Change Log (Hidden Slide)

Release	Date	Owner	Changes
2.7	30-May-2011	Jeff Apcar	Change log added. IPv4 address exhaustion slides removed and migrated to a separate deck to avoid change control problems.
3.0	06-Jun-2011	Jeff Apcar	Sourced from Original Technical Introduction - individual sections broken out into separate presentations covering Routing Protocols., DHCP etc…Renamed to Technical Overview to distinguish it from the original "full" version. Added slide on why 128 bits was chosen for address size
3.1	29-Jun-2011	Jeff Apcar	Internal draft
3.2	29-Jun-2011	Jeff Apcar	Added technical review questions. Removed DNS and DHCP sections to different Module
3.3	08-Jul-2011	Stephan Millet	Peer Review prior to printing
3.4	15-Jul-2011	Jeff Apcar	Updated agenda to reflect contents
3.5	18-Jul-2011	Jeff Apcar	Updated Interface ID slide with "EUI-64 Format" Added Flags to table in Multicast Address slide
3.6	29-Jul-2011	Jeff Apcar	Errors corrected from review by Claud Nelson Mendonca

IPv6 Technical Overview

Version 3.6

Module Objectives

- Provide foundational knowledge required to support IPv6 technologies:
  - Explanation of various IPv6 addresses and packet header
  - Explain the function of ICMPv6 and Neighbour Discovery
  - Explain basic IPv6 interface configuration
- Be able to converse about IPv6 concepts
- Provide the ground work for further learning specific to your technology area

Agenda

- IPv6 Addressing
- IPv6 Header
- ICMPv6 and Neighbor Discovery
- IPv6 Interface Configuration
- Module Review Questions

IPv4/IPv6 Technology Comparison

Service	IPv4	IPv6
Addressing Range	32-bit, NAT	128-bit, Multiple Scopes
IP Provisioning	DHCP	SLAAC, Renumbering, DHCP
Security	IPSec	IPSec
Mobility	Mobile IP	Mobile IP with Direct Routing
Quality-of-Service	Differentiated Service, Integrated Service	Differentiated Service, Integrated Service
Multicast	IGMP/PIM/MBGP	MLD/PIM/MBGP, Scope Identifier

IPv6 Addressing

IPv6 Addresses

- IPv6 addresses are 128 bits long
  - Segmented into 8 groups of four HEX characters (called HEXtets)
  - Separated by a colon (:)
  - Default is 50% for network ID, 50% for interface ID
  - Network portion is allocated by Internet registries 2^64 (1.8 x 1019)
  - Still leaves us with ~ 3 billion network prefixes for each person on earth

Global Unicast Identifier Example

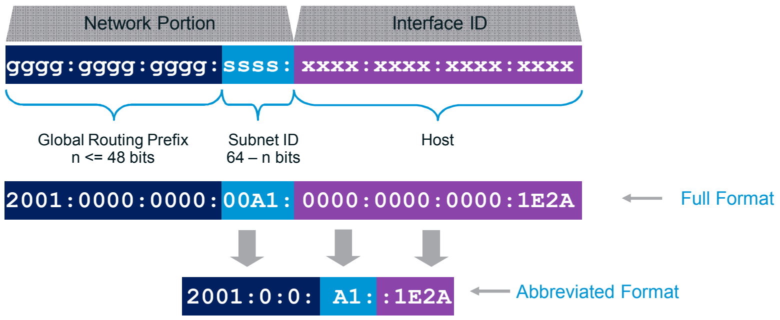

So How Big Is The IPv6 Address Space?

Antares
15th Brightest
star in the sky

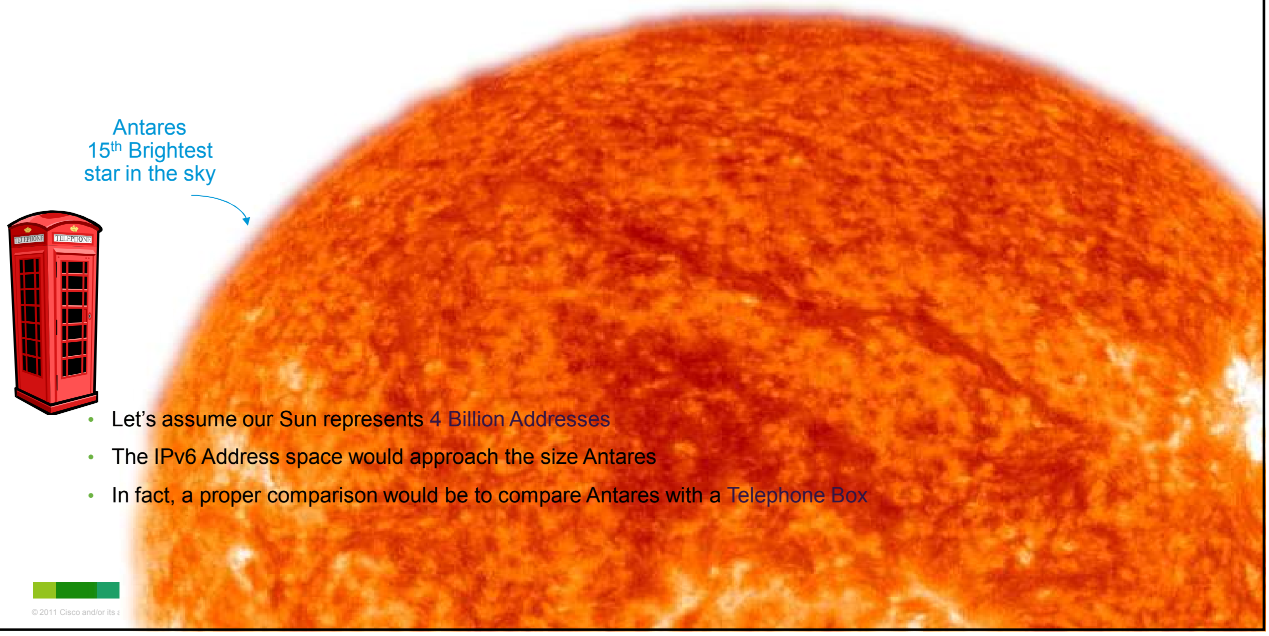

- Let's assume our Sun represents 4 Billion Addresses
- The IPv6 Address space would approach the size Antares
- In fact, a proper comparison would be to compare Antares with a Telephone Box

How Was The IPv6 Address Size Chosen?

- Some wanted fixed-length, 64-bit addresses
  - Easily good for 1012 sites, 1015 nodes, at .0001 allocation efficiency (3 orders of magnitude more than IPv6 requirement)
  - Minimizes growth of per-packet header overhead
  - Efficient for software processing
- Some wanted variable-length, up to 160 bits
  - Compatible with OSI NSAP addressing plans
  - Big enough for auto-configuration using IEEE 802 addresses
  - Could start with addresses shorter than 64 bits & grow later
- Settled on fixed-length, 128-bit addresses

IPv6 Address Format Details

- Hex numbers are not case sensitive
- Abbreviations are possible
  - Leading zeros in contiguous block could be represented by (::)
    - **2001:0db8:0000:130F:0000:0000:087C:140B**
    - **2001:0db8:0:130F::87C:140B**
  - Double colon can only appear once in the address
- IPv6 uses CIDR representation
  - IPv4 address looks like 98.10.0.0/16
  - IPv6 address is represented the same way 2001:db8:12::/48
- Only leading zeros are omitted, trailing zeros cannot be omitted
  - **2001:0db8:0012::/48 = 2001:db8:12::/48**
  - **2001:db80:1200::/48 ≠ 2001:db8:12::/48**

Comparison of Address Type Syntax (CATS)

IPv4 Address	IPv6 Address	Network Operator
		
HOST: 156.50.20.1	FE80::20B:60FF:FEA7:D81A	I need some Catnip…
NET: 156.50.0.0/16	2001:0:0:A1::/64	HEX is a curse…

IPv6 Address Representation

- Loopback address representation
  - `0:0:0:0:0:0:0:1 == ::1`
  - Same as 127.0.0.1 in IPv4
  - Identifies self
- Unspecified address representation
  - `0:0:0:0:0:0:0:0 == ::`
  - Used as a placeholder when no address available
  - (Initial DHCP request, Duplicate Address Detection DAD)
  - NOT the default route
- Default Route representation
  - `::/0`

IPv6 Address Scopes

- Addresses are assigned to interfaces
- An IPv6 interface is "expected" to have multiple addresses and multiple scopes
- Addresses have scope
  - Link Local
  - Unique Local
  - Global
- Addresses have lifetime
  - Valid and preferred lifetime

IPv6 Address Types

★

- Three types of unicast address scopes
  - Link-Local – Non routable exists on single layer 2 domain **(FE80::/64)**

    `FE80:0000:0000:0000: xxxx:xxxx:xxxx:xxxx`

  - Unique-Local – Routable within administrative domain **(FC00::/7)**

    `FCgg:gggg:gggg:ssss: xxxx:xxxx:xxxx:xxxx`

    `FDgg:gggg:gggg:ssss: xxxx:xxxx:xxxx:xxxx`

  - Global – Routable across the Internet **(2000::/3)**

    `2ggg:gggg:gggg:ssss: xxxx:xxxx:xxxx:xxxx`

    `3ggg:gggg:gggg:ssss: xxxx:xxxx:xxxx:xxxx`

- Multicast addresses **(FF00::/8)**

  `FFzs: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx`

  Flags (z) in 3rd nibble (4 bits) Scope (s) into 4th nibble

IPv6 Addressing Types

- Represented in Binary and Hex

Type	Binary	Hex
Global Unicast Address	001	2 or 3
Link Local Unicast Address	1111 1110 10	FE80::/10
Unique Local Unicast Address	1111 1100 1111 1101	FC00::/7 FC00::/8 (registry) FD00::/8 (no registry)
Multicast Address	1111 1111	FF00::/16
Solicited Node Multicast		FF02::1:FF00/104

Interface Address Set

- An interface can have many addresses allocated to it

Address Type	Requirement	Comment
Link Local	Required	Required on all interfaces
Unique Local	Optional	Valid only within an Administrative Domain
Global Unicast	Optional	Globally routed prefix
Auto-Config 6to4	Optional	Used for 2002:: 6to4 tunnelling
Solicited Node Multicast	Required	Neighbour Discovery and Duplicate Detection (DAD)
All Nodes Multicast	Required	For ICMPv6 messages

Global Unicast Addresses

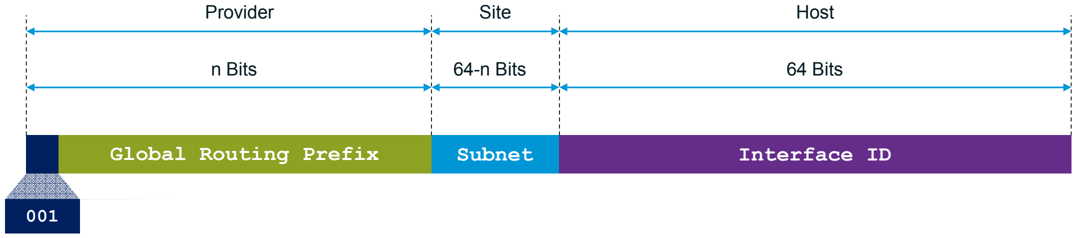

- Addresses for generic use of IPv6
- Structured as a hierarchy to try and keep the aggregation

PI and PA Allocation Process

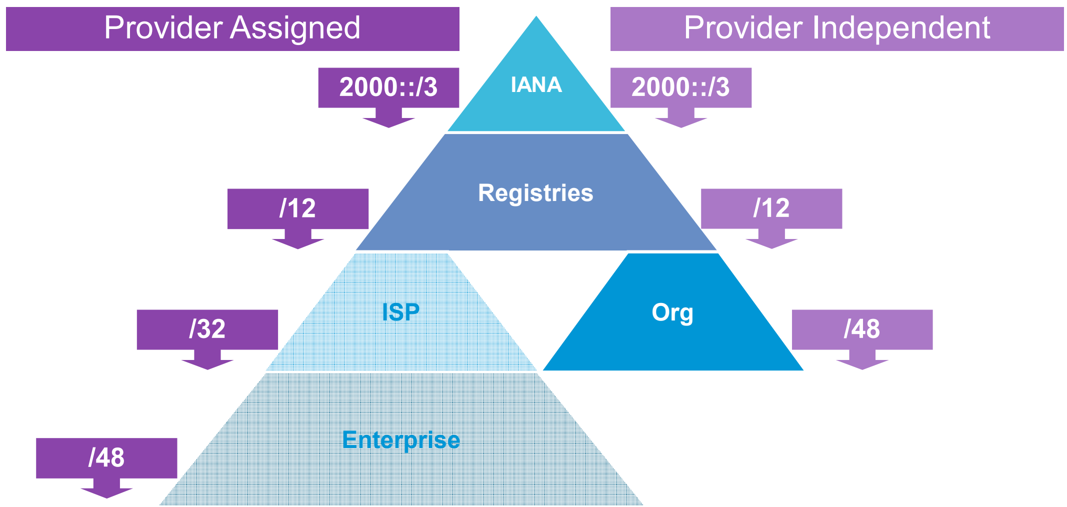

Global Unicast Address Interface ID

- Interface ID unicast address may be assigned in different ways
  - Auto-configured from a 64-bit EUI-64 or expanded from a 48-bit MAC
  - Auto-generated pseudo-random number (to address privacy concerns)
  - Assigned via DHCP
  - Manually configured
- EUI-64 format to do stateless auto-configuration
  - Expands the 48 bit MAC address to 64 bits by inserting FFFE into the middle
  - To ensure 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

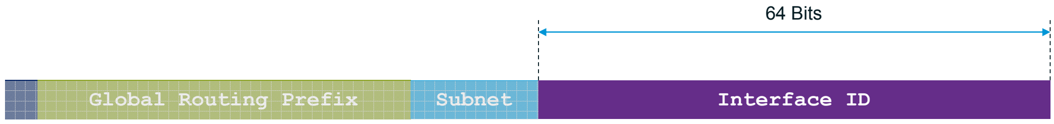

 19

IPv6 Interface Identifier (EUI-64 format)

- Cisco uses the EUI-64 format to do stateless auto-configuration
- 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
- Cisco devices ‘bit-flip’ the 7th bit

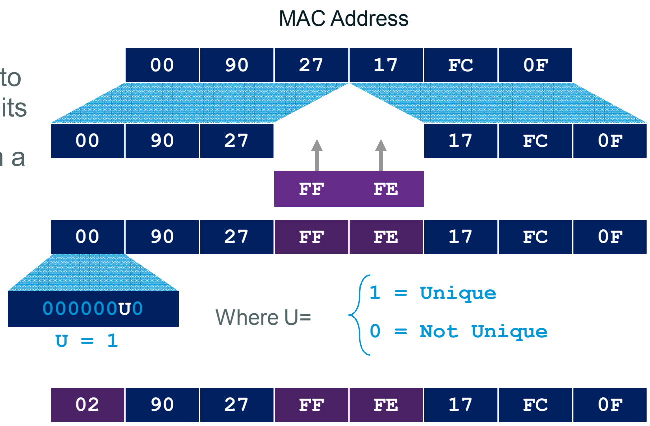

 20

IPv6 over Ethernet

- IPv6 uses Ethernet Protocol ID (0x86DD)

- IPv4 uses Ethernet Protocol ID (0x0800)

Link-Local Address

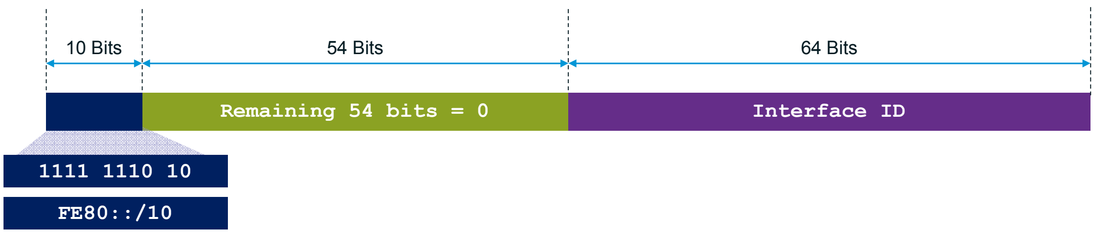

- Mandatory for communication between two IPv6 devices
- Automatically assigned by device using EUI-64
- Also used for next-hop calculation in routing protocols
- Only link specific scope
- Remaining 54 bits could be zero or any manually configured

Unique Local Address (RFC 4193)

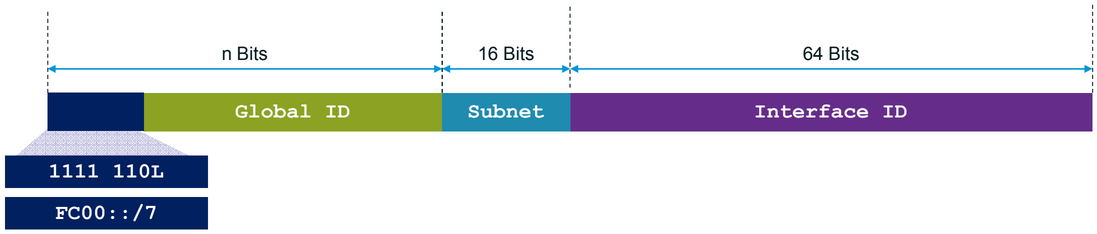

- ULA are "like" RFC 1918 – not routable on Internet
- ULA uses include
  - Local communications
  - Inter-site VPNs (Mergers and Acquisitions)
- FC00::/8 is Registry Assigned (L bit = 0), FD00::/8 is self generated (L bit = 1)
  - Registries not yet assigning ULA space, http://www.sixxs.net/tools/grh/ula/
- Global ID can be generated using an algorithm
  - Low order 40 bits result of SHA-1 Digest {EUI-64 && Time}

IPv6 Multicast Address (RFC 4291)

- An IPv6 multicast address has the prefix FF00::/8 (1111 1111)
  - Second octet defines lifetime and scope

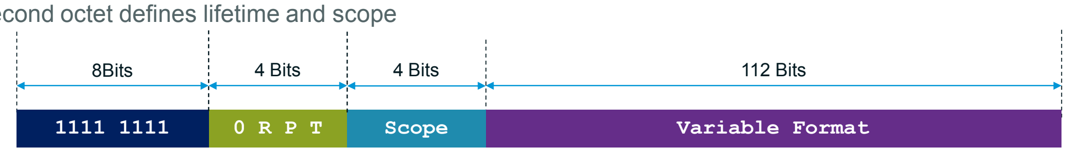

Flags	
R = 0 R = 1	No embedded RP Embedded RP
P = 0 P = 1	Not based on unicast Based on unicast
T = 0 T = 1	Permanent address (IANA assigned) Temporary address (local assigned)

Scope	
1	Node
2	Link
3	Subnet
4	Admin
5	Site
8	Organization
E	Global

IPv6 Multicast Address – Unicast Based (RFC 3306)

- Every Unicast prefix comes with 2^32 multicast group addresses

8Bits	4 Bits	4 Bits	8 Bits	8 Bits	64 Bits	32 Bits
1111 1111	0011	Scope	Reserved	Len	Unicast Prefix	Group ID

Example	
Flags	P=1 (Unicast), T=1 (Temp)
Prefix	2001:db8:cafe:1::
Scope	E (Global – 0011 in binary)
Length	64 bits (0x40)
Group ID	11ff:11ee

ff3e:40:2001:db8:cafe:1:11ff:11ee

Well Known Multicast Addresses

Address	Scope	Meaning
FF01::1	Node-Local	All Nodes
FF01::2	Node-Local	All Routers
FF02::1	Link-Local	All Nodes
FF02::2	Link-Local	All Routers
FF02::5	Link-Local	OSPFv3 Routers
FF02::6	Link-Local	OSPFv3 DR Routers
FF02::1:FFXX:XXXX	Link-Local	Solicited-Node

- "02" means that this is a permanent address (t = 0) and has link scope (2)
- http://www.iana.org/assignments/ipv6-multicast-addresses

Solicited-Node Multicast Address

- For each Unicast and Anycast address configured there is a corresponding solicited-node multicast (Layer 3 address)
- Used in neighbor solicitation (NS) messages
- Multicast address with a link-local scope
- Solicited-node multicast consists of

  FF02::1:FF & {lower 24 bits from IPv6 Unicast interface ID}

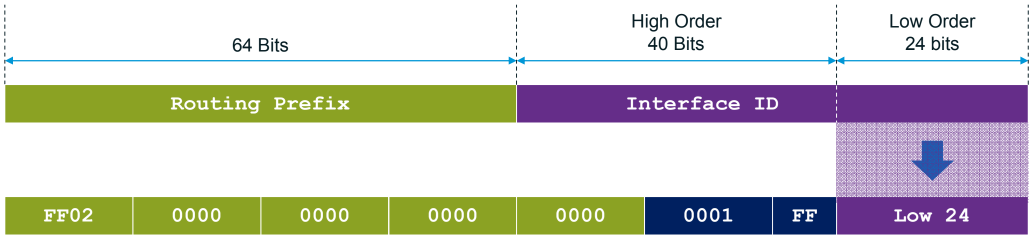

Multicast Mapping over Ethernet (RFC 2464)

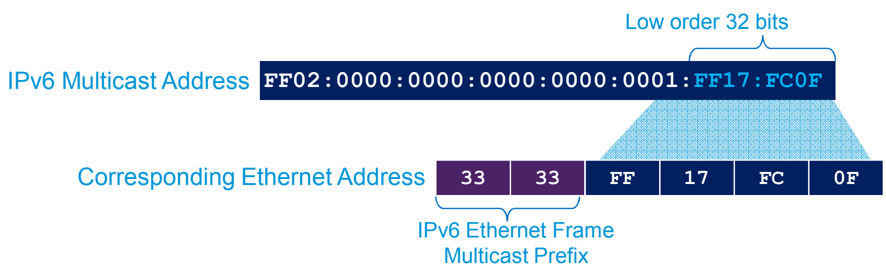

- IPv6 multicast address to Ethernet mapping

  33:33:{Low Order 32 bits of the IPv6 multicast address}

Solicited Node Multicast Address Example

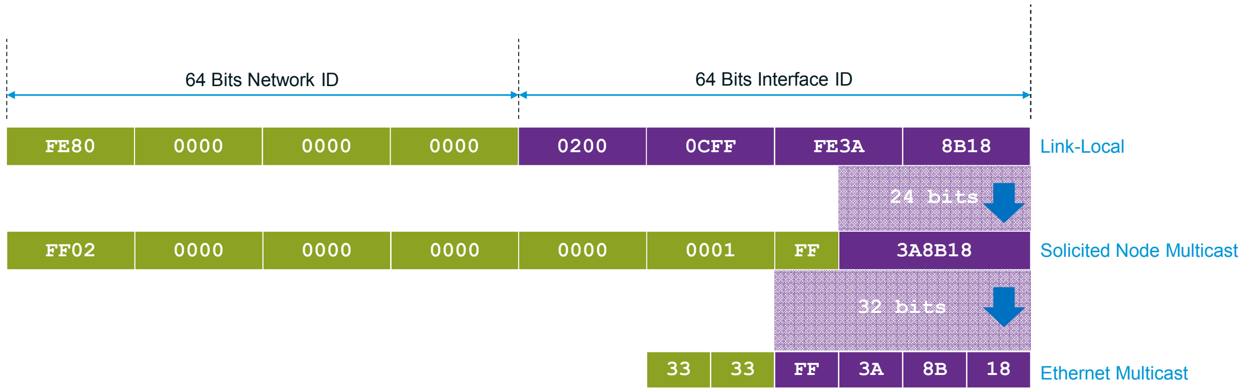

IPv6 Interface Example

```
R1#show ipv6 interface e0
Ethernet0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::200:CFF:FE3A:8B18
  No global unicast address is configured
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF3A:8B18
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
R1#
```

Link-local address (FE80::)
All Nodes
All Routers
Solicited Node Multicast Address

IPv6 Header

IPv4 and IPv6 Header Comparison

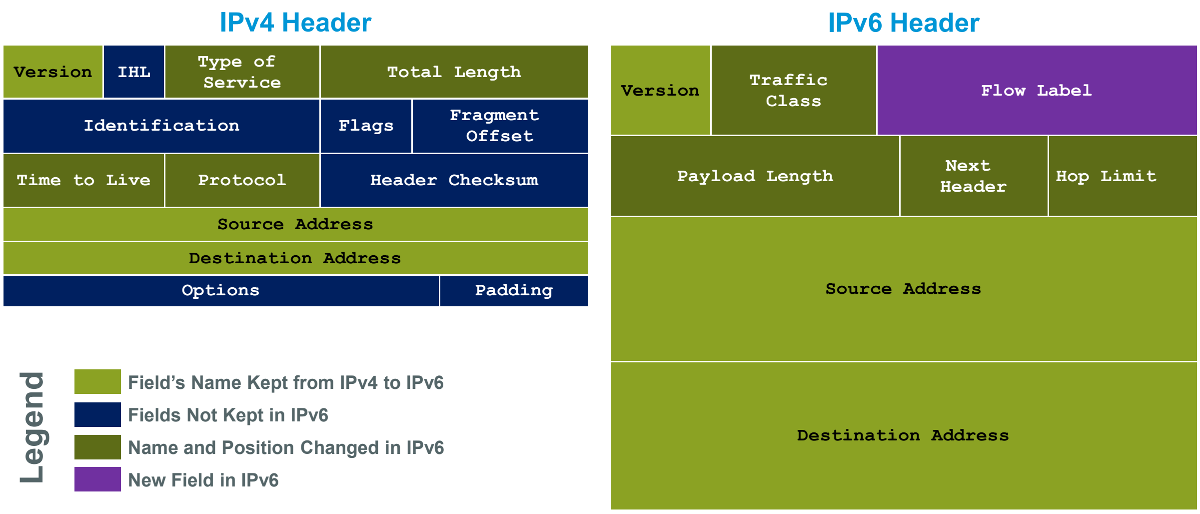

Extension Headers

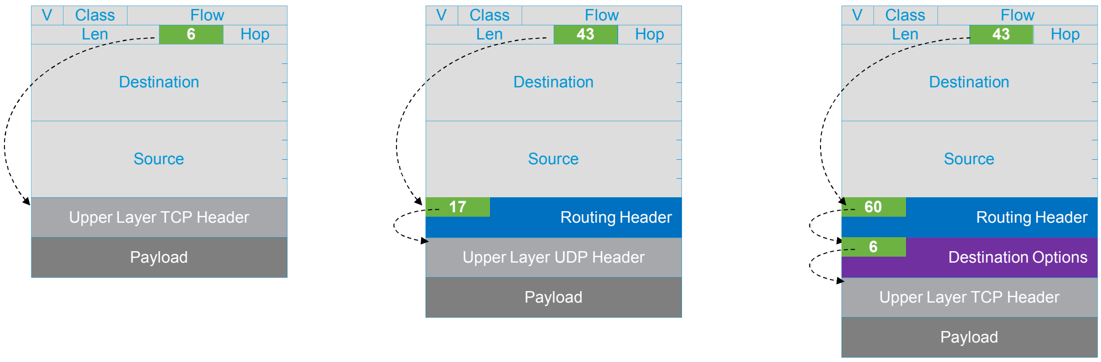

- Extension Headers Are Daisy Chained

Extension Header Order

- Extension headers must be in the following sequence

Order	Header Type	Header Code
1	Basic IPv6 Header	-
2	Hop-by-Hop Options	0
3	Dest Options (with Routing options)	60
4	Routing Header	43
5	Fragment Header	44
6	Authentication Header	51
7	ESP Header	50
8	Destination Options	60
9	Mobility Header	135
-	No Next Header	59
Upper Layer	TCP	6
Upper Layer	UDP	17
Upper Layer	ICMPv6	58

Fragmentation and Path MTU Discovery

Fragmentation in IPv6

- The original large packet consists of two parts
- Unfragmentable part
  - IPv6 header plus any headers that must be processed by the nodes en-route
  - Unfragmentable part is repeated with fragments appended to it following the "fragment header"
- Fragmentable part
  - The headers that need to be processed only by the destination node = the end-to-end headers + upper layer header and data
  - Fragmentable part is divided into pieces with length multiple of 8 octets
- RFC 2460 Section 4.5 defines the fragmentation header
- Minimum MTU for IPv6 is 1280 bytes
  - All links MUST support it

Fragment Header

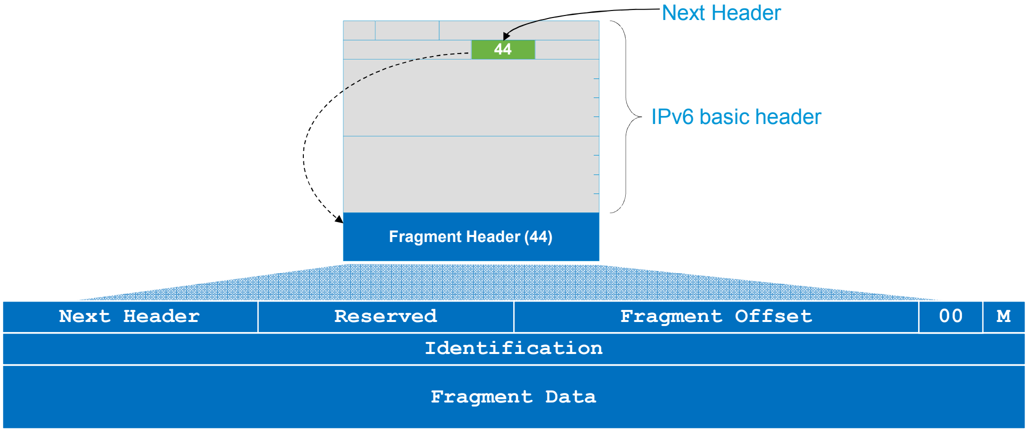

- Fragmentation is left to end devices in IPv6
  - Routers do not perform fragmentation
- Fragment header used when an end node has to send a packet larger than the path MTU

Fragmentation Header

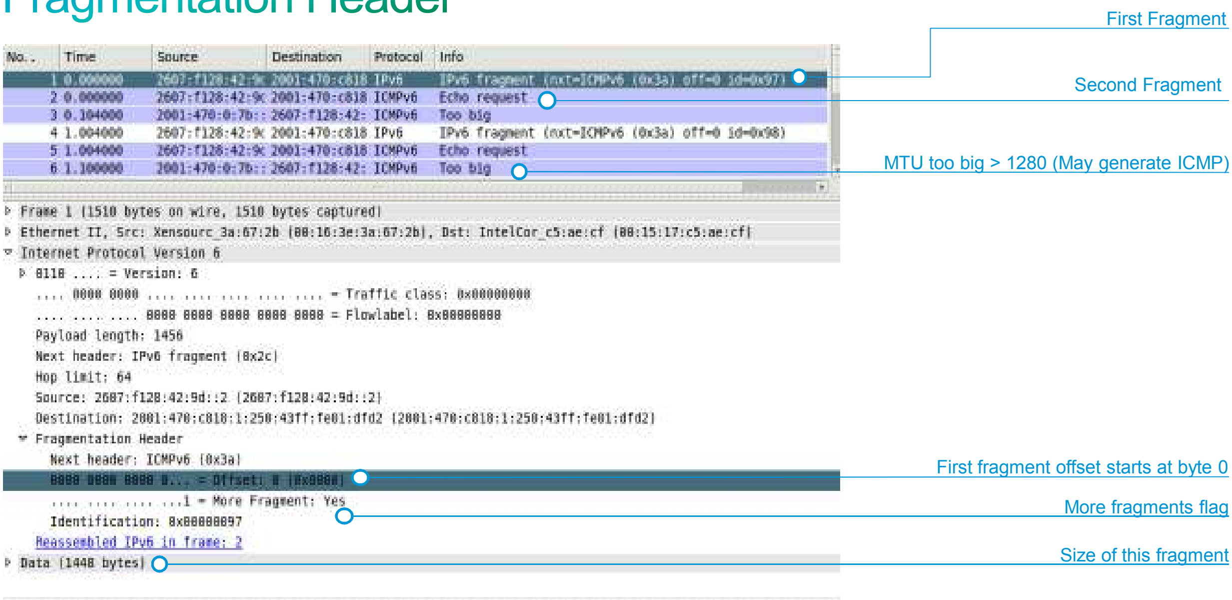

Fragmentation Header - Reassembled

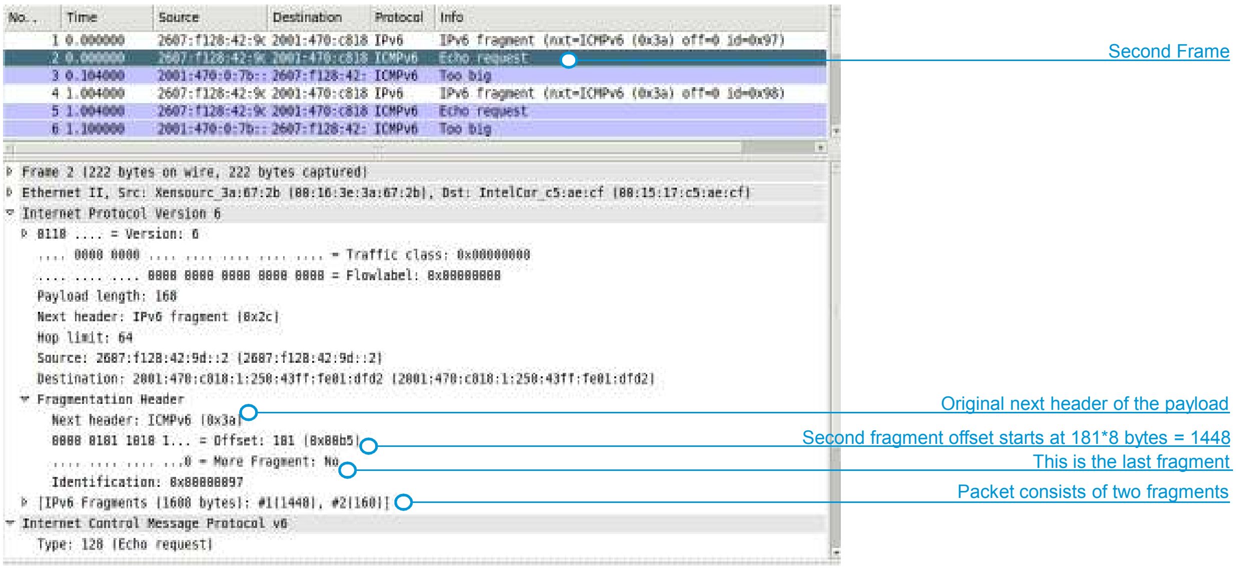

 Cisco Confidential 39

Fragment Header Summary

- 32-bit ID field (similar to IPv4)
- Next Header is the original value of the next protocol, before fragmentation
- Fragment Offset (13 bits)

  Represented in 8-octet units of the data following this header relative to the start of fragmentable part of the packet

  First fragment offset will always be zero
- M = “more fragments” flag

 Cisco Confidential 40

Path MTU Discovery (RFC 1981)

- Each Link has MTU a maximum transmission unit
- Path MTU minimum MTU of all the links in a path between a source and a destination
- Minimum link MTU for IPv6 is 1280 octets
  - In comparison IPv4 minimum MTU is 68 octets
  - If Link MTU < 1280 then fragmentation and reassembly must be used
  - If IPv6 payload > 1280 fragmentation may need to be performed
- PMTU Discovery is expected to be performed by IPv6 end hosts
  - It should only apply if sending packets > 1280 bytes
  - For each destination, start by assuming MTU of first-hop link
  - Exceeding the link MTU invokes ICMP “packet too big” back to source
  - Message includes the offending link MTU value
  - MTU is then cached by source for specific destination

Path MTU Discovery

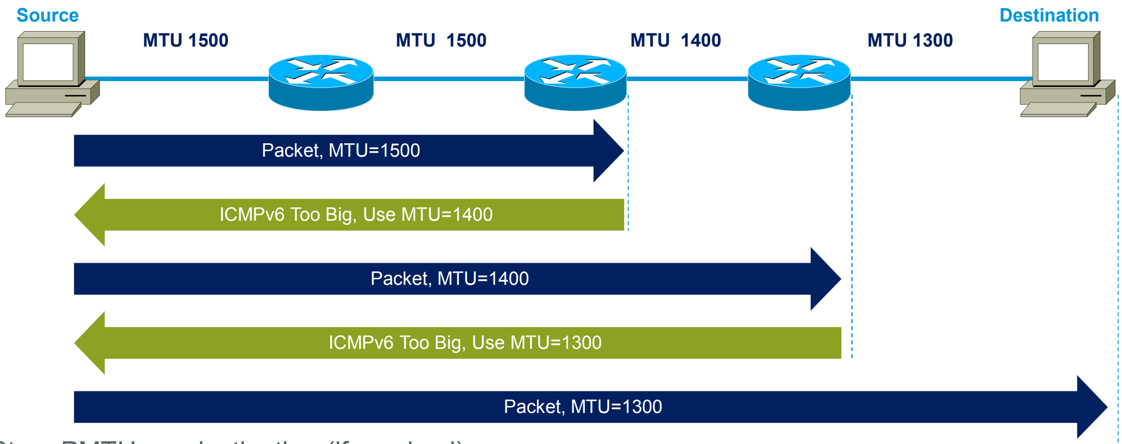

- Store PMTU per destination (if received)
- Age out PMTU (10 mins), reset to first link MTU

ICMPv6 and Neighbor Discovery

ICMPv6 (RFC 2463)

- Internet Control Message Protocol version 6
- Combines several IPv4 functions
  - ICMPv4, IGMP and ARP
- Message types are similar to ICMPv4
  - Destination unreachable (type 1)
  - Packet too big (type 2)
  - Time exceeded (type 3)
  - Parameter problem (type 4)
  - Echo request/reply (type 128 and 129)

ICMPv6 Header

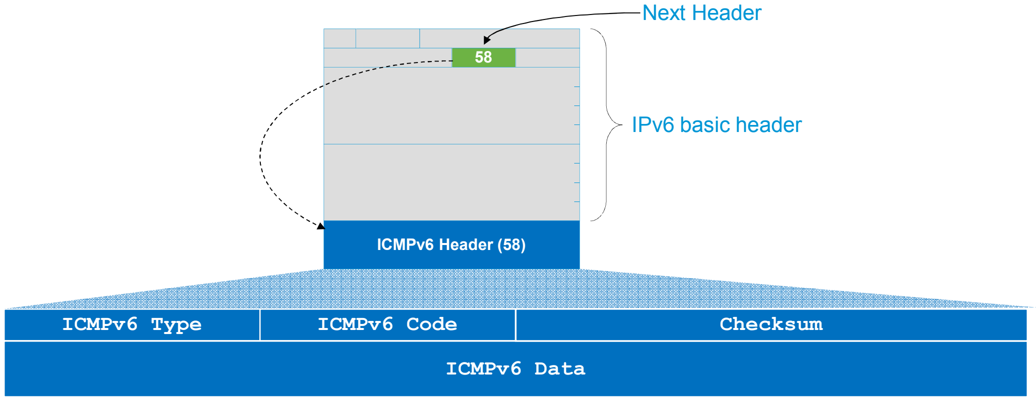

- Also used for Neighbor Discovery, Path MTU discovery and Mcast listener discovery (MLD)
  - Type - identifies the message or action needed
  - Code – is a type-specific sub-identifier.
  - Checksum – computed over the entire ICMPv6

Neighbor Discovery Messages (ND)

- ND uses ICMPv6 messages
  - Originated from node on link local with a hop limit of 255
  - Receivers checks hop limit is still 255 (has not passed a router)
- Consists of IPv6 header, ICMPv6 header, neighbor discovery header, and neighbor discovery options
- Five neighbor discovery messages

Message	Purpose	ICMP Code	Sender	Target
Router Solicitation (RS)	Prompt routers to send RA	133	Nodes	All routers
Router Advertisement (RA)	Advertise default router, prefixes Operational parameters	134	Routers	Sender of RS All routers
Neighbor Solicitation (NS)	Request link-layer of target	135	Node	Solicited Node Target Node
Neighbor Advertisement (NA)	Response to NS (solicited) Advertise link-layer address change (Unsolicited)	136	Nodes	
Redirect	Inform hosts of a better first hop	137	Routers	

ICMPv6 Neighbor Discovery (RFC 4861)

- Replaces ARP, ICMP (redirects, router discovery)
- Uses ICMPv6 header
- Reachability of neighbours
- Hosts use it to discover routers, auto configuration of addresses (SLAAC)
- Duplicate Address Detection (DAD)

IPv4/IPv6 Provisioning Comparison

Function	IPv4	IPv6
Address Assignment	DHCPv4	DHCPv6, SLAAC, Reconfiguration
Address Resolution	ARP RARP	ICMPv6 NS, NA Not Used
Router Discovery	ICMP Router Discovery	ICMPv6 RS, RA
Name Resolution	DNS	DNS

Router Solicitation and Advertisement (RS & RA)

Router Solicitation	
ICMP Type	133
IPv6 Source	A Link Local (FE80::1)
IPv6 Destination	All Routers Multicast (FF02::2)
Query	Please send RA

Router Advertisement	
ICMP Type	134
IPv6 Source	A Link Local (FE80::2)
IPv6 Destination	All Nodes Multicast (FF02::1)
Data	Options, subnet prefix, lifetime, autoconfig flag

- Router solicitations (RS) are sent by booting nodes to request RAs for configuring the interfaces
- Routers send periodic Router Advertisements (RA) to the all-nodes multicast address

Neighbor Solicitation & Advertisement

- Neighbor Solicitation (NS)
  - Used to discover link layer address of IPv6 node

NS Function	Source	Destination
Address resolution	Unicast	Solicited Node Multicast
Node reachability	Unicast	Unicast
Duplicate Address Detection	::0	Solicited Node Multicast

- Neighbor Advertisement (NA)
  - Response to neighbor solicitation (NS) message
  - A node may also send unsolicited Neighbor Advertisements to announce a link-layer address change.

Neighbor Solicitation & Advertisement (NS & NA)

Neighbour Solicitation	
ICMP Type	135
IPv6 Source	A Unicast
IPv6 Destination	B Solicited Node Multicast
Data	FE80:: address of A
Query	What is B link layer address?

Neighbour Advertisment	
ICMP Type	136
IPv6 Source	B Unicast
IPv6 Destination	A Unicast
Data	FE80:: address of B, MAC Address

Contents of Neighbour Solicitation

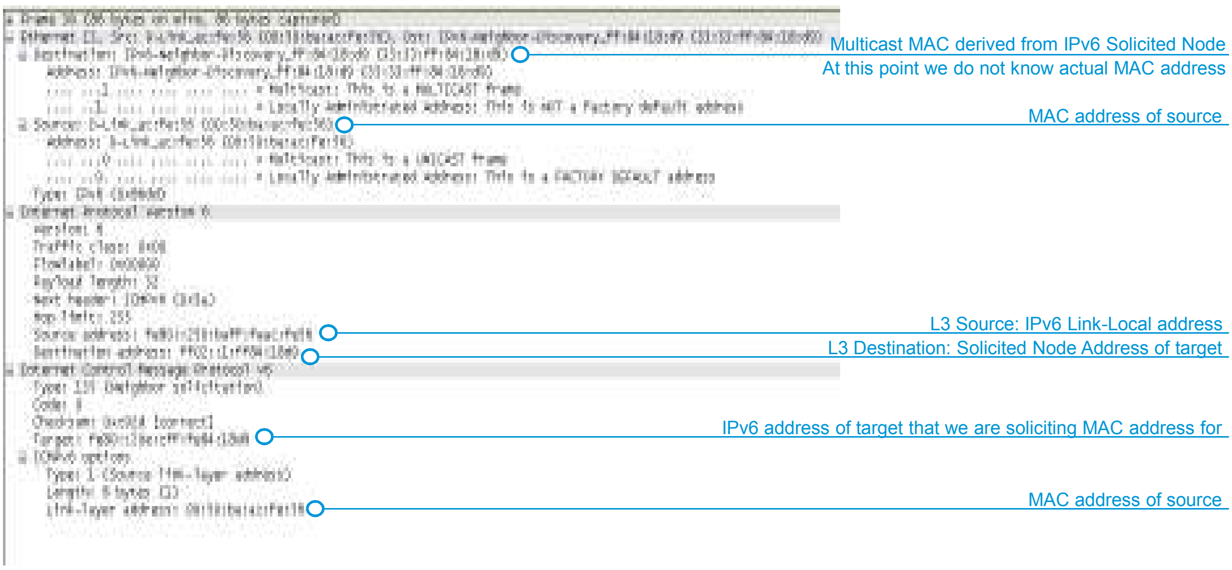

Contents of Neighbor Advertisement

```
⊞ Frame 37 (86 bytes on wire, 86 bytes captured)
⊟ Ethernet II, Src: Intel_84:18:d9 (00:0e:0c:84:18:d9), Dst: D-Link_ac:fe:56 (00:50:ba:ac:fe:56)
  ⊟ Destination: D-Link_ac:fe:56 (00:50:ba:ac:fe:56)
      Address: D-Link_ac:fe:56 (00:50:ba:ac:fe:56)
      .... ...0 .... .... .... .... = Multicast: This is a UNICAST frame
      .... ..0. .... .... .... .... = Locally Administrated Address: This is a FACTORY DEFAULT address
  ⊟ Source: Intel_84:18:d9 (00:0e:0c:84:18:d9)
      Address: Intel_84:18:d9 (00:0e:0c:84:18:d9)
      .... ...0 .... .... .... .... = Multicast: This is a UNICAST frame
      .... ..0. .... .... .... .... = Locally Administrated Address: This is a FACTORY DEFAULT address
    Type: IPv6 (0x86dd)
⊟ Internet Protocol Version 6
    Version: 6
    Traffic class: 0x00
    Flowlabel: 0x00000
    Payload length: 32
    Next header: ICMPv6 (0x3a)
    Hop limit: 255
    Source address: fe80::20e:cff:fe84:18d9
    Destination address: fe80::250:baff:feac:fe56
⊟ Internet Control Message Protocol v6
    Type: 136 (Neighbor advertisement)
    Code: 0
    Checksum: 0xed8b [correct]
  ⊟ Flags: 0x60000000
      0... .... .... .... .... .... .... .... = Not router
      .1.. .... .... .... .... .... .... .... = Solicited
      ..1. .... .... .... .... .... .... .... = Override
    Target: fe80::20e:cff:fe84:18d9
  ⊟ ICMPv6 options
      Type: 2 (Target link-layer address)
      Length: 8 bytes (1)
      Link-layer address: 00:0e:0c:84:18:d9
```

IPv6 source address of Neighbor Advertisement

IPv6 Destination of original Neighbor Solicitation

Link layer address in response to NS

Viewing Neighbors in the Cache

- Neighbors are only considered “reachable” for 30-seconds
- “Stale” indicates that ND packet must be sent again

```
R1#sho ipv6 neighbors
IPv6 Address                              Age Link-layer Addr State Interface
FE80::A8BB:CCFF:FE00:7800                   0 aabb.cc00.7800  STALE Et0/0
FE80::A8BB:CCFF:FE00:7A00                  50 aabb.cc00.7a00  STALE Et0/0
```

Entry STALE due to no contact for > 30 secs (Age 50 secs)

```
R1#ping ipv6
Target IPv6 address: FE80::A8BB:CCFF:FE00:7A00
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands? [no]:
Output Interface: Ethernet0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FE80::A8BB:CCFF:FE00:7A00, timeout is 2 second
s:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 16/24/32 ms
R1#sho ipv6 neighbors
IPv6 Address                              Age Link-layer Addr State Interface
FE80::A8BB:CCFF:FE00:7800                   3 aabb.cc00.7800  STALE Et0/0
FE80::A8BB:CCFF:FE00:7A00                   0 aabb.cc00.7a00  REACH Et0/0
```

After PING entry now reachable again (Age 0 secs)

Neighbour Unreachability Detection

- Neighbor is declared reachable if
  - The connection is making forward progress
  - Previously sent data is known to have been delivered correctly
  - Source receives an NA in response to NS
- If neighbour status unknown then send NS
- Defined in RFC 4861 Section 7.3

 55

Neighbor Cache Entry States

★

- INCOMPLETE
  - Address resolution is in progress and the link-layer address of the neighbor has not yet been determined
- REACHABLE
  - The neighbor is known to have been reachable recently (within tens of seconds ago)
- STALE
  - The neighbor is no longer known to be reachable but until traffic is sent to the neighbor, no attempt should be made to verify its reachability
- DELAY
  - Delay sending probes for a short while in order to give upper layer protocols a chance to provide reachability confirmation
- PROBE
  - The neighbor is no longer known to be reachable, and unicast Neighbor Solicitation probes are being sent to verify reachability

 56

Duplicate Address Detection (DAD)

Neighbour Solicitation	
ICMP Type	135 (Neighbour Solicitation)
Ethernet DA	33-33-FF-52-F9-D8
IPv6 Header	
IPv6 Source	::
IPv6 Destination	FF02::1:FF52:F9D8
Hop Limit	255
NS Header	
Target Address	FE80::260:8FF:FE52:F9D8

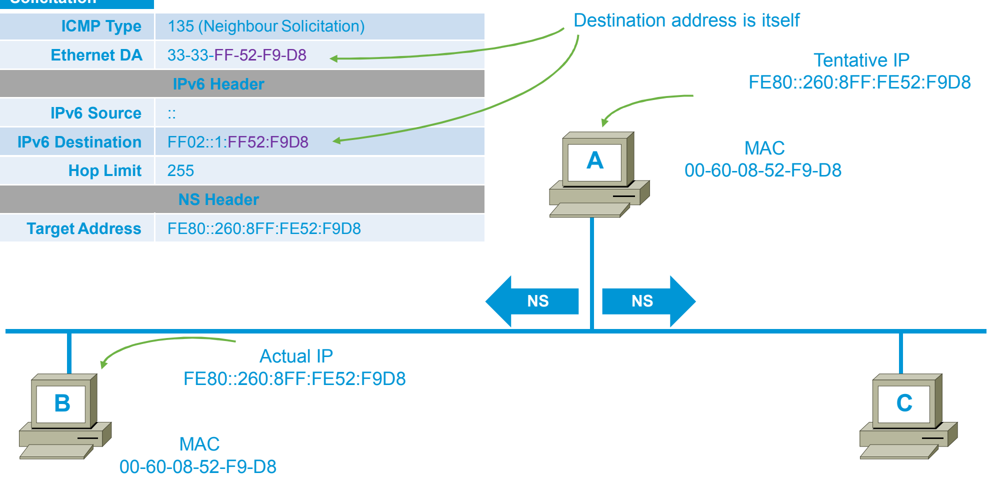

Duplicate Address Detection Response

Neighbour Solicitation	
ICMP Type	135 (Neighbour Solicitation)
Ethernet DA	33-33-00-00-00-01
IPv6 Header	
IPv6 Source	FE80::260:8FF:FE52:F9D8
IPv6 Destination	FF02::1
Hop Limit	255
NA Header	
Target Address	FE80::260:8FF:FE52:F9D8
Neighbour Discovery Option	
Target L2 Address	00-60-08-52-F9-D8

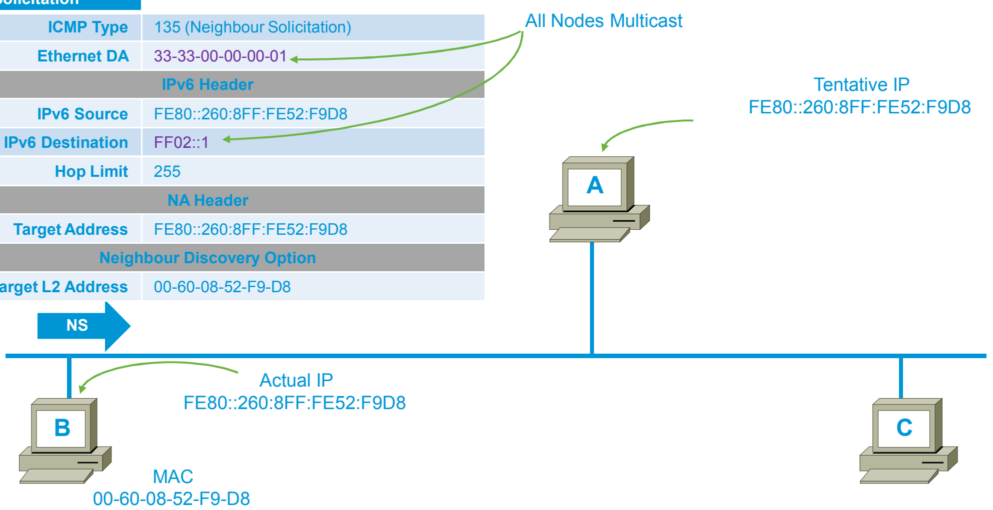

ICMPv6 Redirection

- Redirect is used by a router to informs hosts of a better first hop

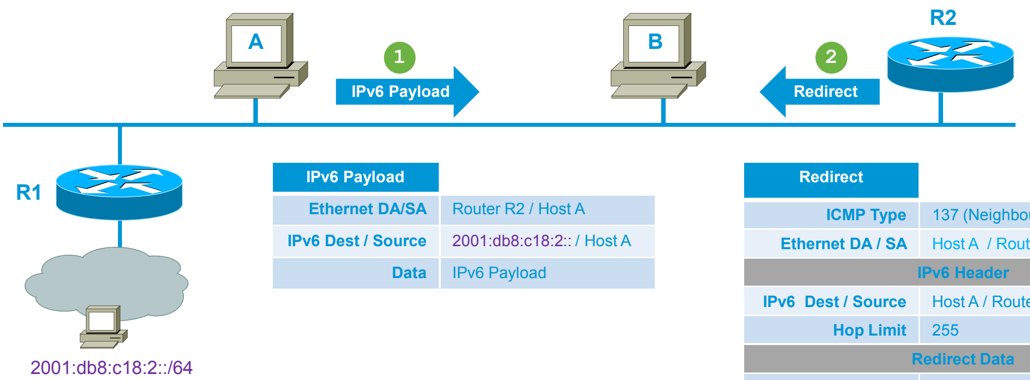

IPv6 Payload	
Ethernet DA/SA	Router R2 / Host A
IPv6 Dest / Source	2001:db8:c18:2:: / Host A
Data	IPv6 Payload

Redirect	
ICMP Type	137 (Neighbour Solicitation)
Ethernet DA / SA	Host A / Router R2
IPv6 Header	
IPv6 Dest / Source	Host A / Router R2
Hop Limit	255
Redirect Data	
Target Address	Router R1
Redirected Prefix	2001:db8:c18:2::/64

Stateless Address Autoconfiguration (RFC4862)

- Autoconfiguration is used to automatically assigned an address to a host "plug and play"
  - Generating a link-local address,
  - Generating global addresses via stateless address autoconfiguration
  - Duplicate Address Detection procedure to verify the uniqueness of the addresses on a link

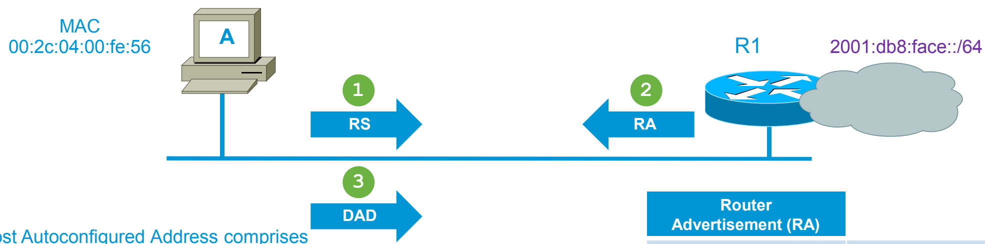

Host Autoconfigured Address comprises Prefix Received + Link-Layer Address if DAD check passes

`2001:db8:face::22c:4ff:fe00:fe56`

Router Advertisement (RA)	
Ethernet DA/SA	Router R2 / Host A
Prefix Information	2001:db8:face::/64
Default Router	Router R1

Prefix Renumbering

- Prefixes can be given a lifetime in RA messages
- Allows seamless transition for renumbering to a new prefix

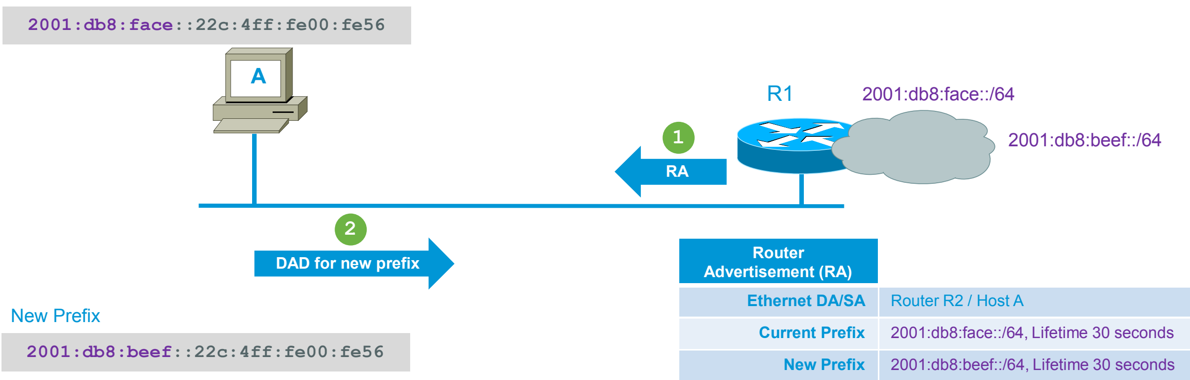

Prefix Renumbering Configuration

```
interface Ethernet0
ipv6 nd prefix 2001:db8:face::/64 43200 0
ipv6 nd prefix 2001:db8:beef::/64 43200 43200
!
! Alternative configuration
!
interface Ethernet0
  ipv6 nd prefix 2001:db8:face::/64 at Jul 31 2008 23:59 Jul 20 2008 23:59
  ipv6 nd prefix 2001:db8:beef::/64 43200 43200
```

- Valid Lifetime for advertised prefixes (default 30 days)
- Preferred lifetime for addresses generated from SLAAC (default 7 days)
- Original advertised prefix - new hosts will not use SLAAC with this prefix
- New advertised prefix - new hosts will use SLAAC with this prefix
- Time based configuration

IPv6 Interface Configuration

Link-Local Configured Interface Identifier Address (IOS)

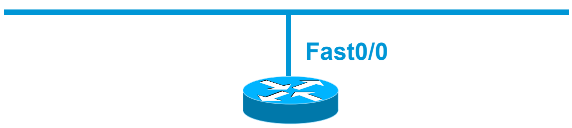


```
ipv6 unicast-routing
!
interface FastEthernet0/0
 ip address 10.151.1.1 255.255.255.0
 ip pim sparse-mode
 duplex auto
 speed auto
 ipv6 enable
 ipv6 nd ra-interval 30
!
```

Enable IPv6 routing (`ipv6 unicast-routing`)

Enable IPv6 on interface and automatically create link-local address (`ipv6 enable`)

IPv6 Interface with Link-Local Address

```
r1#show ipv6 interface fast0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::207:50FF:FE5E:9460
  Global unicast address(es):
    None
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF5E:9460
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 30 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.

r1# show interface fast0/0
FastEthernet0/0 is up, line protocol is up
  Hardware is AmdFE, address is 0007.505e.9460 (bia 0007.505e.9460)
```

- FE80::207:50FF:FE5E:9460 — EUI-64 derived from MAC address 0007.505e.9460
- FF02::1 — Listening for all hosts multicast
- FF02::2 — Listening for all routers multicast
- FF02::1:FF5E:9460 — Solicited Node multicast for link-local address
- 0007.505e.9460 — MAC address 0007.505e.9460

Manually Configured Interface Identifier Address

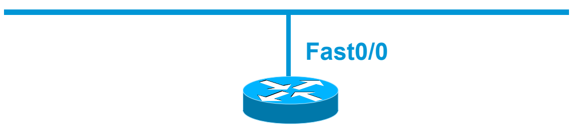


```
ipv6 unicast-routing
!
interface FastEthernet0/0
 ip address 10.151.1.1 255.255.255.0
 ip pim sparse-mode
 duplex auto
 speed auto
 ipv6 address 2006:1::1/64
 ipv6 nd ra-interval 30
!
```

- ipv6 address 2006:1::1/64 — Enables IPv6 and assigns a global prefix and manual interface ID
- ipv6 nd ra-interval 30 — Send RA every 30 seconds

IPv6 Interface with Manual Interface Address

```
r1#show ipv6 interface fast0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::207:50FF:FE5E:9460
  Global unicast address(es):
    2006:1::1, subnet is 2006:1::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:1
    FF02::1:FF5E:9460
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 30 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
r1#
```

Global unicast address with manual interface ID of “1”

Routable /64 subnet

Corresponding Solicited Node multicast address for manual interface ID

Corresponding Solicited Node multicast address for Link-Local interface ID

EUI-64 Configured Interface Identifier Address

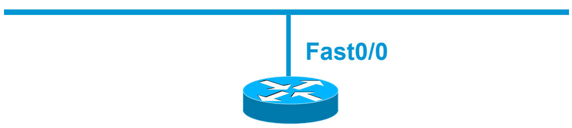


```
ipv6 unicast-routing
!
interface FastEthernet0/0
 ip address 10.151.1.1 255.255.255.0
 ip pim sparse-mode
 duplex auto
 speed auto
 ipv6 address 2006:1::/64 eui-64
 ipv6 nd ra-interval 30
!
```

Enables IPv6 and assigns a global prefix and EUI-64 interface ID

IPv6 Interface with EUI-64 Interface Address

```
r1#show ipv6 interface fast0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::207:50FF:FE5E:9460
  Global unicast address(es):
    2006:1::207:50FF:FE5E:9460, subnet is 2006:1::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF5E:9460
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 30 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.

r1#show interface fast0/0
FastEthernet0/0 is up, line protocol is up
  Hardware is AmdFE, address is 0007.505e.9460 (bia 0007.505e.9460)
```

Link-Local address with EUI-64 interface ID

Manually configured address with EUI-64 Interface ID

Solicited Node multicast for both manual and link-local address

MAC address 0007.505e.9460 used for EUI-64

Review Questions

Question 1

- What type of address is
  2001:0ba0:0000:0000:0000:0000:0000:1234

A. Link Local ✗ Incorrect link-local begin with FE80::/10

B. Provider Independent Multicast ✗ Multicast addresses begin with FF00::/8

C. Global Unicast ✓

D. Unique Local ✗ Unique-local begin with FC00::/7

Question 2

- Which of the following is a valid abbreviation for
  2001:0ba0:0000:0000:0000:0000:0000:1234

A. `2001:0ba0::1234` ✓

B. `2001:ba0:0:0:0:0:0:1234` ✓

C. `2001:0ba::1234` ✗ Incorrect because 0ba0 not equal to 0ba (only leading zeros can be omitted)

D. `2001:0ba0::0:0:0::1234` ✗ Incorrect :: cannot be used more than once

Question 3

- Which of the following is a valid EUI-64 address for an interface with the MAC address `58:b0:35:fe:7e:4a`

A. `2001::5ab0:35ff:fefe:7e4a` ✓

B. `2001::58b0:35ff:fffe:7e4a` ✗ Incorrect FFFE should be inserted in the middle not FFFF

C. `2001::58b0:35ff:feff:7e4a` ✗ Incorrect FE in the MAC address should not be changed to FF

D. `2001::58b0:35ff:fefe:7e4a` ✗ Incorrect because bit 7 was not flipped in EUI-64 interface ID

 Cisco Confidential 73

Question 4

- Which well known multicast addresses are mandatory to have on an interface?

A. Node-Local “All Nodes”, Link-Local “All Nodes” and Link-Local “Solicited-Node” ✗

B. Link-Local “All Nodes” and Link-Local “Solicited Node” ✓

C. Link-Local “All Nodes” only ✗

D. Node-Local “All Nodes” only ✗

 Cisco Confidential 74

Question 5.

- Which type of request does a node use to learn if another node has the same address?

A. Unicast DAD ✗

B. Neighbour Solicitation (NS) ✗
With src = Unicast, dst = solicited node multicast

C. Router Solicitation (RS) ✗

D. Neighbour Solicitation (NS) ✓
With src = ::, dst = solicited node multicast