
TOMORROW
starts here.
CISCO
Cisco live!

Understanding IPv6

BRKRST-1069

Shannon McFarland – CCIE #5245

Principal Engineer

#clmel

Cisco live!

Agenda

- Why IPv6?
- What is IPv6?
- How does IPv6 work?

Why IPv6?

You Have Heard It All Before

- IANA and the RIRs have run out of IPv4 address

- Consumers are generally ambivalent
  - Do not/should not care whether IPv4 or IPv6 is used to deliver their content
  - Won't understand they "Are on the wrong protocol"!

- IPv4 address trading markets
  - Growth, fragmentation, and identity verification of the IPv4 routing table is inevitable
    - /15 IPv4 Addresses For Sale; Asking $9.00/IP
    - /16 IPv4 Addresses For Sale, Asking $9.20/IP
    - /17 IPv4 Addresses For Sale, Asking $9.50/IP

Common Drivers in the Enterprise

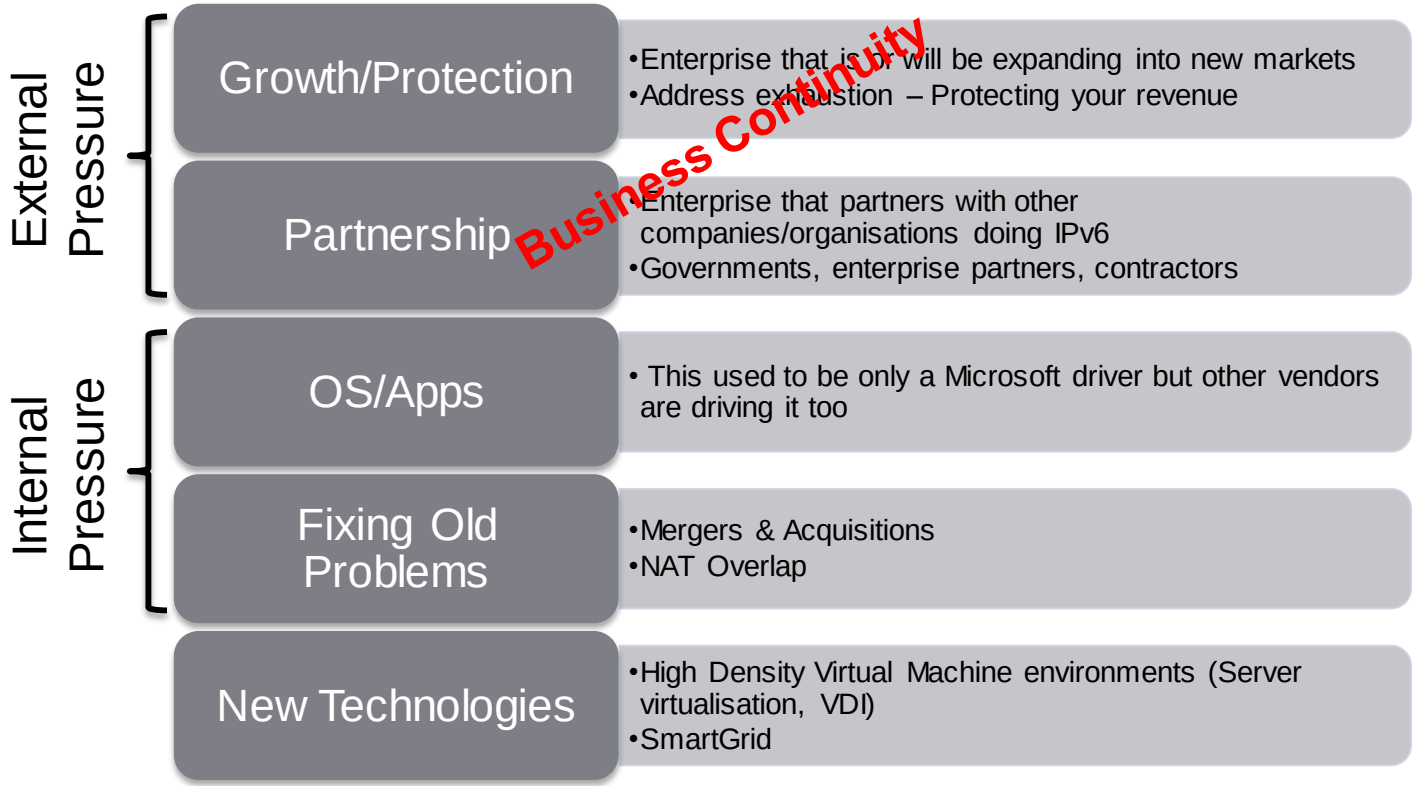

What is IPv6?

What is IPv6?

- 128bit addressing scheme
  - Hexadecimal representation
  - CIDR masking

- Introduces new protocol level behaviours
  - Neighbour Discovery
  - Stateless Addressing
  - No more Broadcast, only Multicast

So How Big Is The IPv6 Address Space?

This one is bigger ☺

340,282,366,920,938,463,463,374,607,432,768,211,456
(IPv6 Address Space - 340 Trillion Trillion Trillion)

vs

4,294,967,296
(IPv4 Address Space - 4 Billion)

IPv4 and IPv6 Header Comparison

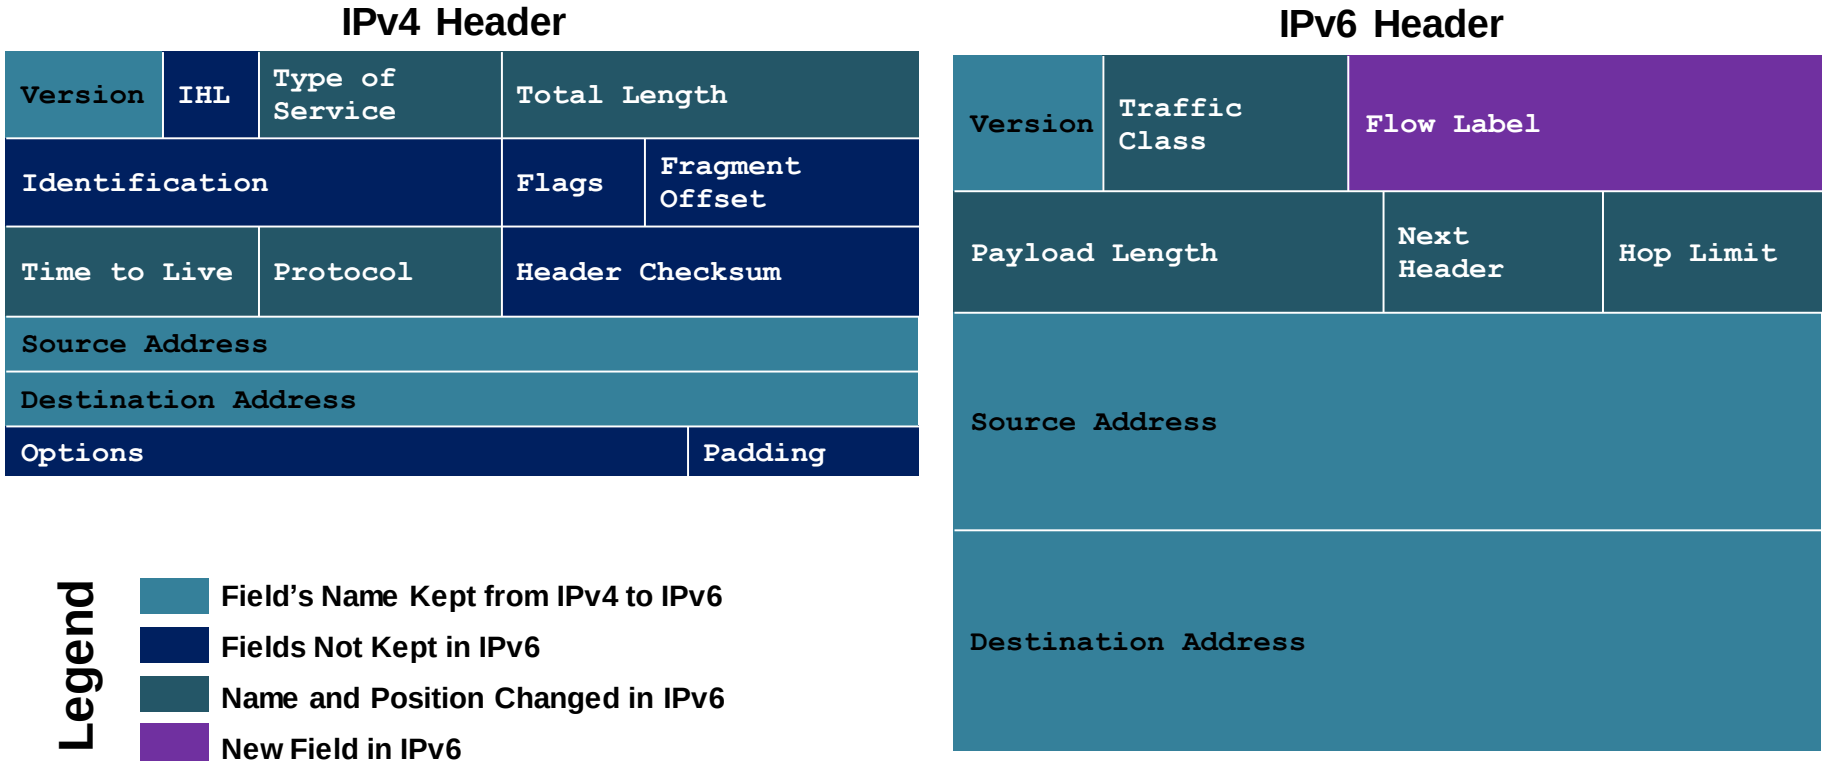

Extension Headers

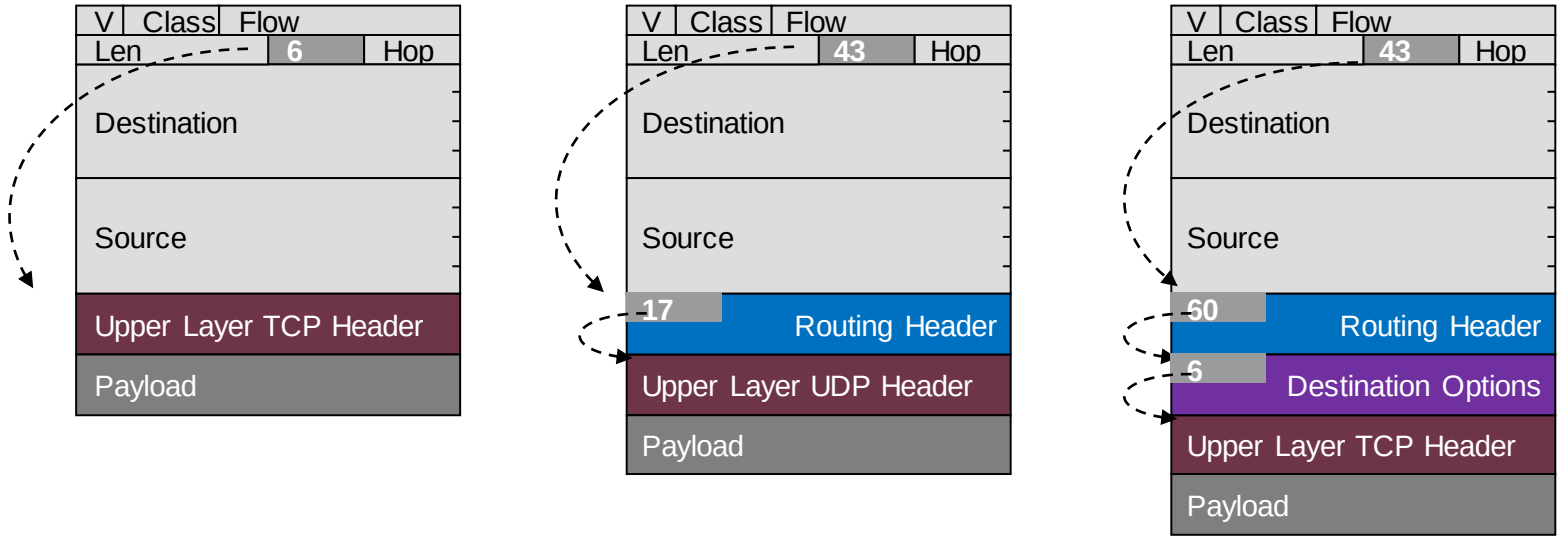

- Extension headers are daisy chained
- Order is important!

IPv6 Protocol Stack

New features

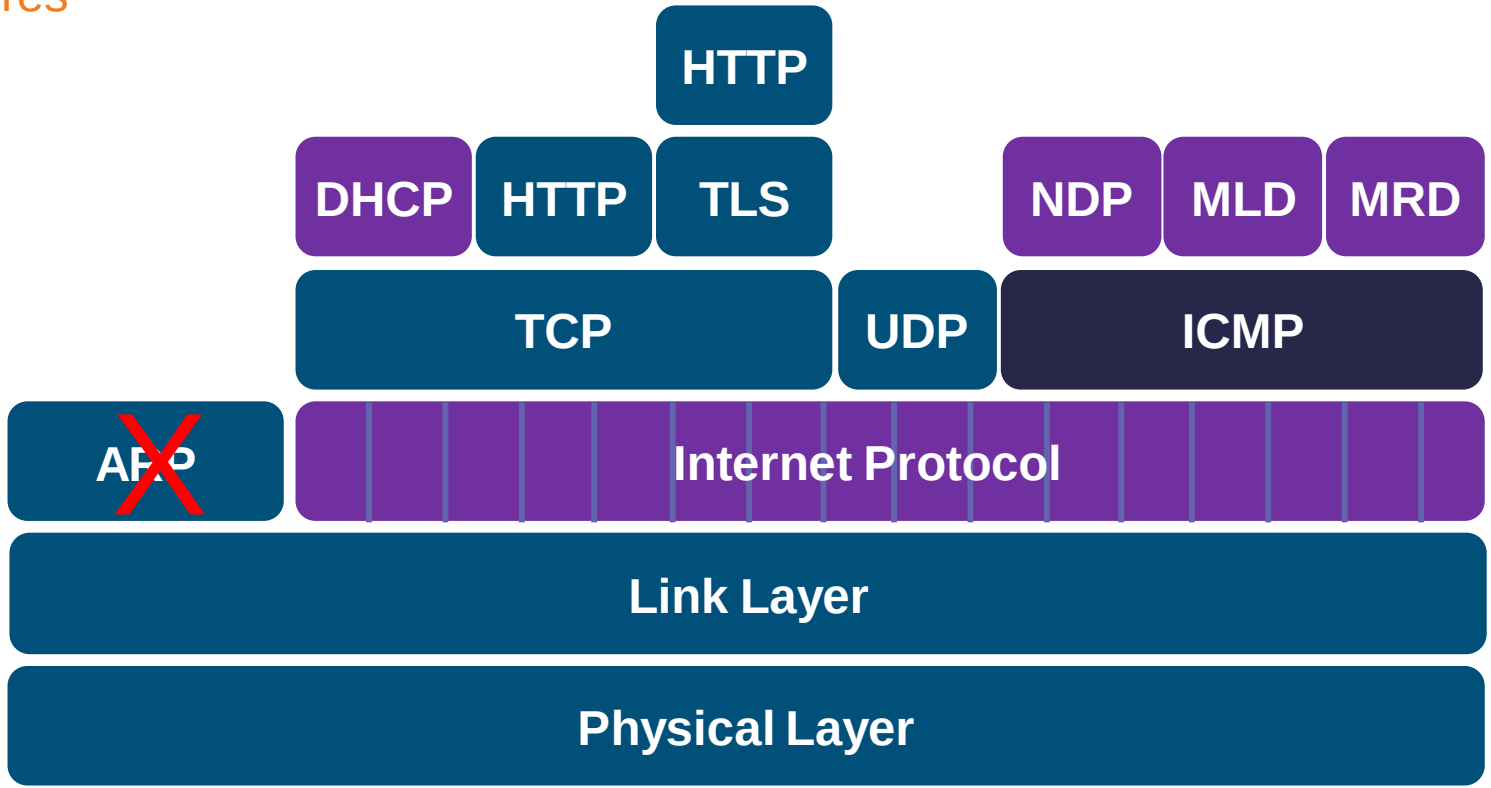

Cisco live!

IPv4/IPv6 Technology Comparison

Service	IPv4	IPv6
Addressing Range	32-bit, NAT	128-bit, Multiple Scopes
IP Provisioning	Manual, DHCP	Manual, SLAAC, DHCP (and renumbering capability)
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 – The First Half

IPv6 Addresses

- IPv6 addresses are 128 bits long
  - Segmented into 8 groups of four HEX characters
  - Separated by a colon (:)

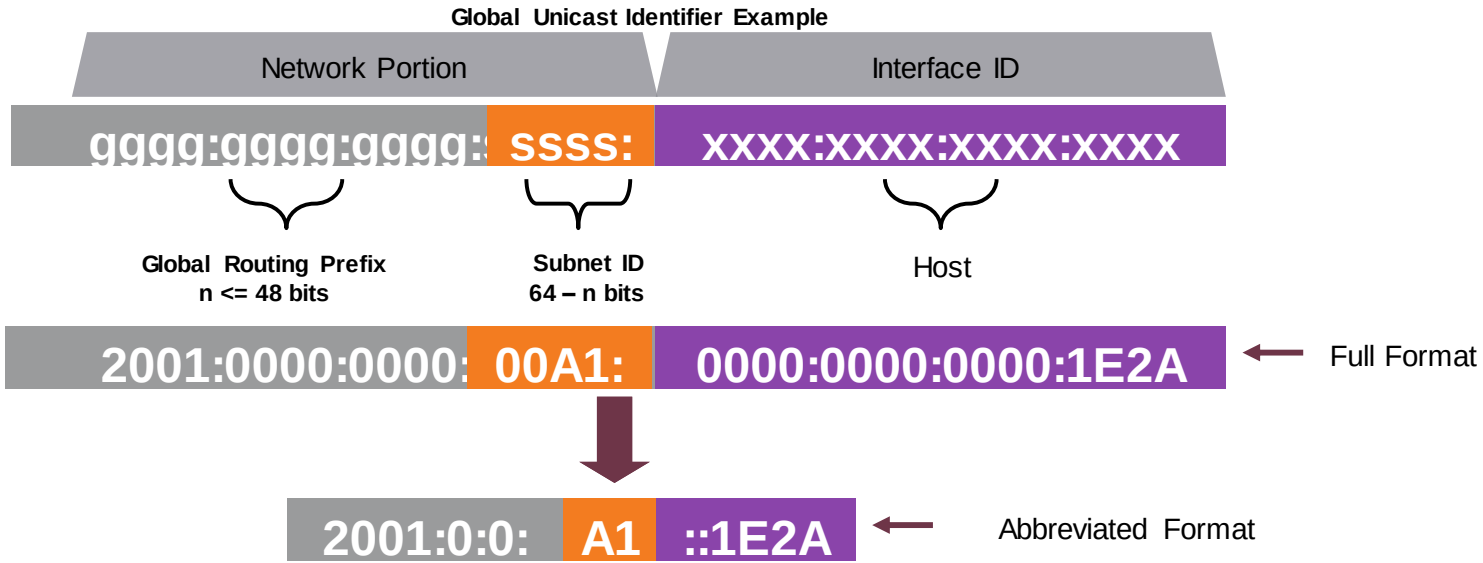

IPv6 Address Syntax

- Hex numbers are not case sensitive
  - **2001:0dB8:0000:130f:0000:0000:087c:AaAa**
- Abbreviations are possible
  - **2001:0db8:0000:130f::87c:aaaa**
    - Zeros in contiguous blocks can be represented by ::
    - Double colon can only appear once in the address
- Only leading zeros can be omitted
  - **2001:db8:0:130f::87c:aaaa**
- IPv6 uses CIDR representation
  - **2001:0db8:0000:130f:0000:0000:087c:aaaa/128**

IPv6 Address Syntax

- Loopback address representation

```
0:0:0:0:0:0:0:1 == ::1
```

  - Same as 127.0.0.1 in IPv4, it identifies self

- Unspecified address representation

```
0:0:0:0:0:0:0:0 == ::
```

  - Initial DHCP request, Duplicate Address Detection DAD

- Default Route representation

```
::/0
```

IPv6 Address Scopes

- An IPv6 interface is “expected” to have multiple addresses and multiple scopes

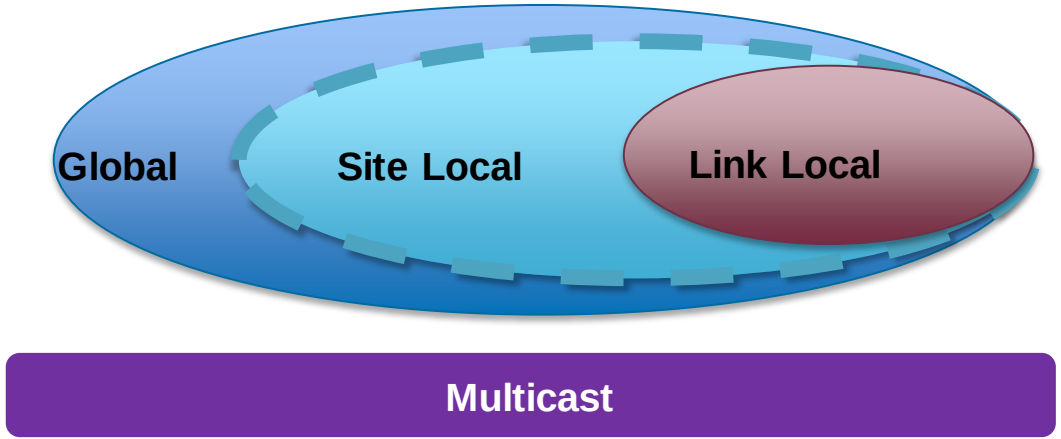

IPv6 Address Types

- Three types of unicast addresses
  - Link-Local – Non routable exists on single layer 2 domain **(fe80::/64)**
  - Unique-Local – Routable within administrative domain **(fc00::/7)**
  - Global – Routable across the Internet **(2000::/3)**
- Multicast addresses **(ff00::/8)**
  - Flags (z) in 3rd nibble (4 bits) Scope (s) into 4th nibble

Link Local Address

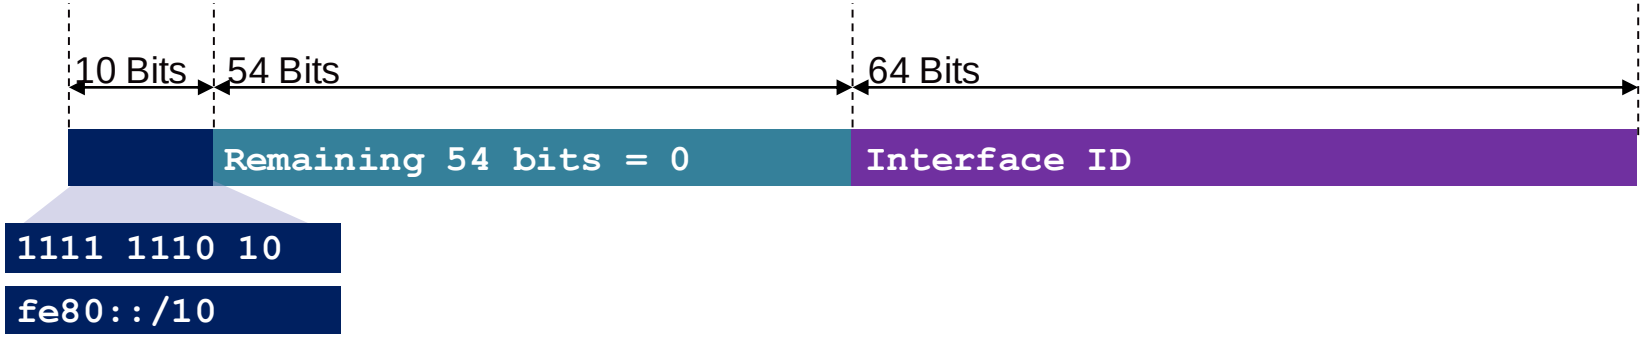

- Mandatory
- Automatically self assigned by the device using EUI-64
- Only link specific scope

Unique Local Address (RFC 4193)

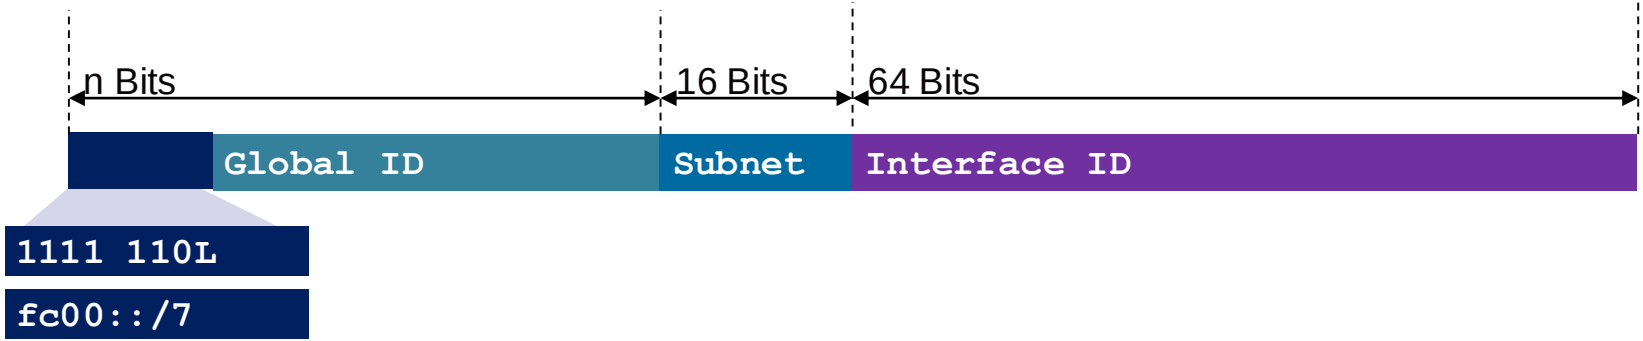

- FC00::/8 is Registry Assigned (L bit = 0), FD00::/8 is self generated (L bit = 1)
  - Registries not yet assigning ULA space
- Global ID can be generated using an algorithm
  - Low order 40 bits result of SHA-1 Digest {EUI-64 && Time}
- Not considered best practice

Global Unicast Addresses

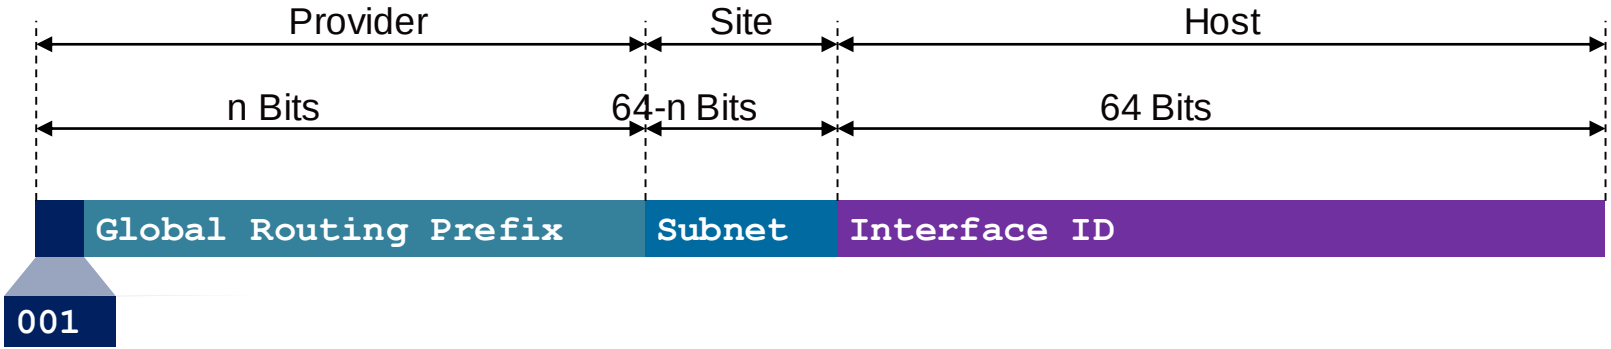

- Globally routable
  - Requires correct border security!!
- Considered best practice for all device numbering
- Common allocation sizes are /32, /48, /52, /56, /64

Interface Address Set

Reference

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

PI and PA Allocation Process

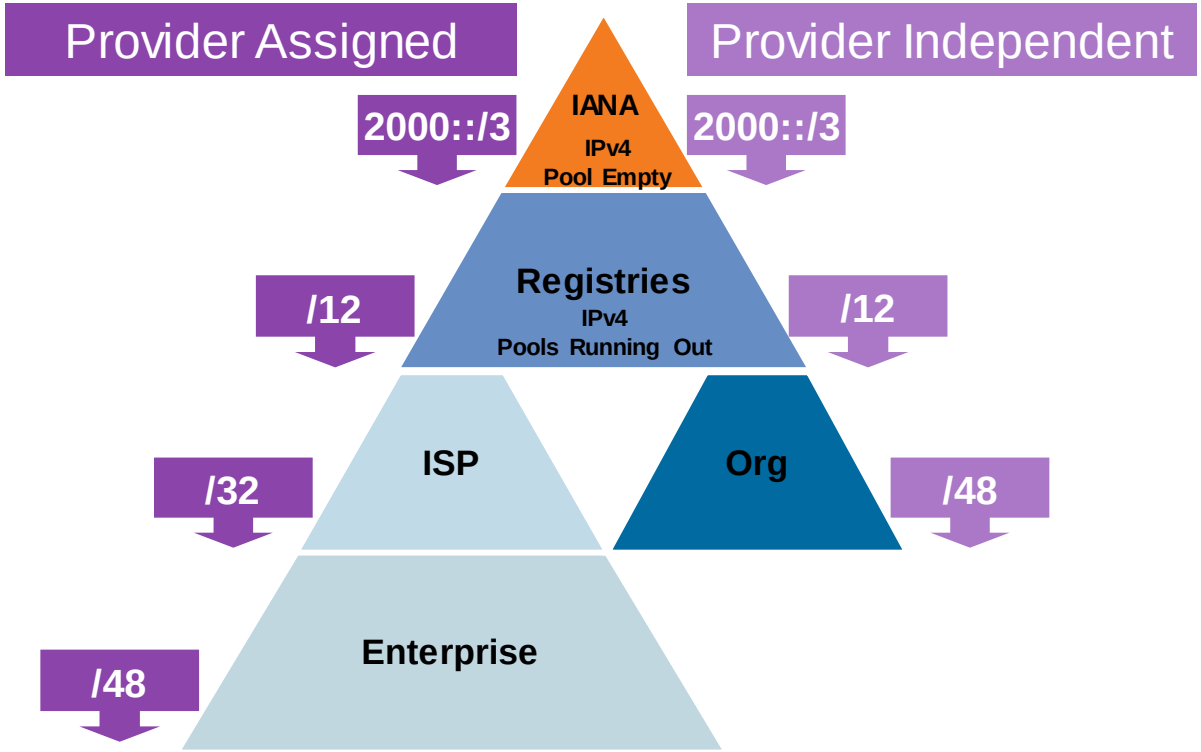


IPv6 Addressing – The Second Half

Address Interface ID

- Interface ID of 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

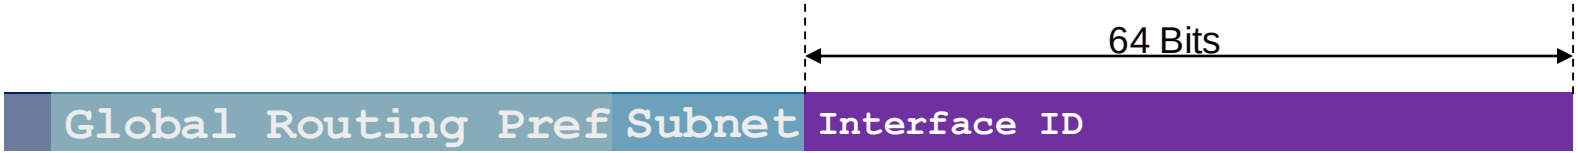

IPv6 Interface Identifier (EUI-64 format)

- This format expands the 48 bit MAC address to 64 bits by inserting FFFE into the middle 16 bits
  - Non-ethernet interfaces use the first MAC address in the pool on the router
  - Cisco devices ‘bit-flip’ the 7th bit

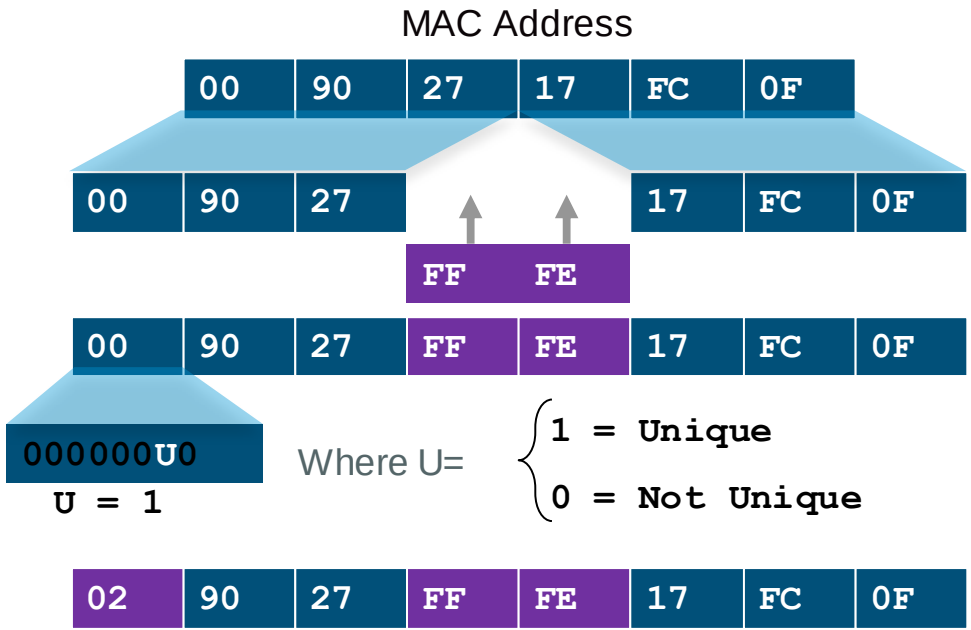

Randomised IID and Privacy Extensions

- Enabled by default on Microsoft Windows
- Enable/disable via GPO or CLI

```
netsh interface ipv6 set global randomizeidentifiers=disabled store=persistent
netsh interface ipv6 set privacy state=disabled store=persistent
```

- Alternatively, use DHCP to a specific pool
- Randomised address are generated for non-temporary autoconfigured addresses including public and link-local
- Randomised addresses engage Optimistic DAD

Link Level - Prefix Length Considerations

64 bits

- Recommended by RFC3177 and IAB/IESG
- Consistency makes management easy
- MUST for SLAAC (MSFT DHCPv6 also)
- Significant address space loss (18.466 Quintillion)

> 64 bits

- Address space conservation
- Special cases:
  /126—valid for p2p
  /127—valid for p2p if you are careful – RFC6164 (RFC3627)
  /128—loopback
- Must avoid overlap with specific addresses:
  Router Anycast (RFC3513)
  Embedded RP (RFC3956)
  ISATAP addresses

- /64 everywhere
- /64 + /126
  - 64 on host networks
  - 126 on P2P
- /64 + /127
  - 64 on host networks
  - 127 on P2P
- /128 on loopback


The Role of Multicast

IPv6 Multicast Address (RFC 4291)

An IPv6 multicast address has the prefix FF00::/8 (1111 1111)

- Second octet defines lifetime and scope

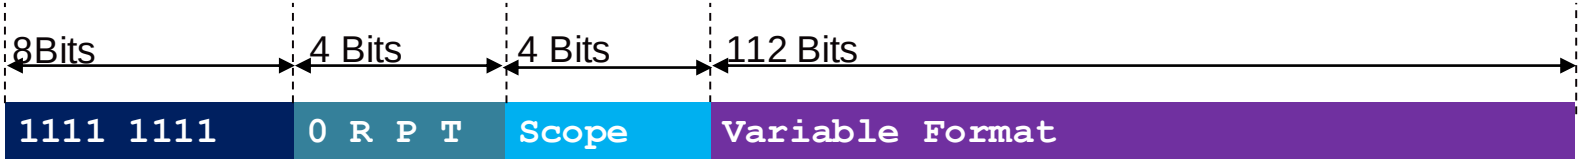

Flags	
R = 0 R = 1	No embedded RP Embedded RP
P = 0 P = 1	Not based on unicast Based on unicast
T = 0 T = 1	Permanent address (IANA assigned) Temporary address (local assigned)

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

Well Known Multicast Addresses

Reference

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
- Used in neighbour solicitation (NS) messages
- FF02::1:FF & {lower 24 bits from IPv6 Unicast interface ID}

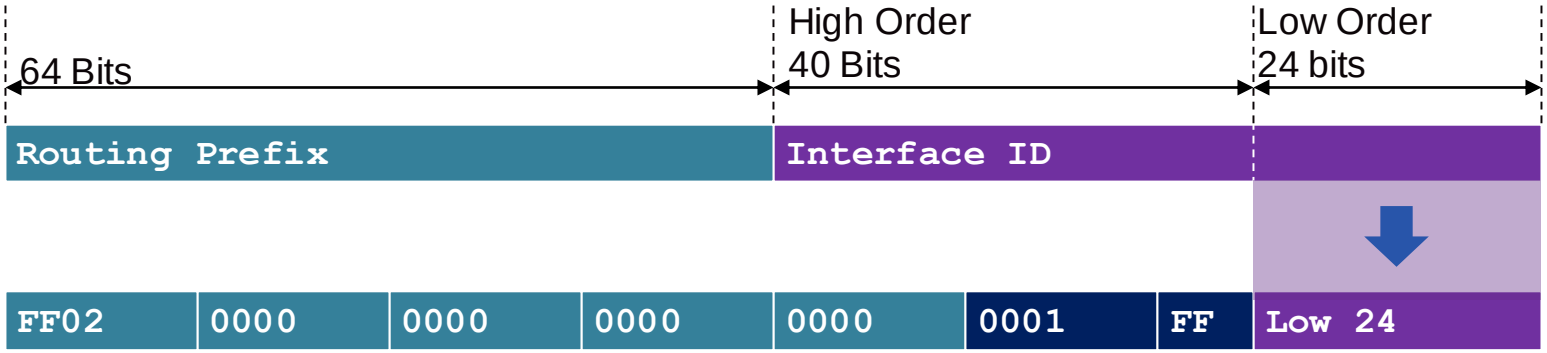

Solicited Node Multicast Address Example

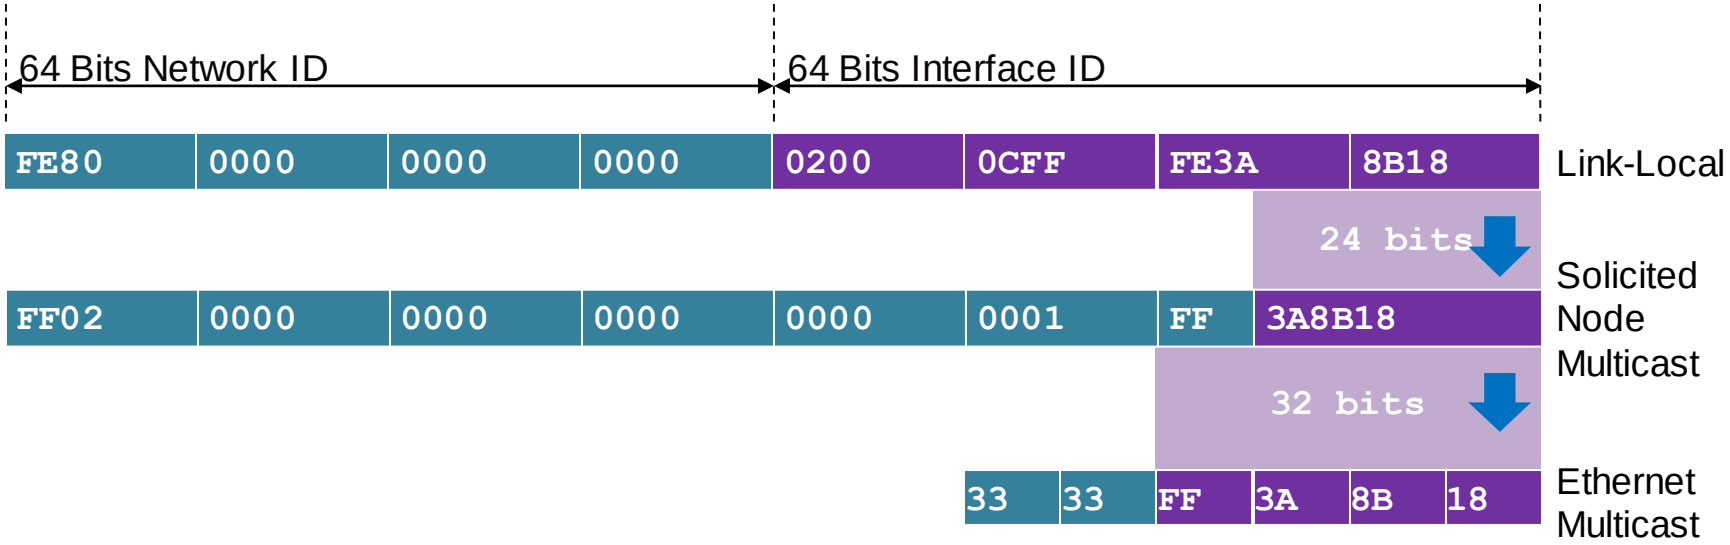

IPv6 Interface Example

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.

Link-local address (FE80::)

All Nodes

All Routers

Solicited Node Multicast Address

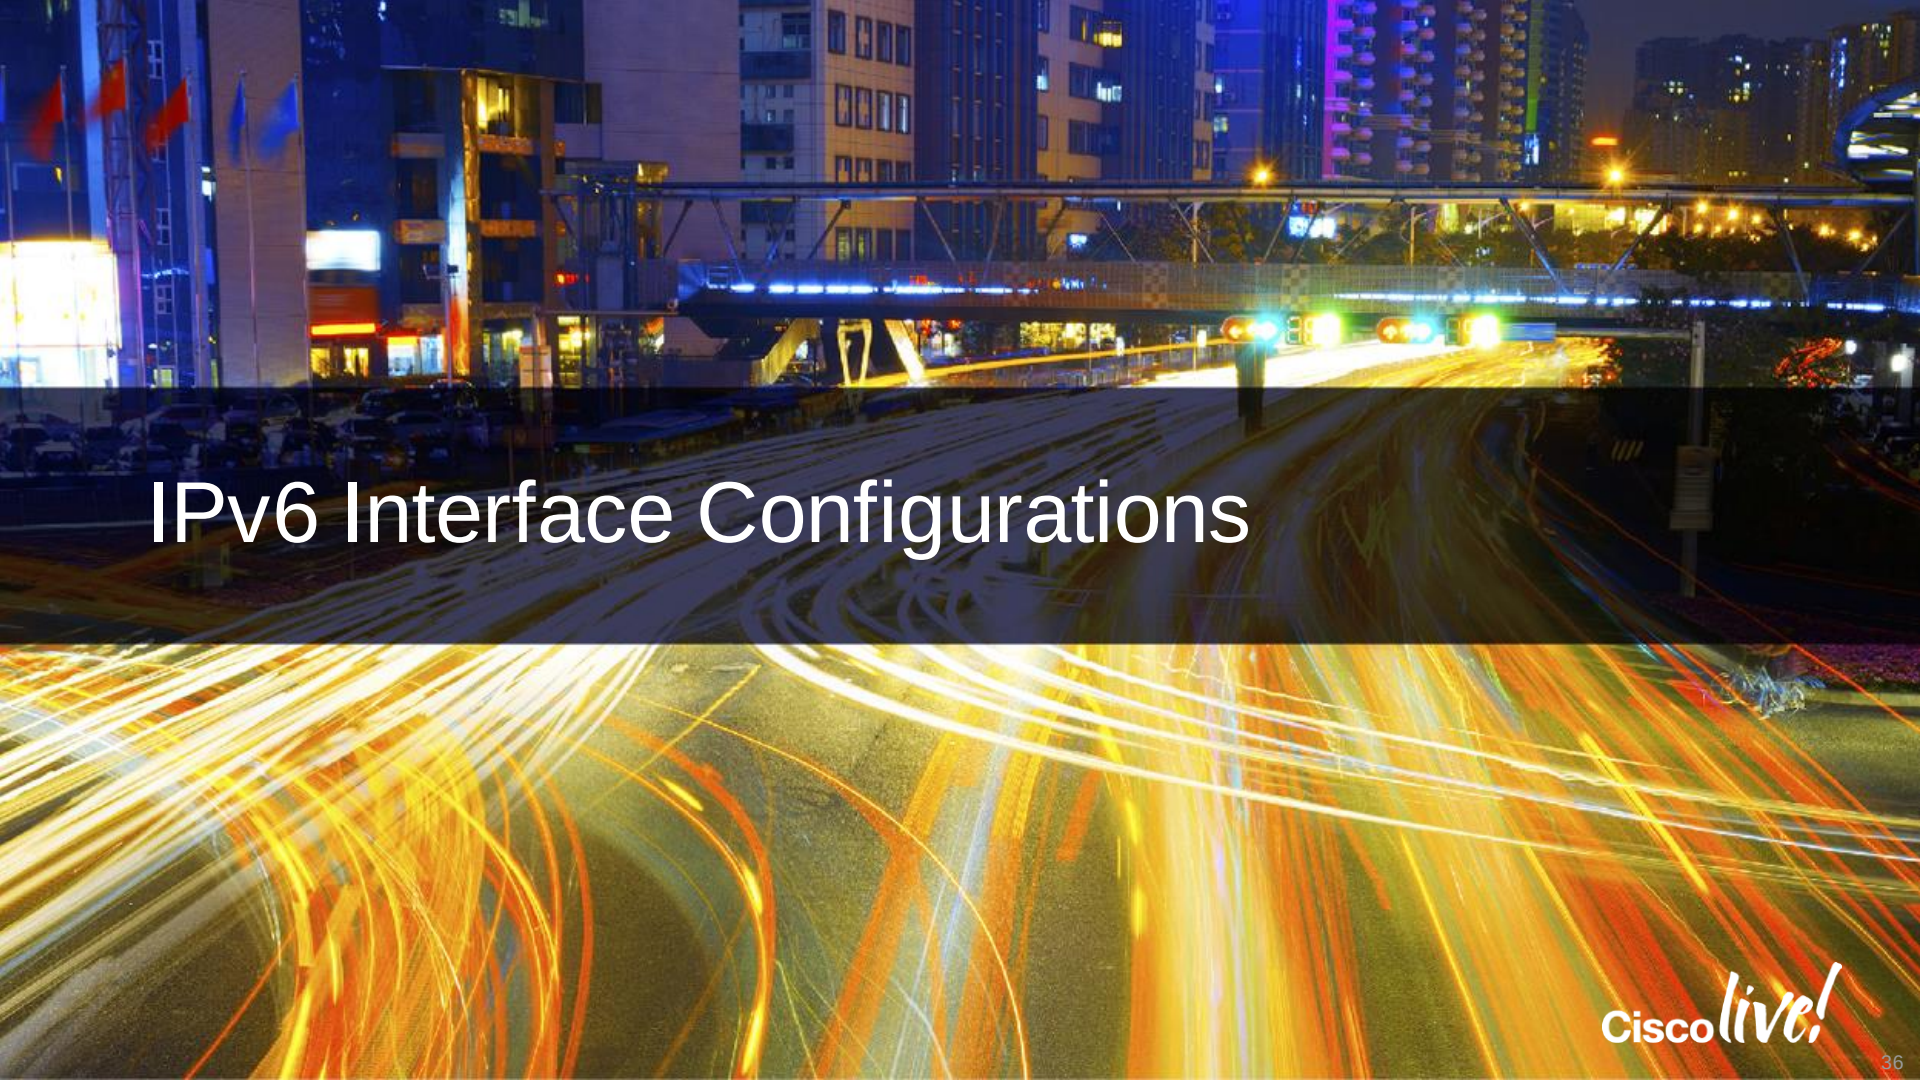

IPv6 Interface Configurations

Link-Local Configured Interface Identifier Address (IOS)

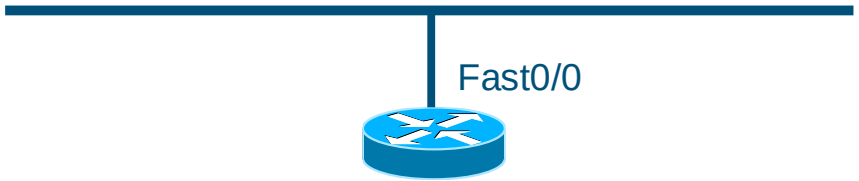


```
ipv6 unicast-routing
!
interface FastEthernet0/0
 ip address 10.151.1.1 255.255.255.0
 duplex auto
 speed auto
 ipv6 enable
!
```

Enable IPv6 routing

Enable IPv6 on interface and automatically create link-local address

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
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)
```

EUI-64 derived from MAC address 0007.505e.9460

Listening for all hosts multicast

Listening for all routers multicast

Solicited Node multicast for link-local address

MAC address 0007.505e.9460

Manually Configured Interface Identifier Address

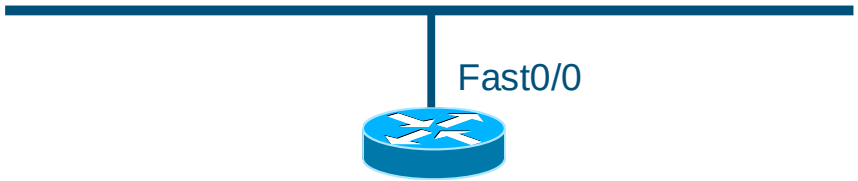


```
ipv6 unicast-routing
!
interface FastEthernet0/0
 ip address 10.151.1.1 255.255.255.0
 duplex auto
 speed auto
 ipv6 address 2001:db8::1/64
!
```

Enables IPv6 and assigns a global prefix and manual interface ID

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):
    2001:db8::1, subnet is 2001:db8::/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
Hosts use stateless autoconfig for addresses.
```

Routable /64 subnet

Global unicast address with manual interface ID of “1”

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

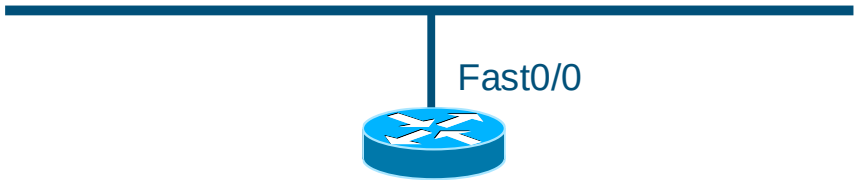


```
ipv6 unicast-routing
!
interface FastEthernet0/0
 ip address 10.151.1.1 255.255.255.0
 duplex auto
 speed auto
 ipv6 address 2001:db8::/64 eui-64
!
```

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):
    2001:db8::207:50FF:FE5E:9460, subnet is 2001:db8::/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
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

ICMPv6 and Neighbour Discovery

ICMPv6

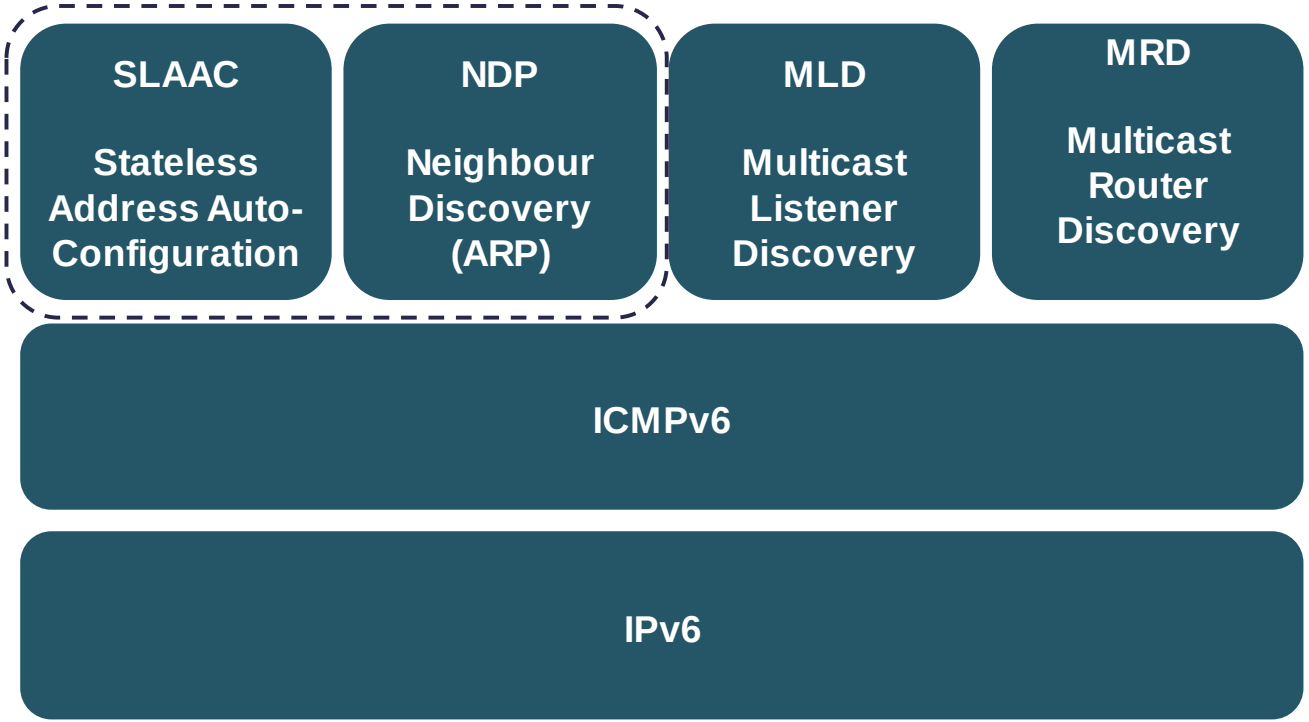

ICMPv6 Header

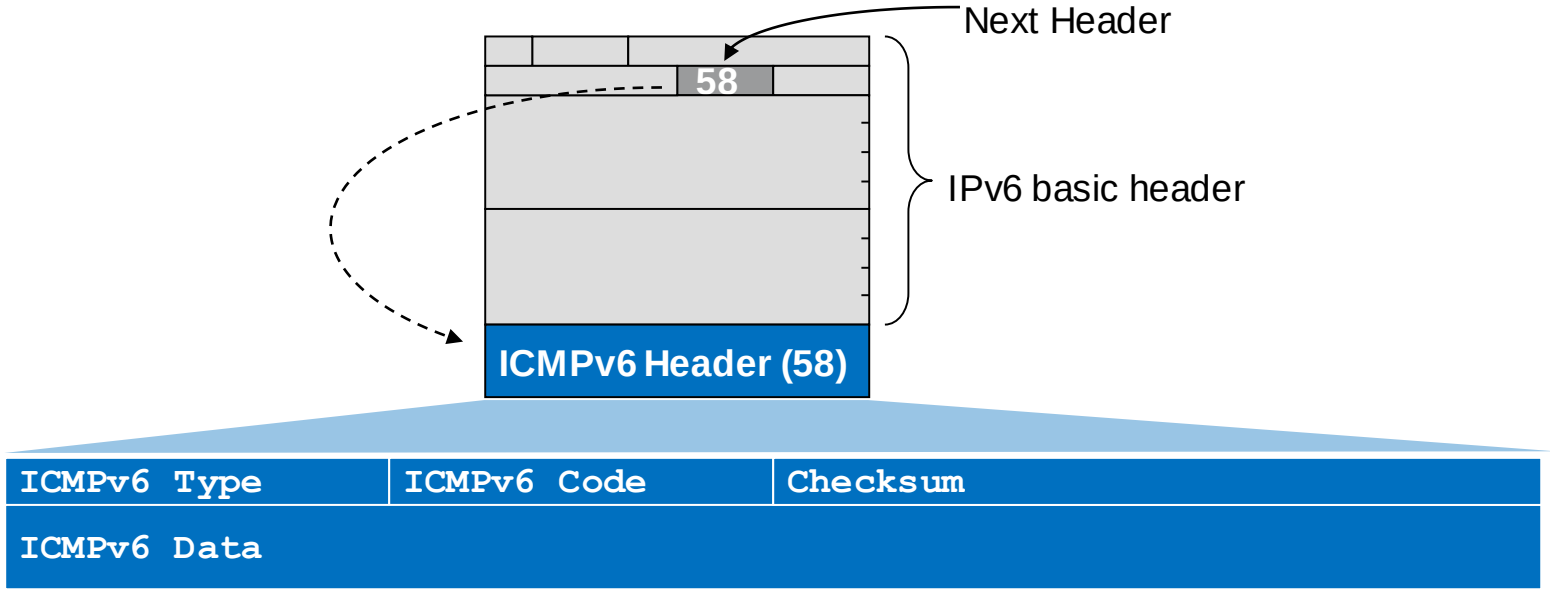

- Also used for Neighbour Discovery, Path MTU discovery and Multicast Listener Discovery (MLD)

Neighbour Discovery Messages (ND)

Reference

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 nodes
Neighbour Solicitation (NS)	Request link-layer of target	135	Node	Solicited Node Target Node
Neighbour Advertisement (NA)	Response to NS (solicited) Advertise link-layer address change (Unsolicited)	136	Nodes	Sender of NS

Cisco live!

Router Solicitation and Advertisement (RS & RA)

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

Router Advertisement	
ICMP Type	134
IPv6 Source	Link Local (FE80::2)
IPv6 Destination	Sender of RS 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

Neighbour Solicitation & Advertisement (NS & NA)

Neighbour Solicitation	
ICMP Type	135
IPv6 Source	A Unicast
IPv6 Destination	B Solicited Node Multicast
Target / Options	B Unicast / 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

Neighbour Cache Entry States

Reference

- INCOMPLETE
  - Address resolution is in progress and the link-layer address of the neighbour has not yet been determined
- REACHABLE
  - The neighbour is known to have been reachable recently (within tens of seconds ago)
- STALE
  - The neighbour is no longer known to be reachable but until traffic is sent to the neighbour, 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 neighbour is no longer known to be reachable, and unicast Neighbour Solicitation probes are being sent to verify reachability

DAD Example

Reference

Tentative IP
FE80::260:8FF:FE52:F9D8

Actual IP
FE80::260:8FF:FE52:F9D8

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

NA	
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
NA Header	
Target Address	FE80::260:8FF:FE52:F9D8
Neighbour Discovery Option	
Target MAC	00-60-08-52-F9-D8

Fragmentation and Path MTU Discovery

Fragmentation in IPv6

- Unfragmentable part
  - IPv6 header plus any headers that must be processed by the nodes en-route
  - 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
- Minimum MTU for IPv6 is 1280 bytes
  - All links MUST support it

Fragment Header

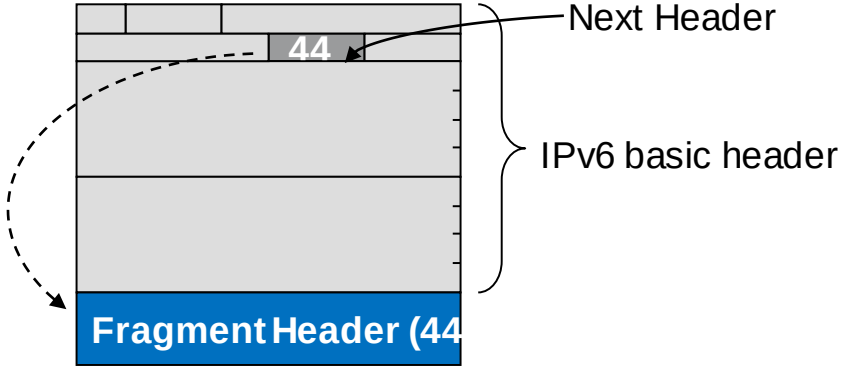

Next Header	Reserved	Fragment Offset	00	M
Identification				
Fragment Data				

- 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

Path MTU Discovery

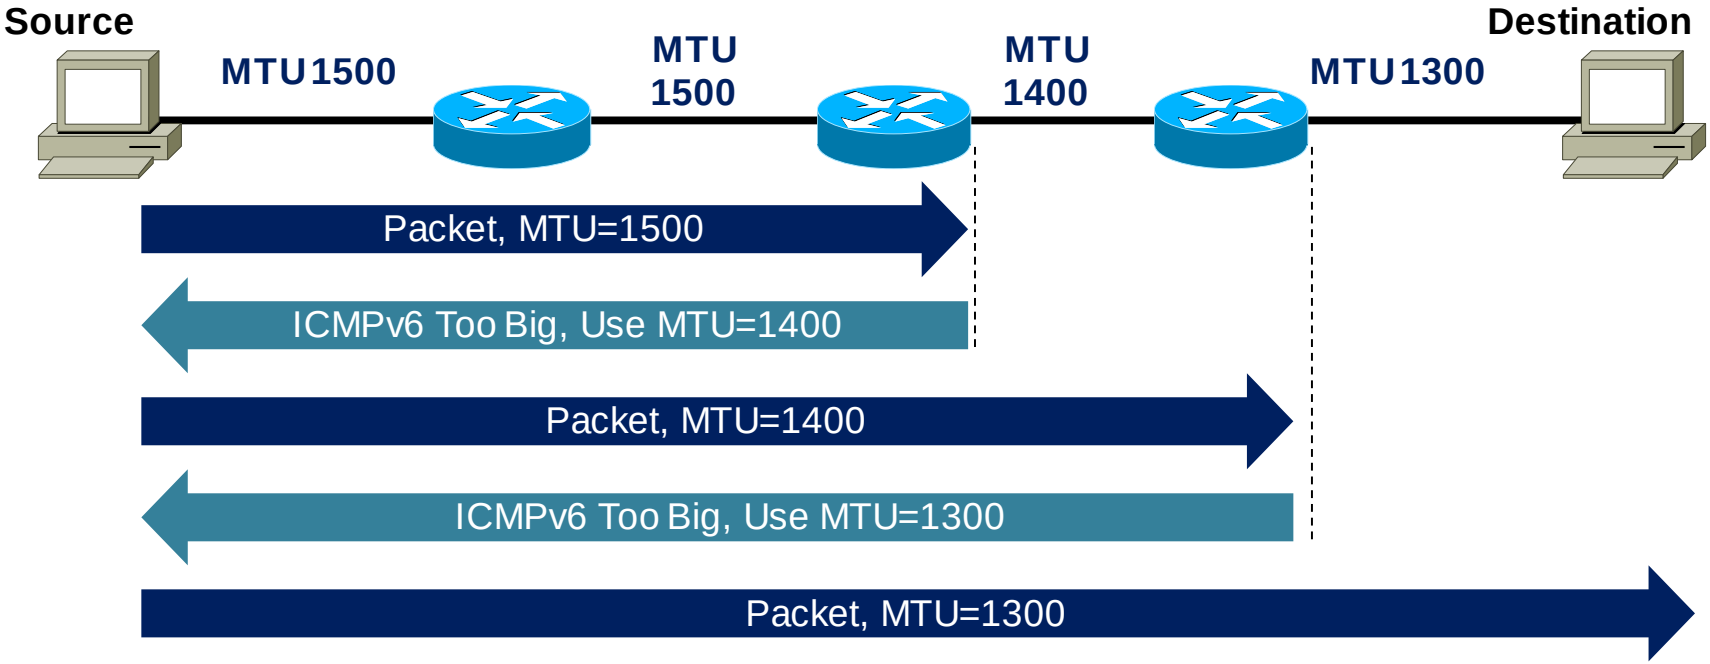

- Store PMTU per destination (if received)
- Age out PMTU (10 mins), reset to first link MTU


Host Address Assignment

IPv6 Host Address Assignment Methods

- Manual Assignment
- Stateless Address Autoconfiguration (SLAAC RFC 4862)
  - Allows auto assignment of address
- Stateful DHCPv6 (RFC 3315)
  - Allows DHCPv6 to allocate IPv6 address plus other configuration
- DHCPv6-PD (RFC 3633)
  - Allows DHCPv6 to allocate entire subnets to a router/CPE device
- Stateless DHCPv6 (RFC 3736)
  - SLAAC for host address allocation and DHCPv6 for other configuration

Stateless Address Autoconfiguration (RFC4862)

- SLAAC is used to automatically assigned an address to a host “plug and play”

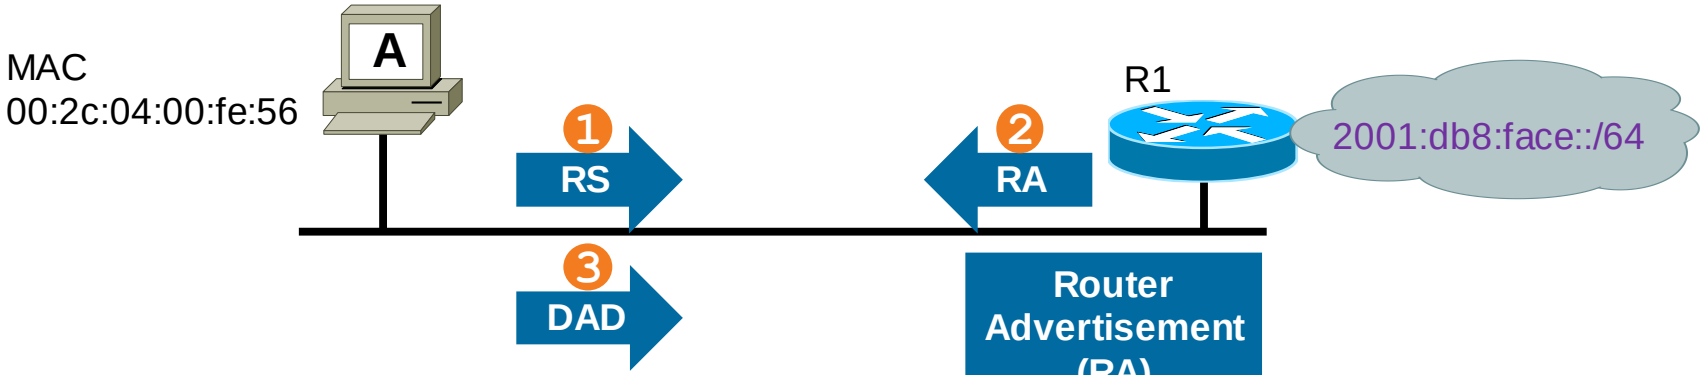

Router Advertisement (RA)	
Ethernet DA/SA	Router R2 / Host A
Prefix Information	2001:db8:face::/64
Default Router	Router R1

Address comprises
Prefix Received + Link-Layer

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

Router Advertisement for Stateful DHCPv6

- RA message contain flags that indicate address allocation combination (A, M and O bits)

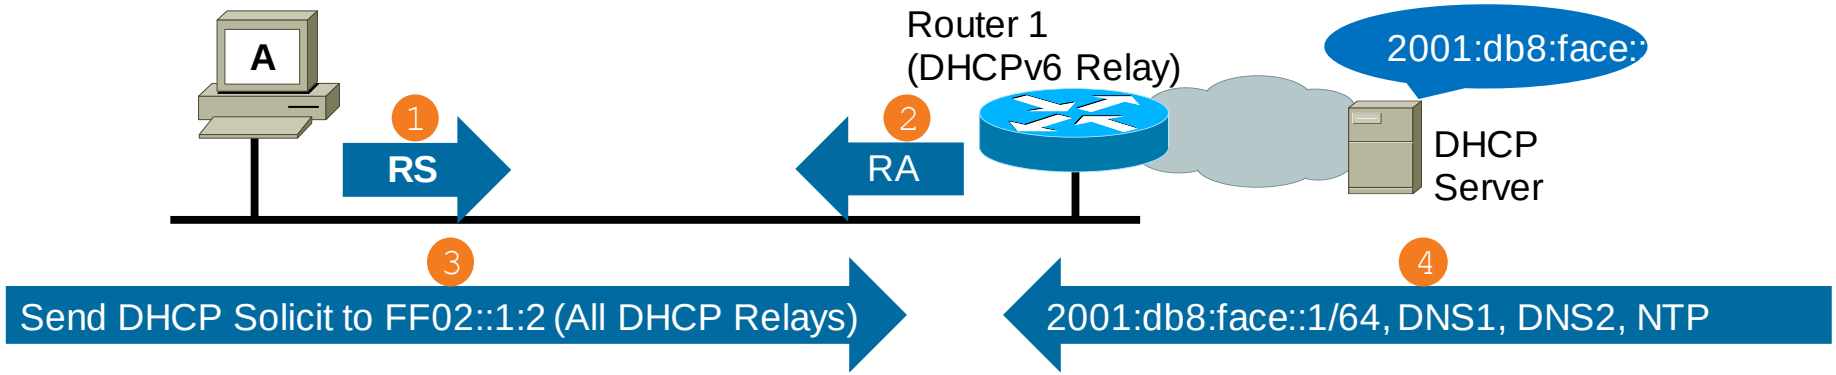

RA	
A bit (Address config flag)	Set to 0 - Do not use SLAAC for host config
M bit (Managed address configuration flag)	Set to 1 - Use DHCPv6 for host IPv6 address
O bit (Other configuration flag)	Set to 1 - Use DHCPv6 for additional info (DNS, NTP)

Router Advertisement for Stateless DHCPv6

- RA message contain flags that indicate address allocation combination (A, M and O bits)

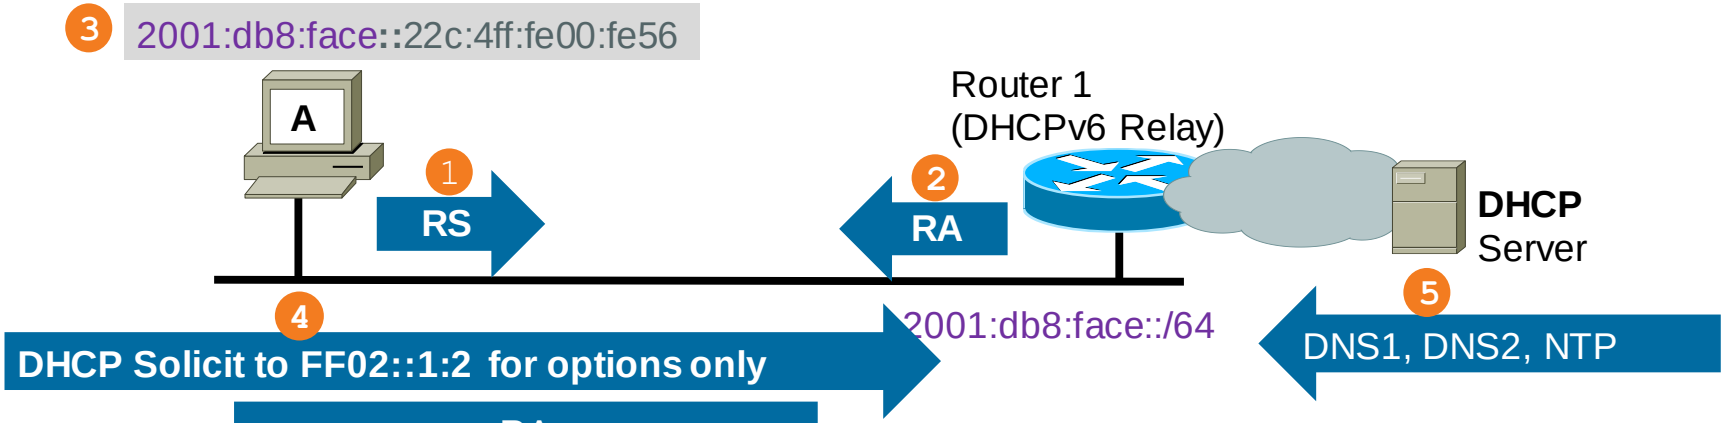

RA	
A bit (Address config flag)	Set to 1 - Use SLAAC for host address config
On-link Prefix	2001:db8:face::/64
M bit (Managed address configuration flag)	Set to 0 - Do not use DHCPv6 for IPv6 address
O bit (Other configuration flag)	Set to 1 - Use DHCPv6 for additional info (DNS, NTP)

DHCPv6 Configuration Options

A bit (default) just use SLAAC

```
interface e0/0
 ipv6 address 2001:db8:1000::1/64
```

Host gets address and other SLAAC options. Nothing else

M bit & O bit (Stateful DHCP)

```
interface e0/0
 ipv6 address 2001:db8:1000::1/64
 ipv6 nd managed-config-flag
 ipv6 nd other-config-flag
 ipv6 dhcp relay destination 2001:db8::10
```

Host gets full stateful config from DHCP server (2001:db8::10)

A bit & O bit (Stateless DHCP)

```
interface e0/0
 ipv6 address 2001:db8:1000::1/64
 ipv6 nd other-config-flag
 ipv6 dhcp relay destination 2001:db8::10
```

Host get address from SLAAC and other config from DHCP server (2001:db8::10)

Default Router Selection

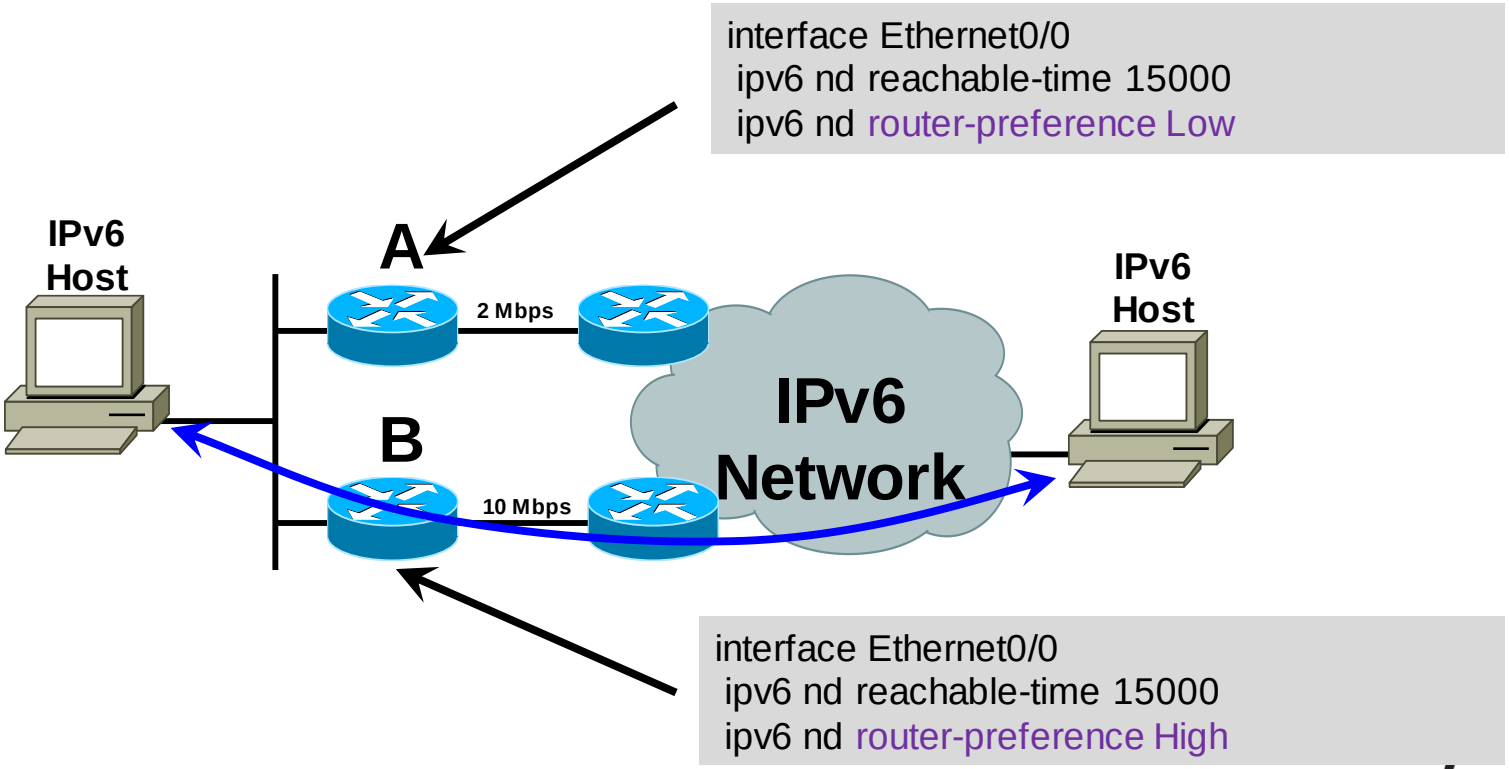

Domain Name System (DNS)

IPv6 and DNS Entries

Function	IPv4	IPv6
Hostname to IP Address	A Record www.abc.test. IN A 92.168.30.1	AAAA Record (Quad A) www.abc.test. IN AAAA 2001:db8:C18:1::2
IP Address To Hostname	PTR Record 1.30.168.192.in-addr.arpa. IN PTR www.abc.test.	PTR Record 2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0.8.1.c.0.8.b.d.0.1.0.0.2.ip6.arpa IN PTR www.abc.test.

Dual Stack Approach and DNS

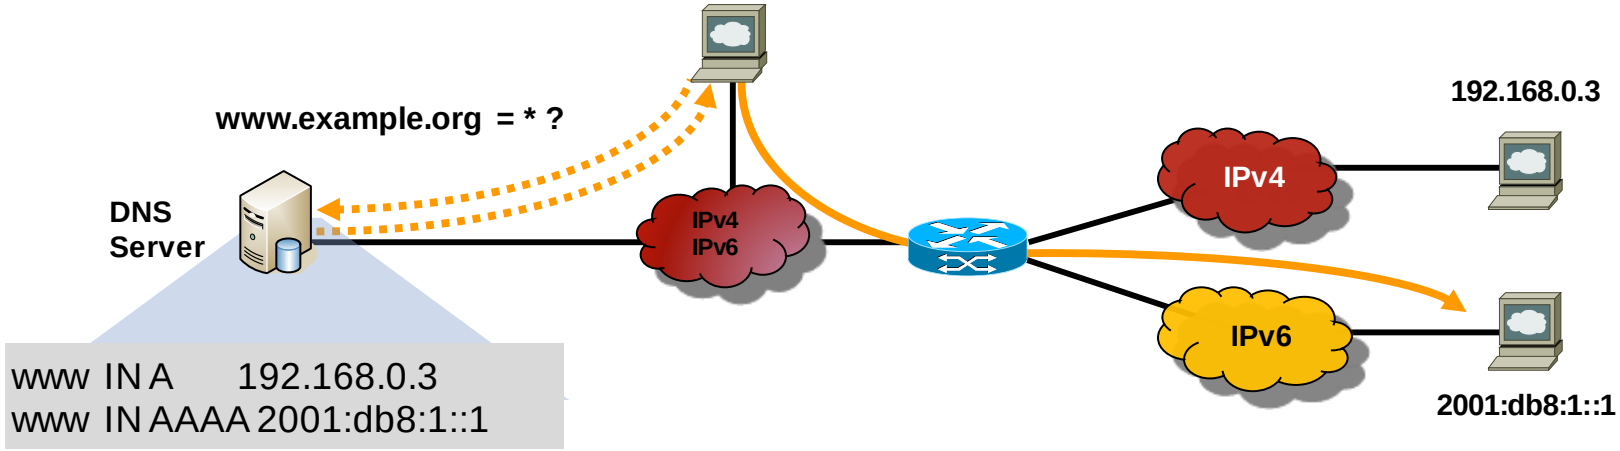

- In a dual stack case an application that:
  - Is IPv4 and IPv6-enabled
  - Can query the DNS for IPv4 and/or IPv6 records (A) or (AAAA) records
  - Chooses one address and, for example, connects to the IPv6 address


Routing IPv6 – How To Tell Everyone About It!

Overview of Routing Protocols in IPv6

- Routing in IPv6 is unchanged from IPv4
  - Still has two families of routing protocols: IGP and EGP
  - Still uses the longest-prefix match routing algorithm
- IGP
  - RIPng (RFC 2080)
  - Cisco EIGRP for IPv6
  - Integrated IS-IS for IPv6 (RFC 5308)
  - OSPFv3 (RFC 5340)
- EGP
  - MP-BGP4 (RFC 4760) and Using MP-BGP for IPv6 (RFC 2545)
- Cisco IOS supports all IPv6 routing protocols

Static Routing

- Similar to IPv4
- Next hop / interface is required

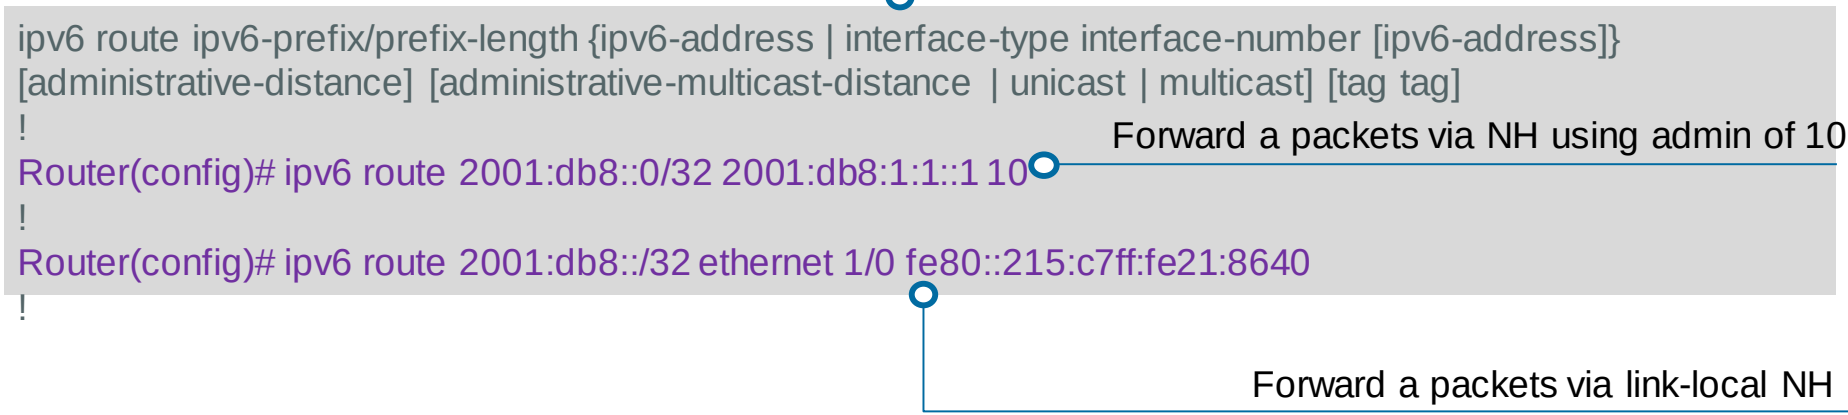

Default Routing Example

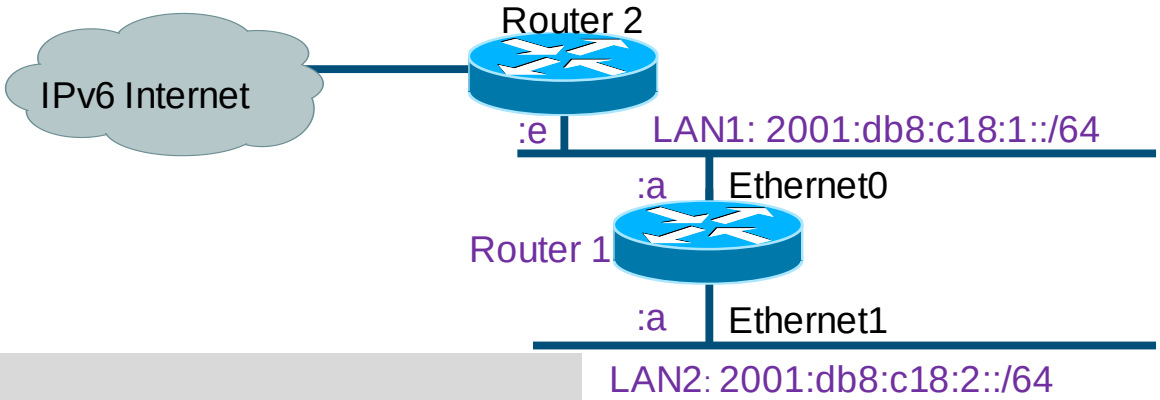


```
router 1#
!
ipv6 unicast-routing
!
interface Ethernet0
 ipv6 address 2001:db8:c18:1::a/64
!
interface Ethernet1
 ipv6 address 2001:db8:c18:2::a/64
!
ipv6 route ::/0 2001:db8:c18:1::e
```

Default route to Router 2

EIGRP for IPv6 Features

- Three new TLVs introduced
- Hello messages use FF02::A (all EIGRP routers)
- Automatic summarisation is disabled by default for IPv6 (unlike IPv4)
- Check “shutdown” mode
- RID stays at 32 bits

EIGRP for IPv6 Configuration

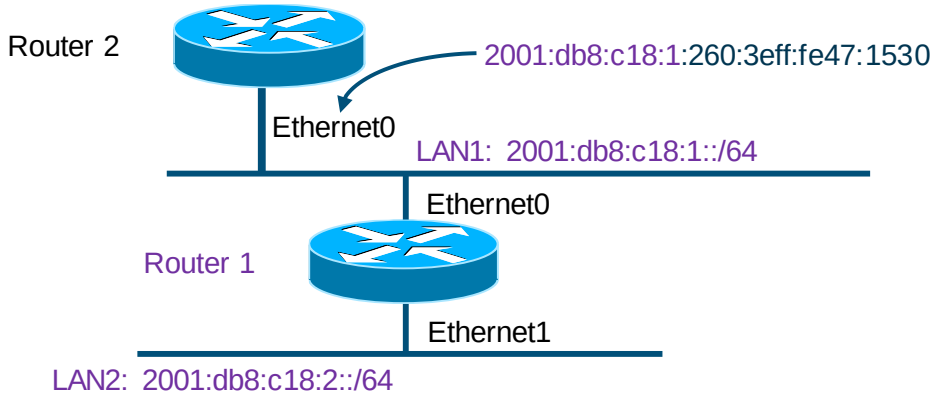


```
Router2#
!
ipv6 router eigrp 100
 eigrp router-id 10.10.10.1
!
interface Ethernet0
 ipv6 address 2001:db8:c18:1::/64  eui-64
 ipv6 eigrp 100
!
```

```
Router1# show ipv6 eigrp neighbor
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   FE80::260:3eff:fe47:1530  E0              14 00:01:43    1     4500  0   1

Router1# show ipv6 eigrp topology all-links

P 2001:db8:c18:1::/64, 1 successors, FD is 28160, serno 1
        via Connected, Ethernet0
        via FE80::260:3eff:fe47:1530 (30720/28160), Ethernet0
```

Neighbours and next hops are identified by link-local address

OSPFv3 Overview

- OSPFv3 is OSPF for IPv6 (RFC 5340)
- Based on OSPFv2 with enhancements
- Distributes IPv6 prefixes only
- Ships-in-the-night with OSPFv2
- No in-protocol Authentication

OSPFv3 Configuration Example

Classic IOS syntax

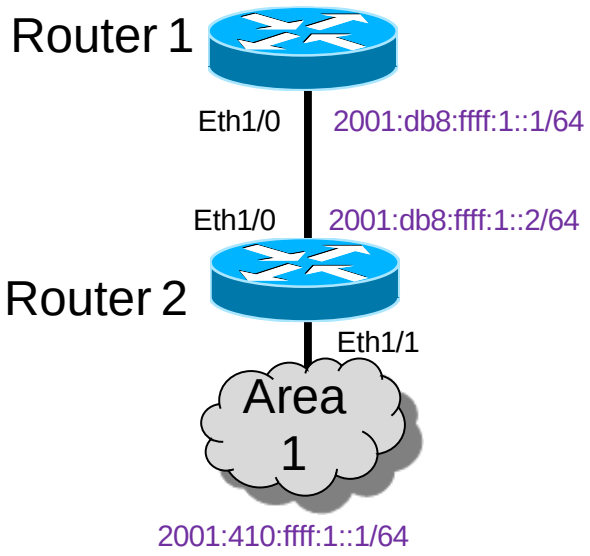


```
Router1#
interface Ethernet1/0
 ipv6 address 2001:db8:ffff:1::1/64
 ipv6 ospf 100 area 0
!
ipv6 router ospf 100
  router-id 10.1.1.3
!
```

Interlink connection

OSPFv3 process

```
Router2#
interface Ethernet1/0
 ipv6 address 2001:db8:ffff:1::2/64
 ipv6 ospf 100 area 0
!
interface Ethernet1/1
 ipv6 address 2001:db8:cafe::1/48
 ipv6 ospf 100 area 1
!
ipv6 router ospf 100
  router-id 10.1.1.4
```

Interlink connection

Enables IPv6 facing Area 1

32 bit ID specified in dotted decimal notation

OSPFv3 Configuration Example

Unified IOS syntax

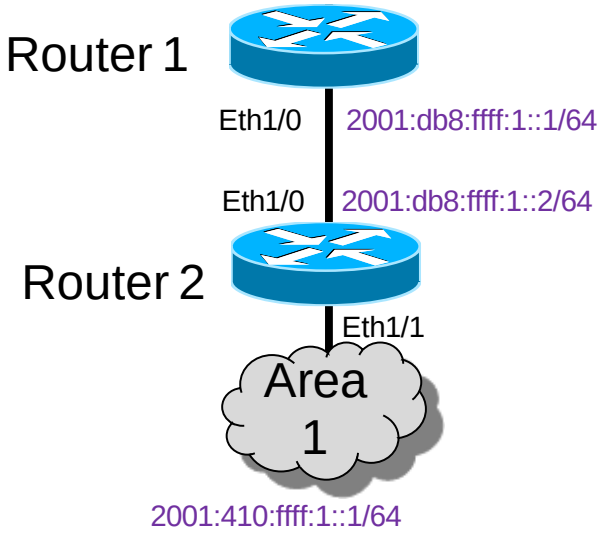


```
Router1#
interface Ethernet1/0
 ipv6 address 2001:db8:ffff:1::1/64
 ospfv3 100 area 0 ipv6
!
router ospfv3 100
  router-id 10.1.1.3
!
```

Interlink connection

OSPFv3 process

```
Router2#
interface Ethernet1/0
 ipv6 address 2001:db8:ffff:1::2/64
ospfv3 100 area 0 ipv6
!
interface Ethernet1/1
 ipv6 address 2001:db8:cafe::1/48
ospfv3 100 area 1 ipv6
!
router ospfv3 100
  router-id 10.1.1.4
```

Interlink connection

Enables IPv6 facing Area 1

32 bit ID specified in dotted decimal notation

MP-BGP for IPv6 Overview

- TCP Interaction
  - BGP-4 runs over a TCP (179) session using IPv4 or IPv6
  - The NLRI BGP carried (IPv4, IPv6, MPLS) is agnostic of the session protocol
- Router ID
  - BGP router-id must still exist is in 32 bit dotted decimal notation
- Next-hop contains a global IPv6 address (or potentially a link local address)
- Link local address as a next-hop is only set if the BGP peer shares the subnet with both routers (advertising and advertised)

BGP IPv6 Configuration Global Address Peering

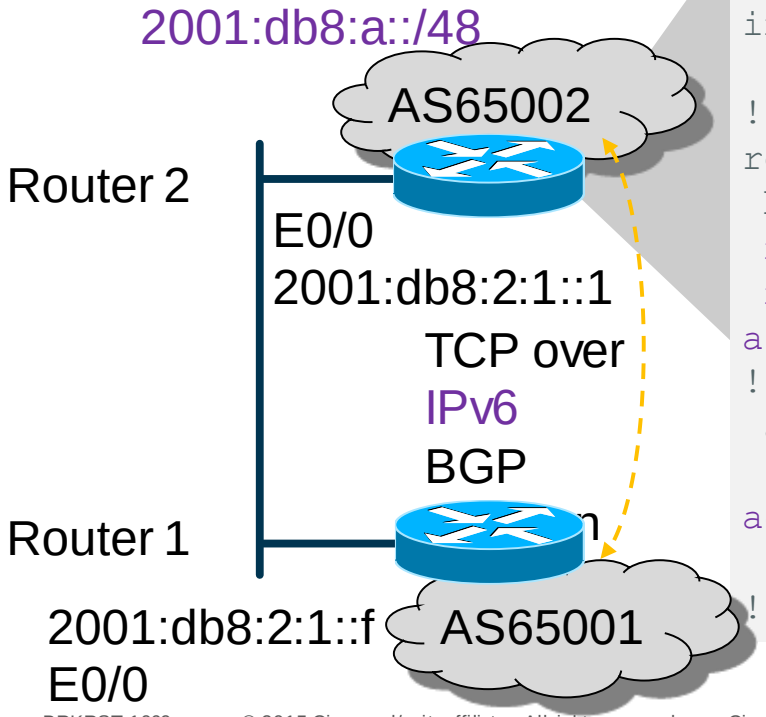


```
Router2#
!
interface Ethernet0/0
 ipv6 address 2001:db8:2:1::1/64
!
router bgp 65002
 bgp router-id 10.10.10.1
 no bgp default ipv4-unicast
 neighbor 2001:db8:2:1::f remote-
as 65001
!
 address-family ipv6
  neighbor 2001:db8:c18:2:1::f
activate
  network 2001:db8:a::/48
!
```

Router ID in dotted decimal notation

Disable default IPv4 behaviour

Use IPv6 address family

Activate IPv6 session

IPv6 prefix to be advertised

Cisco live!


Putting It All Together

Dual Stack Configuration

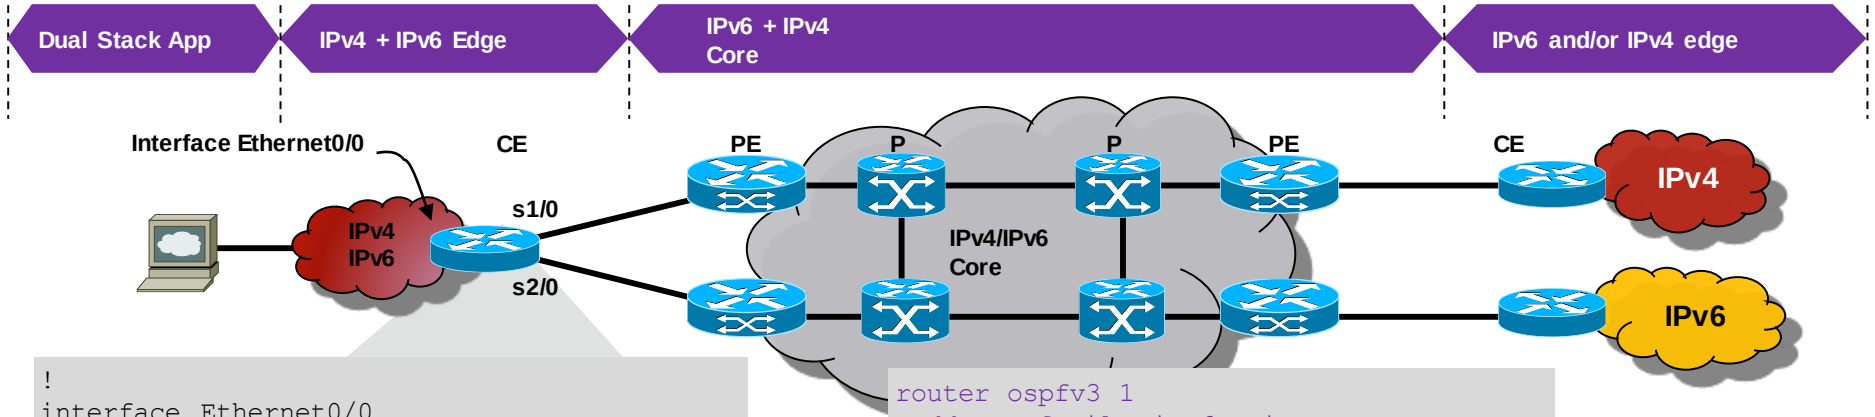


```
!
interface Ethernet0/0
 ip address 192.168.99.1 255.255.255.0
 ipv6 address 2001:db8:213:1::1/64
 ospfv3 1 area 0 ipv6
!
interface Serial 1/0
 ip address 192.0.2.1 255.255.255.252
 ipv6 address 2001:db8:ffff:1::1/64
!
interface Serial 2/0
 ip address 192.0.2.5 255.255.255.252
 ipv6 address 2001:db8:ffff:2::1/64
!
```

```
router ospfv3 1
 address-family ipv6 unicast
 exit-address-family
!
router bgp 65000
...
address-family ipv4
neighbor 192.0.2.2 activate
neighbor 192.0.2.6 activate
!
address-family ipv6
neighbor 2001:db8:ffff:1::2 activate
neighbor 2001:db8:ffff:2::2 activate
...
```


Q & A
Cisco live!

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2015 T-Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site http://showcase.genie-connect.com/clmelbourne2015
- Visit any Cisco Live Internet Station located throughout the venue

T-Shirts can be collected in the World of Solutions on Friday 20 March 12:00pm - 2:00pm

Learn online with Cisco Live!
Visit us online after the conference for full access to session videos and presentations. www.CiscoLiveAPAC.com

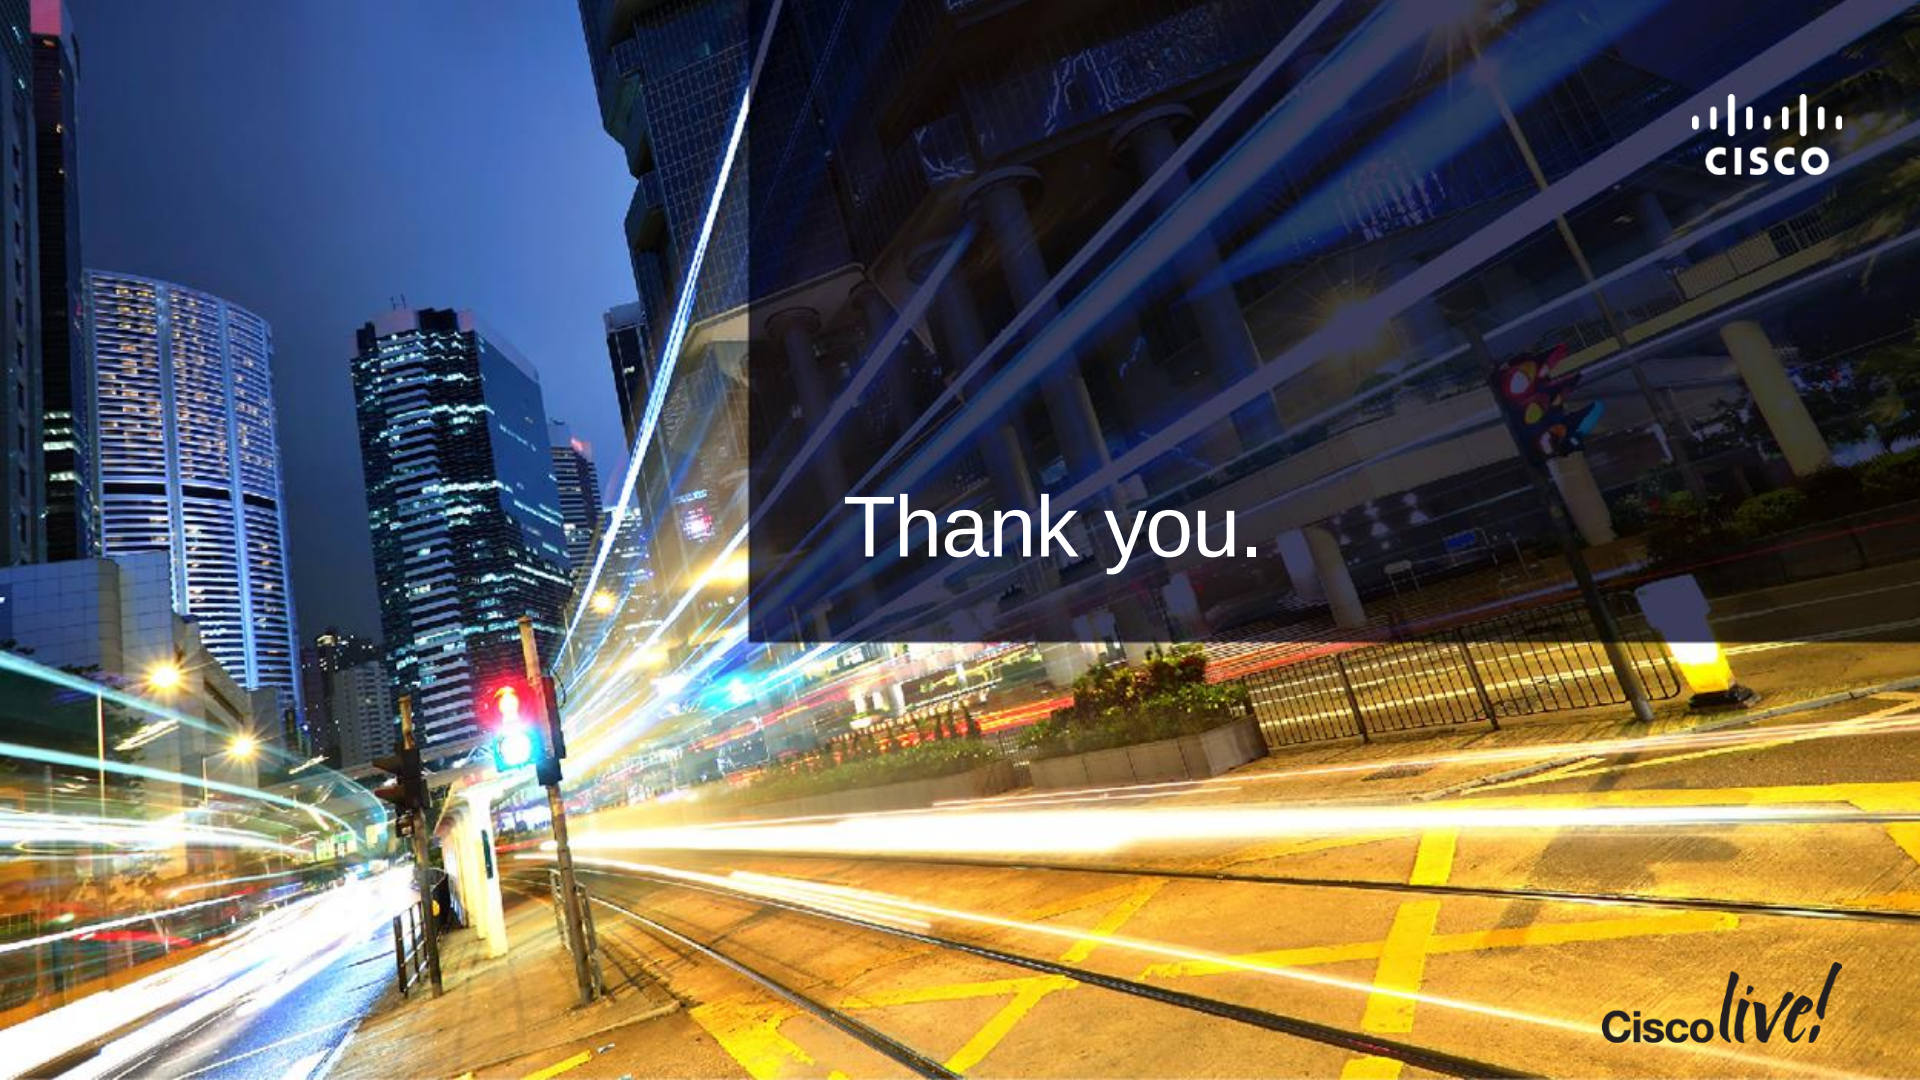
Thank you.
CISCO
Cisco live!


CISCO