be aware of the security issues if split-tunneling is used
  - Attacker can come in IPv6 interface and jump on the IPv4 interface (encrypted to enterprise)
- Remember that the IPv6 tunneled traffic is still encapsulated as a tunnel WHEN it leaves the VPN device
- Allow IPv6 tunneled traffic across any access lists (Protocol 41)

Required Stuff: Client Side

Cisco.com

- Client operating system with IPv6
  - Microsoft Windows XP SP1 (Supports Configured/ISATAP)
  - Linux (7.3 or higher)—USAGI port required for ISATAP
  - Mac OS X (10.2 or higher)—Currently need a VPN device on client network
  - SunOS (8 or higher)—Currently need a VPN device on client network
  - See reference slide for links/OS listing
- Cisco VPN Client 4.0.1 and higher for configured/ISATAP

- Cisco VPN Client 3.x for configured ONLY
- Cisco HW VPN Client 3002—recommended for Mac/Sun clients until virtual adapter support is available

IPv6 Using Cisco VPN Client

Example: Client Configuration (Windows XP): ISATAP

Cisco.com

- Microsoft Windows XP (SP1 or higher)
- IPv6 must be installed
  - `C:\>ipv6 install`
- XP will automatically attempt to resolve the name "ISATAP"
  - Local host name
  - Hosts file—SystemRoot\system32\drivers\etc
  - DNS name query
  - NetBIOS and Lmhosts
- Manual ISATAP router entry can be made
  - `netsh interface ipv6 isatap set router 20.1.1.1`
- Key fact here is that NO additional configuration on the client is needed again!
- USE PREVIOUS ISATAP CONFIGURATIONS SHOWN FOR ROUTER-SIDE

Note: ISATAP is supported on some versions of Linux/BSD (manual router entry is required)

Does It Work?

Cisco.com

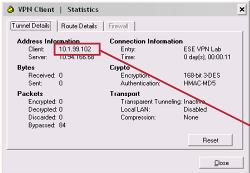

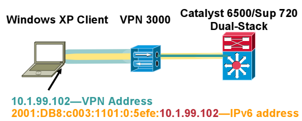


```
Interface 2: Automatic Tunneling Pseudo-Interface

Addr Type  DAD State  Valid Life   Pref. Life   Address
---------  ---------- ------------ ------------ -----------------------------
Public     Preferred   29d23h56m5s  6d23h56m5s 2001:db8:c003:1101:0:5efe:10.1.99.102
Link       Preferred   infinite     infinite   fe80::5efe:10.1.99.102
```

```
netsh interface ipv6>show route
Querying active state...

Publish  Type       Met  Prefix                    Idx  Gateway/Interface Name
-------  --------  ----  ------------------------  ---  ---------------------
no       Autoconf     9  2001:db8:c003:1101::/64     2  Automatic Tunneling Pseudo-Interface
no       Manual       1  ::/0                        2  fe80::5efe:20.1.1.1
```


Cisco.com

SERVICE PROVIDER DEPLOYMENT

Start Here: Cisco IOS Software Release Specifics for IPv6 Features
http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/ipv6_c/ftipv6s.htm

IPv6 in the SP: What Does It Do for Me?

Cisco.com

- Benefits for the ISP (short term):
  - Expanded private use address pool for internal devices
  - Ability to acquire large enough address blocks to avoid impeding rollout/subscriber-growth business plans
  - Not lose existing or new customers due to lack of support
- Benefits for the ISP (long term):
  - Reduction in ‘application failure’ related support calls caused by IPv4/NAT
  - Ability to remove customer-managed infrastructure component (NAT) from the path, improving application support
  - Ability to deploy new service offerings into the home without dealing with translation issues and address constraints

Today's Network Infrastructure

Cisco.com

- Service Providers core infrastructure are basically following two paths
  - MPLS with its associated services
    - MPLS/VPN, L2 services over MPLS, QoS,…
  - Native IPv4 core with associated services
    - L2TPv3, QoS, Multicast,…
- IP services portfolio—Access
  - Enterprise: Lease lines
  - Home Users/SOHO: ADSL, FTTH, Dial
  - Data Center: Web hosting, servers,…
- Next step—The integration of IPv6 services

Note: Don't classify IPv6 tunneled traffic as "undetermined" (Protocol 41)

Dual-Stack IPv4-IPv6

Cisco.com

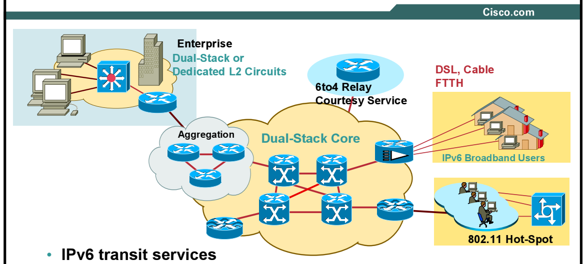

- IPv6 transit services
- IPv6 enabled on Core routers
- Enterprise and consumer IPv6 access
- Additional services
  - IPv6 multicast for streaming

IPv6 over MPLS

Cisco.com

- Many service providers have already deployed MPLS in their IPv4 backbone for various reasons
- MPLS can be used to facilitate IPv6 integration
- Multiple approaches for IPv6 over MPLS:
  - IPv6 over L2TPv3
  - IPv6 over EoMPLS/AToM
  - IPv6 CE-to-CE IPv6 over IPv4 Tunnels
  - IPv6 Provider Edge Router (6PE) over MPLS
  - IPv6 VPN Provider Edge (6VPE) over MPLS
  - Native IPv6 over MPLS

IPv6 QoS

IPv6 QoS: Header Fields

Cisco.com

- **IPv6 traffic class**
  - Exactly the same as TOS field in IPv4
- **IPv6 flow label (RFC 3697)**
  - A new 20-bit field in the IPv6 basic header which:
    - Labels packets belonging to particular flows
    - Can be used for special sender requests
  - Per RFC, Flow Label must not be modified by intermediate routers

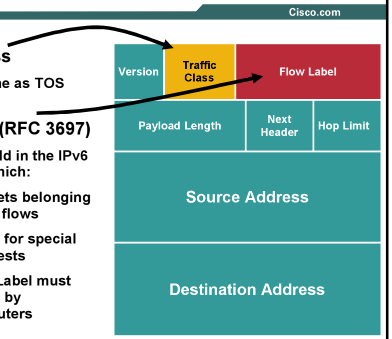

 © 2005 Cisco Systems, Inc. All rights reserved.

IPv6 QoS Syntax Changes

Cisco.com

- **IPv4 syntax has used “ip” following match/set statements**
  - Example: `match ip dscp, set ip dscp`
- **Modification in QoS syntax to support IPv6 and IPv4**
  - New `match` criteria
    - `match dscp`—Match DSCP in v4/v6
    - `match precedence`—Match Precedence in v4/v6
    - `match protocol ipv6`—Match on IPv6 Protocol
  - New `set` criteria
    - `set dscp`—Set DSCP in v4/v6
    - `set precedence`—Set Precedence in v4/v6
- **Additional support for IPv6 does not always require new Command Line Interface (CLI)**
  - Example—WRED

 © 2005 Cisco Systems, Inc. All rights reserved.

Simple QoS Example: IPv4 and IPv6

Cisco.com

```
class-map match-any BRANCH-BULK-DATA
 match access-group name BULK-DATA-IPV6
 match access-group name BULK-DATA
class-map match-all BULK-DATA
 match  dscp af11
!
policy-map RBR-WAN-EDGE
 class BULK-DATA
  bandwidth percent 4
  random-detect
!
policy-map RBR-LAN-EDGE-IN
 class BRANCH-BULK-DATA
  set dscp af11
!
ip access-list extended BULK-DATA
 permit tcp any any eq ftp
 permit tcp any any eq ftp-data
!
ipv6 access-list BULK-DATA-IPV6
 permit tcp any any eq ftp
 permit tcp any any eq ftp-data
```

ACL Match To Set DSCP (If Packets Are Not Already Marked)

`service-policy input RBR-LAN-EDGE-IN`

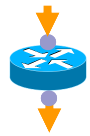

`service-policy output RBR-WAN-EDGE`

ACLs to Match for Both IPv4 and IPv6 Packets

IPv6 QoS: Support

Cisco.com

- Cisco's current IPv6 QoS implementation supports:
  - Packet classification
  - Queuing—(does support LLQ)—excluding PQ/CQ
  - Traffic shaping
  - WRED
  - Class-based packet marking
  - Policy-based packet marking
- Cisco's current IPv6 QoS implementation does not support:
  - Compressed Real-Time Protocol (CRTP)
  - Network-Based Application Recognition (NBAR)
  - Committed Access Rate (CAR)
  - Priority Queuing (PQ)
  - Custom Queuing (CQ)

Cisco.com

CISCO SYSTEMS

IPv6 MANAGEMENT

IETF IPv6 MIBs

Cisco.com

- Standardization status at IETF:
  - At the beginning:
    - IPv4 and IPv6 MIBs dissociated
  - Today:
    - Unified MIBs are on standardization track

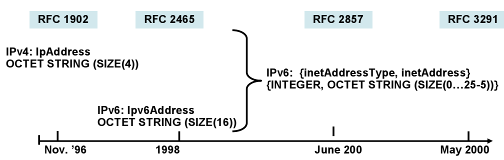

IETF IPv6 MIB History

Cisco.com

- Initial IPv6 MIBs were developed for IPv6 only

RFC 2465	IPv6 MIB
RFC 2466	ICMPv6
RFC 2452	IPv6 TCP MIB
RFC 2454	IPv6 UDP MIB

The IETF intends to deprecate the old MIBs when the "new" MIBs are complete

IETF statement: The IPV6-TCP-MIB defined in RFC 2452 (and others) has been moved to Historic since the approach of having separate IP version specific tables is not followed anymore. Implementation of RFC 2452 is thus not suggested anymore.

IETF MIB Update Status

Cisco.com

- draft-ietf-ipv6-rfc2011-update-10.txt (IP MIB)
  - Obsoletes RFCs 2011, 2465 and 2466
  - Proposed standard; in the RFC Editor's queue
- draft-ietf-ipv6-rfc2096-update-07.txt (IP Forwarding Table MIB)
  - Proposed standard; in the RFC Editor's queue
- RFC 4022 (TCP MIB)
- draft-ietf-ipv6-rfc2013-update-04.txt (UDP MIB)
  - proposed standard; in the RFC Editor's queue

IETF MIB Update Status

Cisco.com

- draft-ietf-ipv6-inet-tunnel-mib-03.txt (IP Tunnel MIB)
- draft-ietf-idr-bgp4-mibv2-05.txt (BGP MIB)
  - Document is now obsolete; was different from draft-ietf-idr-bgp4-mib-14.txt, which is just an update (mainly cleanup) to RFC 1657
- draft-ietf-ospf-ospfv3-mib-08.txt (OSPFv3 MIB)
  - On-going
- MLDv1 MIB (RFC 3019)
  - But MLDv2 (RFC 3810) is now published and implemented

Cisco IOS IPv6 MIBs

Cisco.com

Cisco IOS IPv6 MIB Implementation

- Cisco IOS implementation—Cisco-IETF-IP and Cisco-IETF-IP-Forwarding MIB—is compliant with -00 of these new drafts
- Cisco IOS 12.2(14)S, 12.0(22)S, and 12.2(15)T/12.3(1)M/12.3(2)T
- No support of Interface Stats table—IPv4/IPv6 traffic stats at the interface level
- IPv4/IPv6 stats can be monitored from CLI “show interface accounting” on most platforms
- Plan is to update these MIBs when finally published
- Interface Stats table will be added

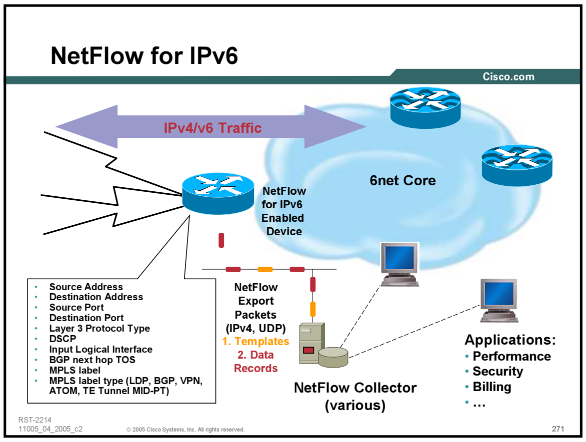
NetFlow for IPv6
Cisco.com
IPv4/v6 Traffic
6net Core
NetFlow for IPv6 Enabled Device
Source Address
Destination Address
Source Port
Destination Port
Layer 3 Protocol Type
DSCP
Input Logical Interface
BGP next hop TOS
MPLS label
MPLS label type (LDP, BGP, VPN, ATOM, TE Tunnel MID-PT)
NetFlow Export Packets (IPv4, UDP)
1. Templates
2. Data Records
NetFlow Collector (various)
Applications:
Performance
Security
Billing
…
RST-2214
11005_04_2005_c2
271

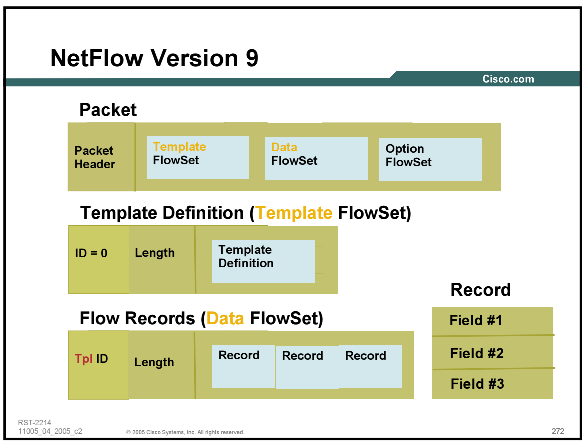
NetFlow Version 9
Cisco.com
Packet
Packet Header
Template FlowSet
Data FlowSet
Option FlowSet
Template Definition (Template FlowSet)
ID = 0
Length
Template Definition
Record
Field #1
Field #2
Field #3
Flow Records (Data FlowSet)
Tpl ID
Length
Record
Record
Record
RST-2214
11005_04_2005_c2
272

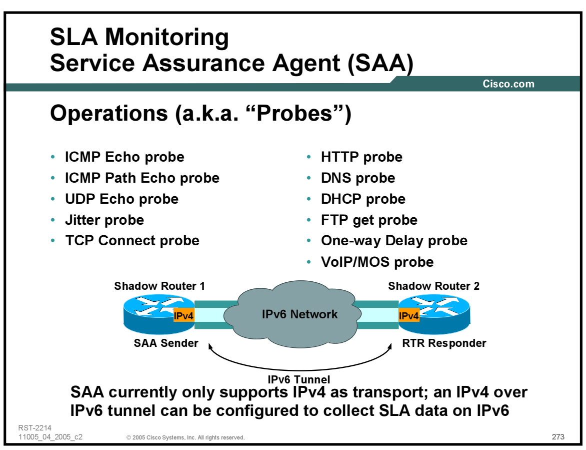
SLA Monitoring
Service Assurance Agent (SAA)
Cisco.com
Operations (a.k.a. "Probes")
ICMP Echo probe
ICMP Path Echo probe
UDP Echo probe
Jitter probe
TCP Connect probe
HTTP probe
DNS probe
DHCP probe
FTP get probe
One-way Delay probe
VoIP/MOS probe
Shadow Router 1
Shadow Router 2
IPv4
IPv6 Network
IPv4
SAA Sender
RTR Responder
IPv6 Tunnel
SAA currently only supports IPv4 as transport; an IPv4 over IPv6 tunnel can be configured to collect SLA data on IPv6
RST-2214
11005_04_2005_c2
273

Management Platforms

Cisco.com

- Commercial Network Management platforms versus GPL or Home made tools
  - HP-OV proposes a version with IPv6 features: NNM 7.0 (September 2003)
  - Netview (IBM) doesn't propose any IPv6 features?
  - Tivoli: no information …
  - Infovista: « no IPv6 plan at the moment »
  - MRTG IPv6 capable

Cisco IPv6 NMS Applications Preview

Cisco.com

- LNS 2.5—CiscoWorks adds IPv6 support
  - CiscoView
  - PathTrace
  - User Tracking
- Netflow Collector 5.0

Configuration Management Diff Viewer

Cisco.com

Compare archive versions with IPv6 address and commands

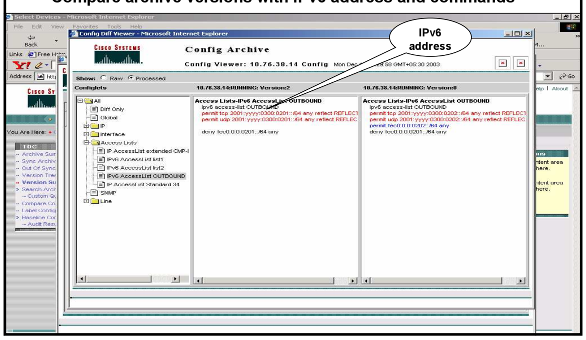

IPv6 Management: Traffic Monitor

Cisco.com

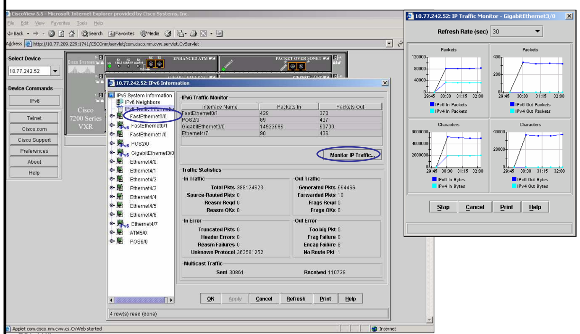

Highlighting IPv6 Devices and Links

Cisco.com

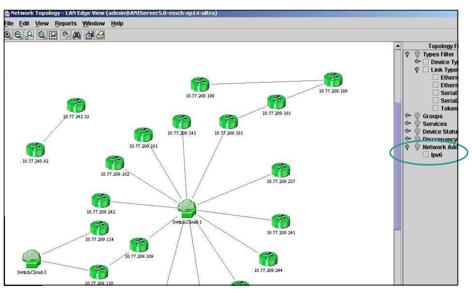

IPv6 Path Trace

Cisco.com

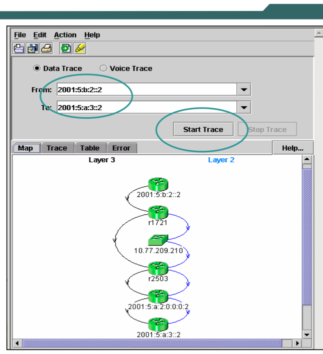

User Tracking End Host Report

Cisco.com

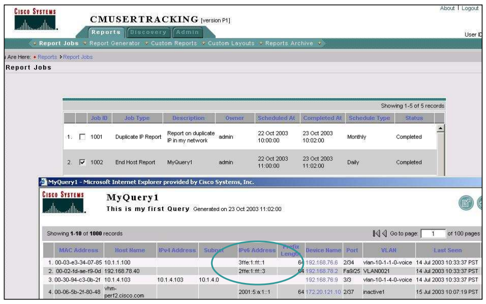

Cisco.com

IPv6 SECURITY
SEE SEC-2003

www.cisco.com/security_services/ciag/documents/v6-v4-threats.pdf

IPv6 Security

Cisco.com

- RFC “mandates” privacy and encryption
- Same IPSec you already know
- Two security extension headers defined; all implementations required to support (IPSec)
  - Authentication Header (AH)
  - Encapsulating Security Payload (ESP)
  - Key distribution protocols are under development
  - Support for manual key configuration required
- IPv6 Security is MORE THAN IPSec!
- New concept of privacy addressing
  - On by default in Microsoft XP SP1+
  - Randomly generated address
- Nearly impossible to perform successful network scans

IPv6 Protocol Challenges

Cisco.com

- Inherits many challenges found in IPv4
  - Same applications
  - Same TCP, UDP layers
- Many new features
  - Autoconfiguration (router advertisements)
  - ND—Neighbor Discovery (altering ICMPv6 packets)
  - DAD—Multiple (bad) addresses
  - Mobile IPv6—binding update, etc.

RST-2214
11005_04_2005_c2 283

IPv6 Security Considerations

Cisco.com

- If all hosts are performing encryption, what happens to…
  - Intrusion detection
  - Intrusion prevention (inline filtering)
  - Virus protection
  - Deep packet inspection
  - Proxies
- The real world will likely implement…
  - Decoupling of end to end encryption (terminate connections on a bulk encryption device)
  - Use of authentication headers providing packet integrity, but not encryption
  - Extensive use of personal (host-based) firewalls and host-based IDS (Cisco Security Agent) to augment network-based security tools

RST-2214
11005_04_2005_c2 284

IPv6 Transition Mechanism Challenges

Cisco.com

- Dual stack
  - Consider security for both protocols
  - Cross v4/v6 abuse
  - Resiliency (shared resources)
- Tunnels
  - Bypass firewalls (protocol 41)
  - Relayed DoS attacks from v6 to v4 and vice versa
- Translation mechanisms
  - Prevent end-to-end network and transport layer security

RST-2214
11005_04_2005_c2 285

Basic IPv6 Packet Filtering
(Access Control List)

Cisco.com

When Used for Traffic Filtering, IPv6 Access Control Lists (ACL) Offers the Same Level of Support as in IPv4

- Every IPv6 ACL has `implicit permit icmp any any nd-na` and permit icmp any any nd-ns
- Implicit deny `all` at the end of access list

```
interface FastEthernet0/0
 ipv6 address 2001:DB8:C003:1101::1/64
 ipv6 traffic-filter V6FILTER in
!
ipv6 access-list V6FILTER
 permit tcp any host 2001:DB8:C003:1102::10 eq web
!
```

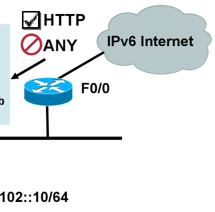

RST-2214
11005_04_2005_c2

Cisco IOS IPv6 Firewall Feature Set
Example: Nothing New from IPv4

Cisco.com

- Cisco IOS Firewall released 12.3(7)T

```
ipv6 unicast-routing
ipv6 cef
!
ipv6 inspect audit-trail
ipv6 inspect max-incomplete low 150
ipv6 inspect max-incomplete high 250
ipv6 inspect one-minute low 100
ipv6 inspect one-minute high 200
ipv6 inspect name V6FW tcp timeout 300
ipv6 inspect name V6FW udp
ipv6 inspect name V6FW icmp
!
interface FastEthernet0/0
ipv6 address 2001:DB8:C003:1112::2/64
ipv6 cef
ipv6 traffic-filter EXAMPLE in
ipv6 inspect V6FW in
!
ipv6 access-list EXAMPLE
permit tcp any host 2001:DB8:C003:1113::2 eq www
permit tcp any host 2001:DB8:C003:1113::2 eq ftp
deny ipv6 any any log
```

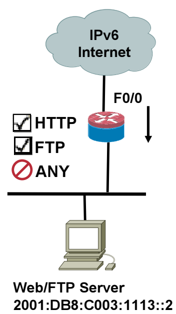

http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/ps5761/index.html

RST-2214
11005_04_2005_c2

PIX 7.0: ACL
Very Similar to Cisco IOS

Cisco.com

```
interface Ethernet0
 nameif outside
 ipv6 address 2001:db8:c000:1051::37/64
 ipv6 enable
interface Ethernet1
 nameif inside
 ipv6 address 2001:db8:c000:1052::1/64
 ipv6 enable

ipv6 unicast-routing

ipv6 route outside ::/0 2001:db8:c000:1051::1

ipv6 access-list SECURE permit tcp any host 2001:db8:c000:1052::7 eq telnet
ipv6 access-list SECURE permit icmp6 any 2001:db8:c000:1052::/64

access-group SECURE in interface outside
```

RST-2214
11005_04_2005_c2

11005_04_2005_c2.scr

PIX 7.0 and Stateful Inspection

Cisco.com

```
pixA# show conn
4 in use, 7 most used
ICMP out fe80::206:d7ff:fe80:2340:0 in fe80::209:43ff:fea4:dd07:0
  idle 0:00:00 bytes 16
UDP out 2001:db8:c000:1051::138:53 in 2001:db8:c000:1052::7:50118
  idle 0:00:02 flags -
TCP out 2001:200:0:8002:203:47ff:fea5:3085:80 in
  2001:db8:c000:1052::7:11009 idle 0:00:14 bytes 8975 flags UfFRIO
TCP out 2001:db8:c000:1051::1:11008 in 2001:db8:c000:1052::7:23
  idle 0:00:04 bytes 411 flags UIOB
```


Conclusion

Cisco.com

- **Start now rather than later**
  - **Purchase for the future**
  - **Start moving legacy application towards IPv6 support**
- **Integration can be done per Application (Dual Stack or Tunneled)**
- **Things to consider:**
  - **EIGRP for IPv6 (EFT)**
  - **IPv6 HSRP (EFT)**
  - **IPv6 IPSec (EFT)**
  - **DNS**
  - **Enterprise products/features—(Voice, CDN, Advanced Security)**
  - **Full-scale management of IPv6—Appendix Section**
  - **ISP multihoming solutions (Multi6 WG)—"Goals for IPv6 Site-Multihoming Architectures" (RFC 3582)—http://www.ietf.org/html.charters/multi6-charter.html**
- **Enterprise and SP Deployment Scenarios**
  - **<draft-ietf-v6ops-ent-scenarios-05.txt>**
  - **<draft-ietf-v6ops-bb-deployment-scenarios-01.txt>**
  - **RFC4029—Scenarios and Analysis for Introducing IPv6 into ISP Networks**