You make **possible**

Introduction to IPv6:

Connecting to IPv6 access networks

Nicole Wajer, Tim Martin

TECRST-1991

Barcelona | January 27-31, 2020

Cisco Webex Teams

Questions?

Use Cisco Webex Teams to chat
with the speaker after the session

How

1. Find this session in the Cisco Events Mobile App
2. Click “Join the Discussion”
3. Install Webex Teams or go directly to the team space
4. Enter messages/questions in the team space

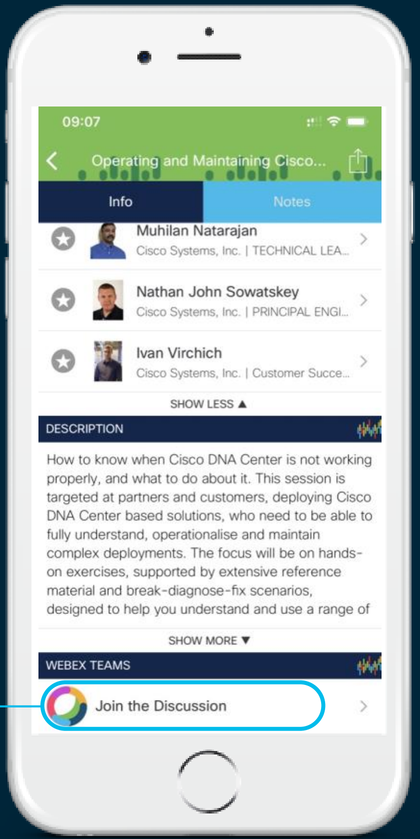

Agenda

- IPv6 the Protocol
- IPv6 Address Planning
- IPv6 Access Layer Design
- Conclusion

IPv6:

The Protocol

Tim Martin

TECRST-1991

Follow

@bckcntryskr

CISCO CCIE EMERITUS #2020

IPv6 FORUM CERTIFIED CERTIFICATION

Hall of Fame Speaker Cisco live!

ciscopress.com Author

IPv6 Node Types

- Node: Any device that implements an IPv6 protocol stack
- Host: A device with one or more interfaces participating in IPv6 network
- Router: A device that forwards packets and provides provisioning services

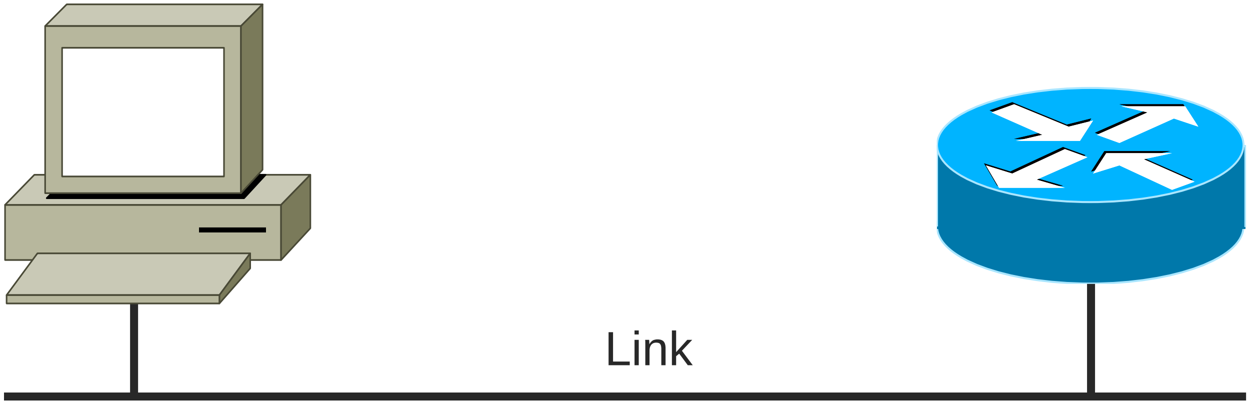

Address’s & Headers

IPv6 Addressing

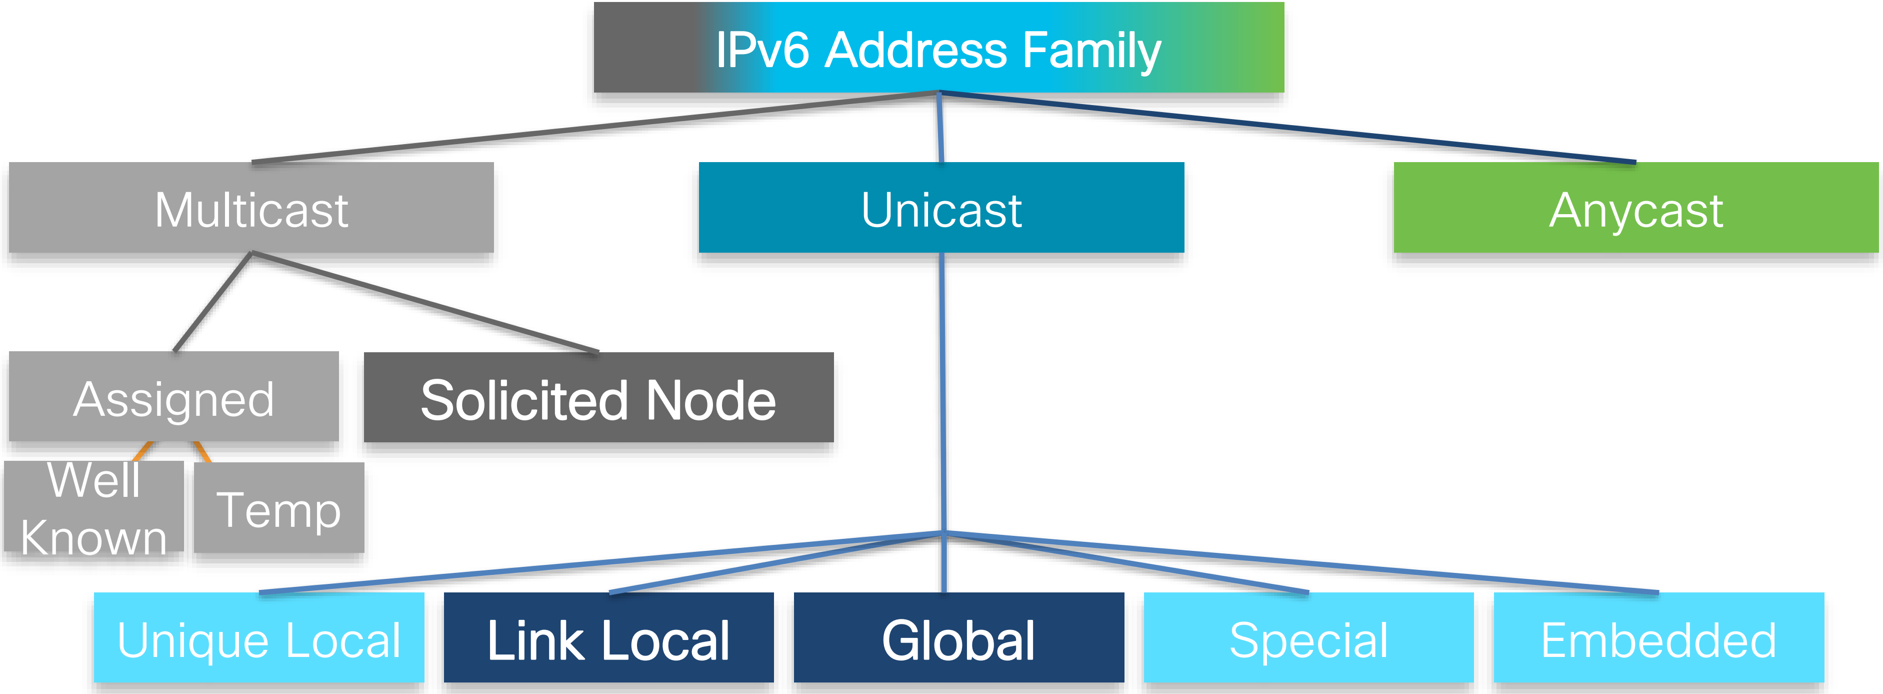

*IPv6 does not use broadcast addressing

Hexadecimal Is Really Not That Difficult

- Widely used in computing and programming
  - Hex is a base 16 numerical system
  - Typically expressed by 0x, i.e. 0x34
- Every nibble is a Hex character
  - 4 bits have 16 combinations
  - Easier than high school algebra

100s	10's	1's
0	5	2
1	7	2
5	8	9

256's	16's	1's
	3	4
	a	c
2	4	d

Binary	Hex	Decimal
0000	0	0
0001	1	1
0010	2	2
0011	3	3
0100	4	4
0101	5	5
0110	6	6
0111	7	7
1000	8	8
1001	9	9
1010	A	10
1011	B	11
1100	C	12
1101	D	13
1110	E	14
1111	F	15

IPv6 Address Format

- IPv6 addresses are 128 bits long (32 hex characters)
  - 8 groups (words, quads) of 16 bits separated by (:)
  - RFC5952 – lower case, leading zeros, zero compression

2001:0db8:0046:a1d1:0000:0000:0000:0001

↓

2001:db8:46:a1d1:0:0:0:1

↓

2001:db8:46:a1d1::1

Global Route Prefix | **Subnet Id** | **Interface Id**

2001	:	0db8	:	0046	:	a1d1	:	0000	:	0000	:	0000	:	0001
16 bits		16 bits		16 bits		16 bits		16 bits		16 bits		16 bits		16 bits

Explaining BIG Numbers With Math

- The standard LAN size has been set at a /64
  - 18,446,744,073,709,600,000 IPv6 addresses
- Let’s attempt to exhaust all of the available addresses
  - We will allocate 10,000,000 addresses per second
  - Hint: there are 31,536,000 seconds per year
  - 10,000,000 x 31,536,000 = 315,360,000,000,000

18,446,744,073,709,600,000

/ 315,360,000,000,000

= 58,494 years

Attribution: Ed Horley

IPv6 Unicast Address Types

Link-Local – Non routable exists on a layer 2 domain (**fe80**::/10)

fe80:0000:0000:0000:xxxx:xxxx:xxxx:xxxx

Unique-Local – Routable within administrative domain (**fc**00::/7)

fc00:~~~~:~~~~:**:**xxxx:xxxx:xxxx:xxxx

fd00:~~~~:~~~~:**:**xxxx:xxxx:xxxx:xxxx

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

2000:**:****:****:**xxxx:xxxx:xxxx:xxxx

3fff:**:****:****:**xxxx:xxxx:xxxx:xxxx

IPv6 over Ethernet

- IPv6 has a specific Ether type id
- IPv6 relies heavily on Multicast

Destination Ethernet Address	Source Ethernet Address	0x0800	IPv4 Header and Payload
Destination Ethernet Address	**Source Ethernet Address**	**0x86DD**	**IPv6 Header and Payload**

33	33	xx	xx	xx	xx

0011 00UG

0 = Universal/unique

1 = Local/not unique

I/G bit = Multicast/Broadcast

U/L bit = Universal/Local

IPv4 and IPv6 Header Comparison

IPv4 Header (20-60)

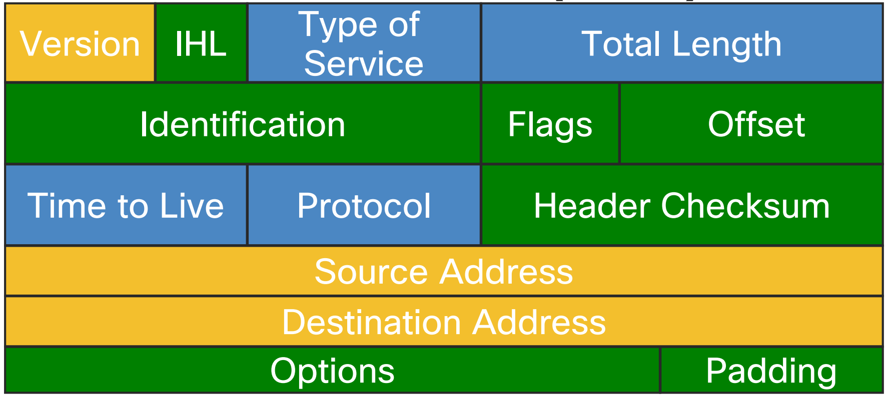

IPv6 Header (40)

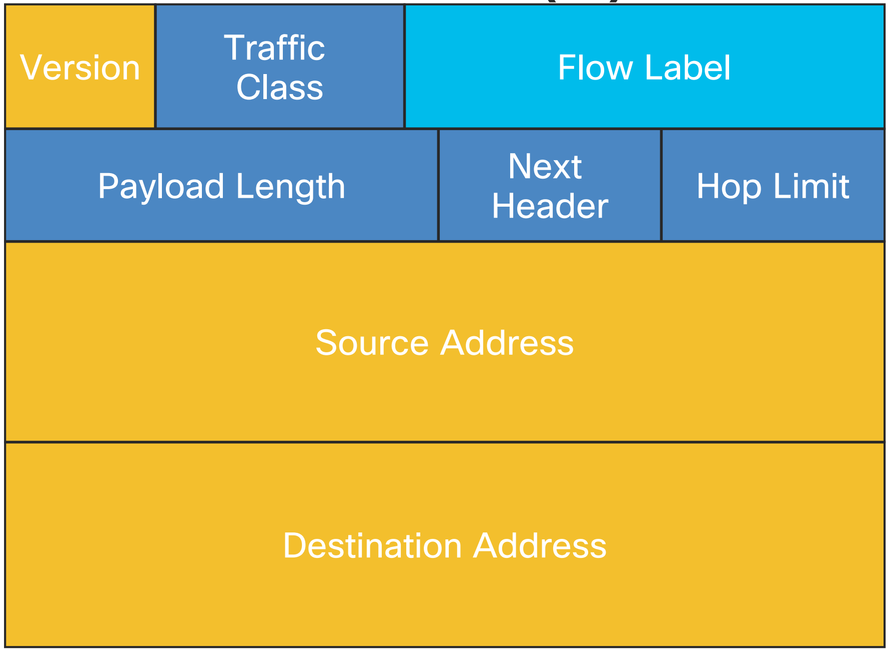

- Length was variable
- Fields in green are removed
- Options appear in extension headers
- Upper layer checksums use Pseudo Header format: SRC/DST + Next Header

Extension Headers (~ Layer 3.5)

- EH are daisy chained, processed in order
- Length is variable, end on 64-bit boundary
- EHs have a Next Header field
- All EHs must be in the initial fragment

Extension Header	Type
Hop-by-Hop Options	0
Destination Options	60
Routing Header	43
Fragment Header	44
Authentication Header	51
ESP Header	50
Destination Options	60
Mobility Header	135
Experimental	253,254
No Next Header	59

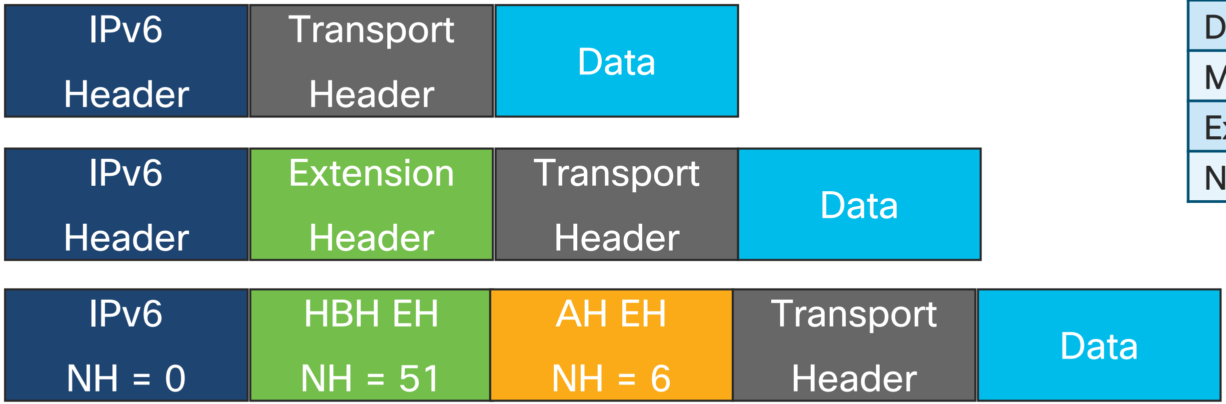

ICMPv6 Messages

- Neighbor or router discovery (133-137)
- Multicast Listener Discovery (130-132, 143)
- Diagnostics using Ping, Traceroute (128, 129)
- Destination Unreachable(1)
- Packet Too Big (2)
- Time Exceeded (3)
- Parameter Problem (4)

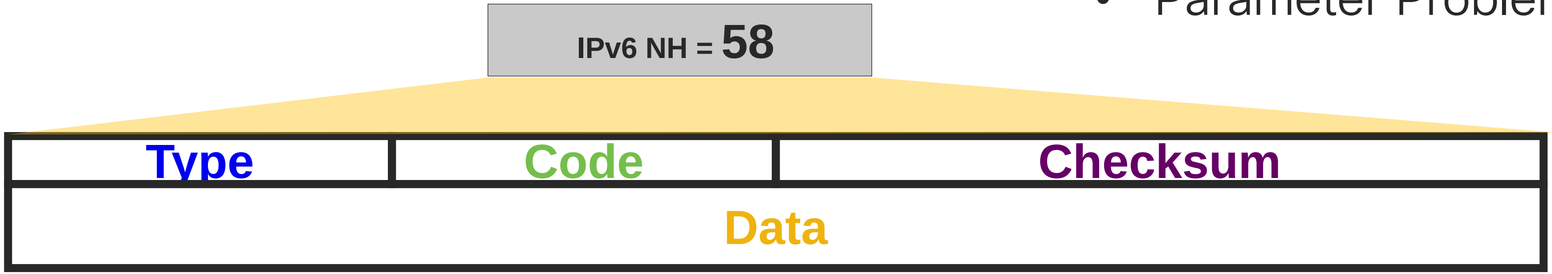

- Type – (0-127) = Error messages, 128-255 Informational messages
- Code – More granularity within the type
- Checksum – Computed over the entire ICMPv6 & pseudo header
- Data – Contents of "offending", filled to 1280B (error) or specific message format (info)

Path MTU Discovery

Next Header	Reserved	Fragment Offset	Res	M
Identification				

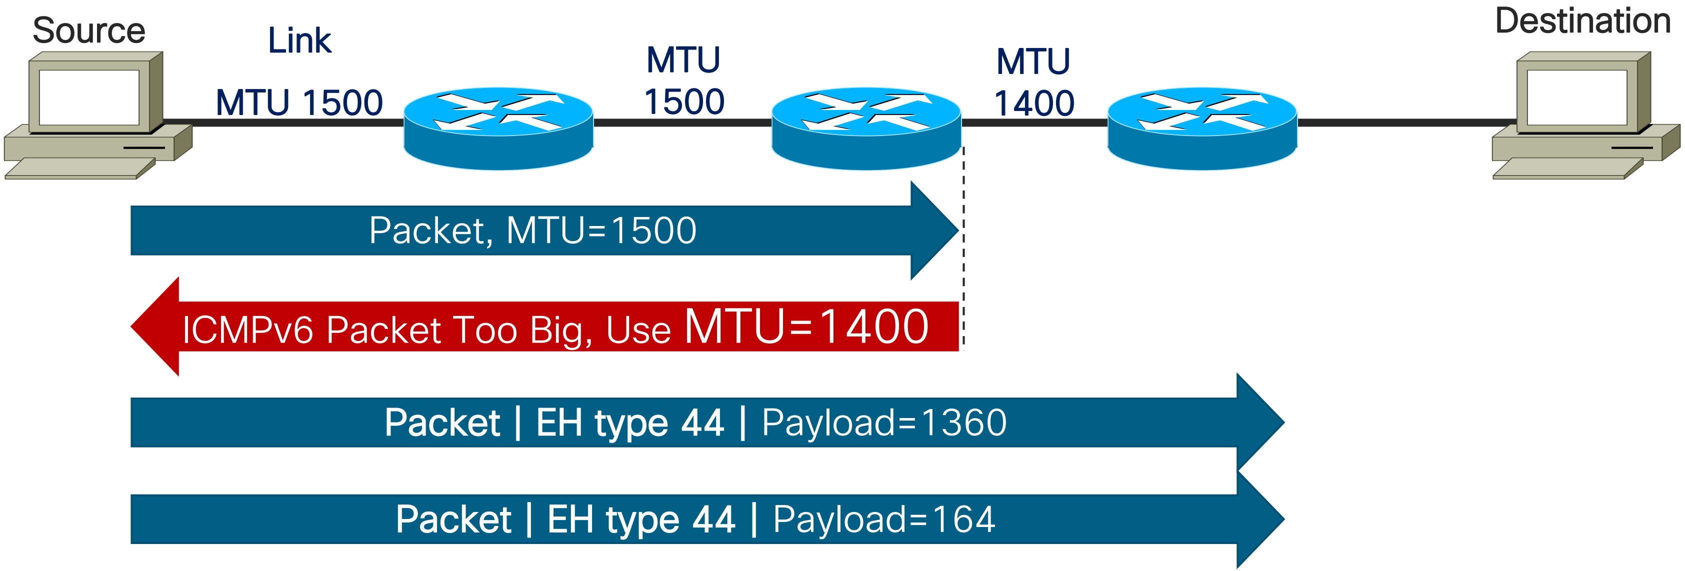

IPv6 Multicast Address (RFC 4291)

- Prefix ff00::/8
- Changes based on flag settings
  - Typically the last 32 bits of host's unicast address

8-bits	4-bits	4-bits	112-bits
1111 1111	0 R P T	Scope	Variable format

Flags	
O	Reserved
R = 0 R = 1	No embedded RP Embedded RP
P = 0 P = 1	Without Prefix Address based on Prefix
T = 0 T = 1	Well Known Address Temporary address

Scope	
1	Node
2	Link
3	Realm
4	Admin
5	Site
8	Organization
e	Global

Unicast:
2001:0db8::426:c001

Multicast:
ff0e::426:c001

Special Use Addresses (RFC 5156)

- Localhost
  - 0:0:0:0:0:0:0:1=> ::1
- Unspecified address
  - 0:0:0:0:0:0:0:0=> 0::0 => :: => ::/128
- Documentation Prefix
  - 2001:0db8::/32
- Discard Prefix
  - 0100::/64
- 6to4 Automatic Tunneling
  - 2002::/16
- Default Route
  - ::/0

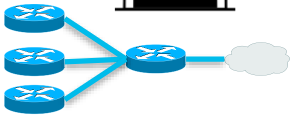

Interface ID’s

cisco Live!

IPv6 Interface id Assignment

Similar to IPv4	New in IPv6
Manually configured 	StateLess Address AutoConfiguration SLAAC EUI64 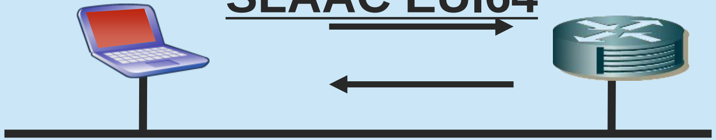
Assigned via DHCPv6 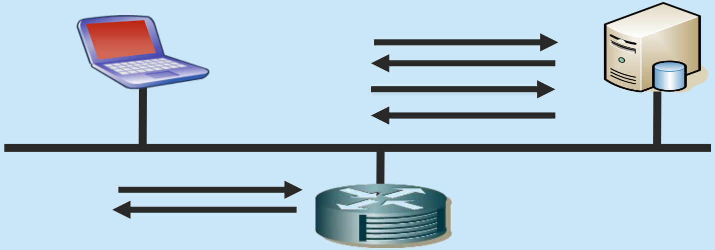	SLAAC Privacy Extensions 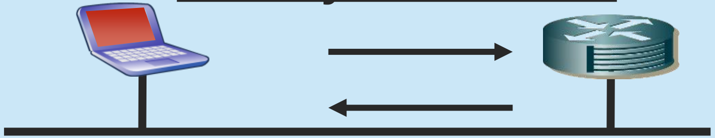

*Secure Neighbor Discovery SeND

Extended Unique Identifier (EUI64)

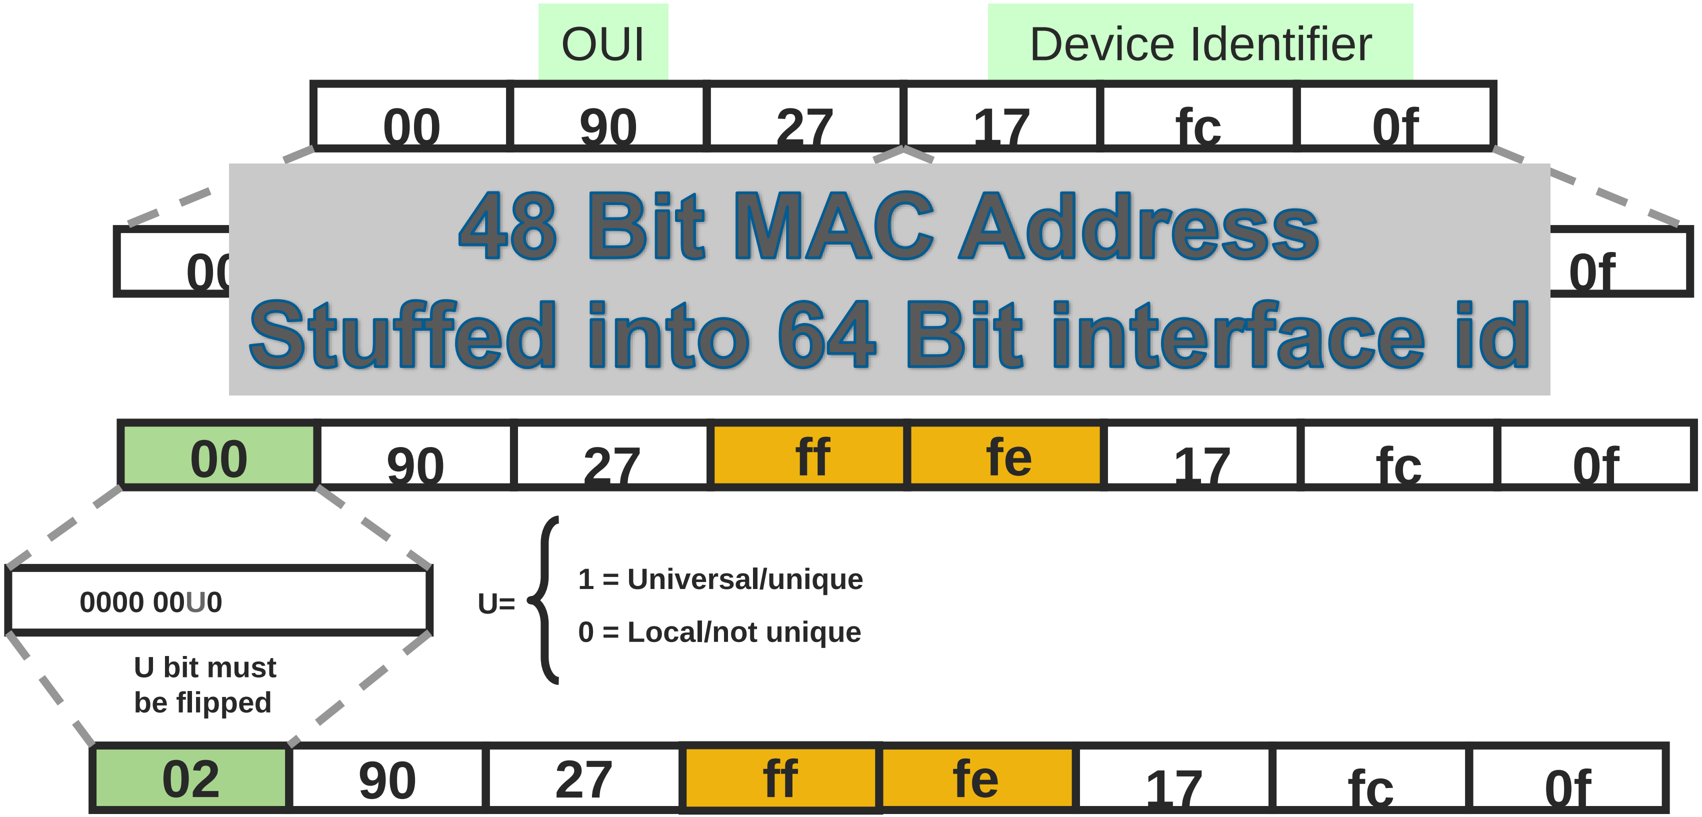

IPv6 Privacy Extensions (RFC 4941)

- Generated on unique 802 using MD5, then stored for next iteration
- **Enabled by default** in Windows, Android, iOS, Mac OS/X, Linux
- **Temporary or Ephemeral addresses** for client application (web browser)

- RFC 7217
- Generate IID's that are Stable/Constant for Each Network Interface
- IID's Change As Hosts Move From One Network to Another

	/32	/48	/64	Randomly generated /64
2001	**db8**	**4646**	**1234**	**ed34:fc01:93d2:630c**

DHCPv6

- Source – fe80::1234, Destination – ff02::1:2
- Client UDP 546, Server UDP 547
- DUID – Different from v4, used to identify clients
- Original Multicast encapsulated in unicast **(relay)**
- *ipv6 dhcp relay destination* 2001:db8::feed:1

SOLICIT (any servers) →
← ADVERTISE (want this address)
REQUEST (I want that address) →
← REPLY (It's yours)

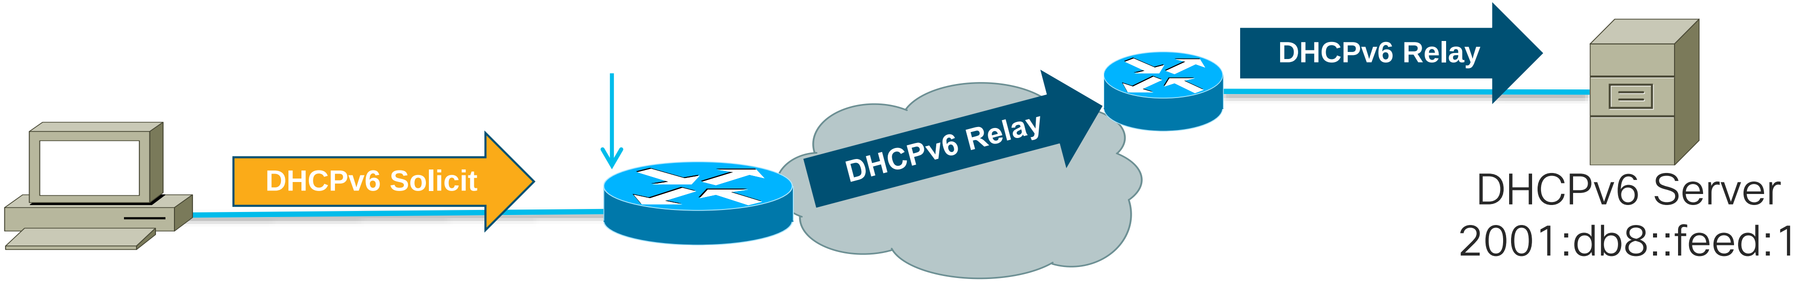

Link Operations

Neighbor Discovery Protocol – (NDP)

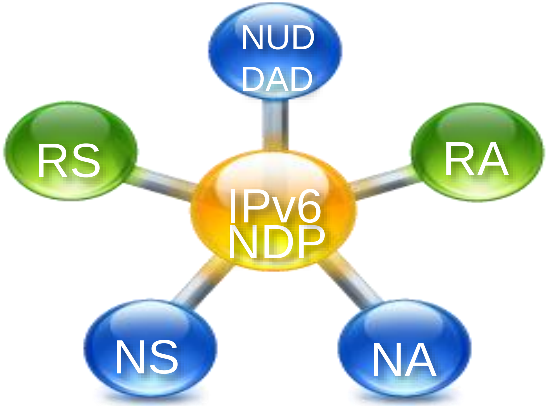

- Should use Link Local (fe80::/64) as its source
- Hop Limit must be set to 255
  - Generalized TTL Security Mechanism

- 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)

IPv4	IPv6
ARP Request	Neighbor Solicitation
ARP Reply	Neighbor Advertisement
Broadcast	Solicited Node Multicast

Solicited Node Multicast

- Required & special form of multicast used for neighbor resolution
- Every Unicast address must
  - Create corresponding solicited node multicast (ff02::1:ff00:0/104)
- Any Layer 3 multicast must
  - Map to corresponding Layer 2 multicast (33-33-xx-xx-xx-xx)

IPv6 Source	fe80::04cb:57ff:fe**3c:deca**
IPv6 Destination	ff02::1:ff**3c:deca**
Ethernet Destination	33-33-FF-**3C-DE-CA**
Ethernet Source	02-CB-57-**3C-DE-CA**

Solicited Node Multicast Example

```
R1#sh ipv6 int e0
Ethernet0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::200:CFF:FE3A:8B18
  Global unicast address(es):
    2001:DB8:46:1234::1 subnet is 2001:DB8:46:1234::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:1
    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 router advertisements are sent every 200 seconds
```

Solicited-Node Multicast Address*

*If EUI format is used then the 1rst solicited node mcast addr is used for both the LL & GU

Neighbor Solicitation & Advertisement

- Node A needs to resolve node B's link address, Map's L3 to L2
- Multicast for resolution (new), Unicast for reachability (cache)
- Node B will add node A to its neighbor cache during this process w/o sending NS

ICMPv6 Type	135 NS
IPv6 Source	fe80::a
IPv6 Destination	**ff02::1:ff**00:b
Hop Limit	255
Target Address	2001:db8:46:46::b
Query	**What is B link layer address?**
Opt. 1 SLLA	A's Link Layer Address

ICMPv6 Type	136 NA
IPv6 Source	fe80::b
IPv6 Destination	fe80::a
Target Address	2001:db8:46:46::b
Option 2 TLLA	**B's Link Layer Address**
***Flags**	R = Router S = Response to Solicitation O = Override cache information

Duplicate Address Detection (DAD)

- Unspecified Source (::), No Option 1 SLLA
- Probing the Local Link to Verify Address Uniqueness
- Microsoft uses a variant known as Optimistic DAD

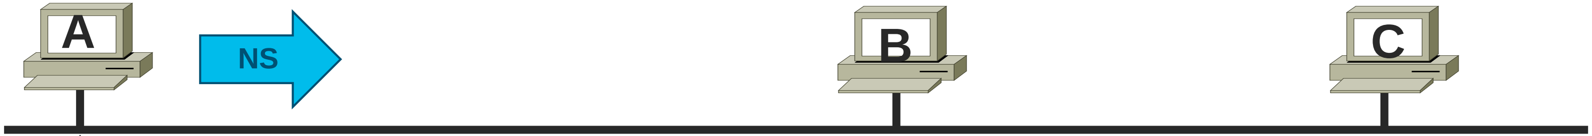

ICMPv6 Type	135 NS
IPv6 Source	UNSPEC = ::
IPv6 Dest.	A Solicited Node Multicast ff02::1:ff00:a
Query	Anyone Using “a”

Router Solicitation and Advertisement

- Router solicitations (RS) are sent by nodes at boot up
- Routers forward packets as well as provide provisioning services

RS	
ICMP Type	133
IPv6 Source	fe80::a
IPv6 Destination	ff02::2
Opt. 1 SLLA	SRC Link Layer Address

RA	
ICMP Type	134
IPv6 Source	fe80::1
IPv6 Destination	fe80::a
Data	Options, subnet prefix, lifetime, autoconfig flag

Router Advertisement

- M-Flag – Stateful DHCPv6 to acquire IPv6 address
- O-Flag – Stateless DHCPv6 in addition to SLAAC
- Preference Bits – Low, Med, High
- Router Lifetime – Must be >0 for Default
- Options - Prefix Information, Length, Flags
- L bit – Only way a host get a On Link Prefix
- A bit – Set to 0 for DHCP to work properly

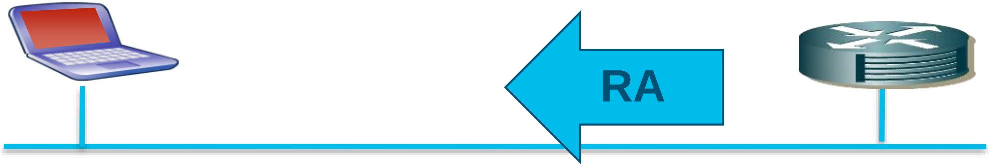

Type: 134 (RA)
Code: 0
Checksum: 0xff78 [correct]
Cur hop limit: 64
∞ Flags: 0x84
    1… …. = Managed **(M flag)**
    .0.. …. = Not other **(O flag)**
    ..0. …. = Not Home (H flag)
    …0 1… = Router pref: High
Router lifetime: (s)**1800**
Reachable time: (ms) 3600000
Retrans timer: (ms) 1000
ICMPv6 Option 3 (Prefix Info)
Prefix length: 64
∞ Flags: 0x84
    1… …. = On link (L Bit)
    …. 1… = No Auto (A Bit)
Prefix: 2001:0db8:4646:1234::/64

Host OS's

Router Advertisement Sent

```
▽ Internet Control Message Protocol v6
    Type: Router Advertisement (134)
    Code: 0
    Checksum: 0x1a4d [correct]
    Cur hop limit: 64
  ▽ Flags: 0x00
      0... .... = Managed address configuration: Not set
      .0.. .... = Other configuration: Not set
      ..0. .... = Home Agent: Not set
      ...0 0... = Prf (Default Router Preference): Medium (0)
      .... .0.. = Proxy: Not set
      .... ..0. = Reserved: 0
    Router lifetime (s): 180
    Reachable time (ms): 0
    Retrans timer (ms): 0
  ▷ ICMPv6 Option (Recursive DNS Server 2001:558:feed::1)
  ▽ ICMPv6 Option (Prefix information : 2001:db8:46:1::/64
      Type: Prefix information (3)
      Length: 4 (32 bytes)
      Prefix Length: 64
    ▷ Flag: 0xc0
      Valid Lifetime: 300
      Preferred Lifetime: 300
      Reserved                                       (2001:db8:46:1::)
      Prefix: 2001:db8:46:1::
  ▷ ICMPv6 Option (Source link-layer address : 00:50:f1:00:00:00)
```

OSX Host Address Acquisition

- Effect of the Router Advertisement from previous slide
  - Preferred & valid lifetimes
  - DNS server information

```
tmartin# ifconfig -L
en0: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
      ether b8:e8:56:19:f3:8a
      inet6 fe80::bae8:5642:ce19:f38a%en0 prefixlen 64 scopeid 0x4
      inet6 2001:db8:46:1:1809:5618:fa19:f38a prefixlen 64 autoconf pltime 267 vltime 267
      inet6 2001:db8:46:1:883e:b6a2:863:e31b prefixlen 64 autoconf temp pltime 267 vltime 267
      nd6 options=201<PERFORMNUD ,DAD>
```

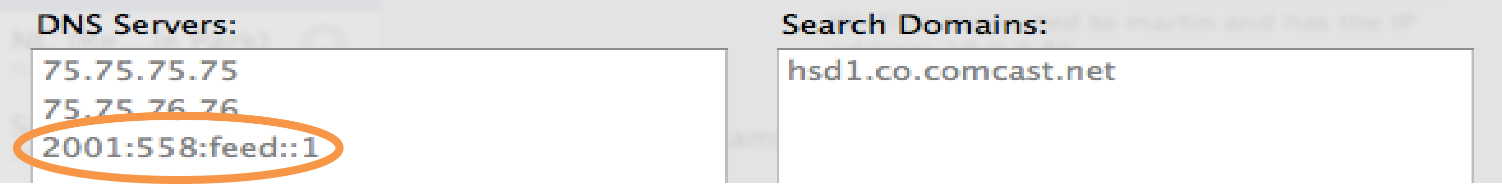

Linux, Ubuntu IPv6 Basics

- Check an IPv6 address
  - ip -6 addr show dev eth0
- Check an IPv6 address
  - ifconfig eth0 | grep “inet6 addr:”
- Check for IPv6 neighbors
  - ip -6 neigh show
- Ping6 2001:db8:46:1::c001:d00d
  - Windows:> ping fe80::250:e1ff:fec3:1**%17**
  - Unix$ ping6 fe80::250:e1ff:fec3:1**%en0**

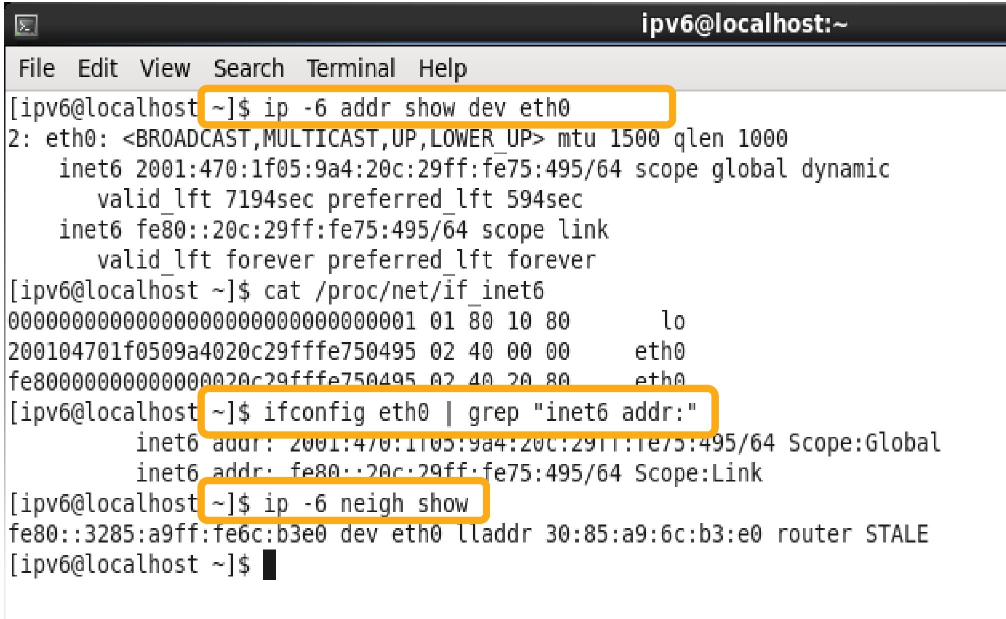

Windows Host Address Acquisition

```
C:\Documents and Settings\>netsh
netsh>interface ipv6
netsh interface ipv6>show address
Querying active state...
Interface 5: Local Area Connection
Addr Type  DAD State   Valid Life    Pref. Life    Address
---------  ----------  ------------  ------------  -----------------------------
Public     Preferred   29d23h58m25s   6d23h58m25s  2001:db8:4646:1:4f02:8a49:41ad:a136
Temporary  Preferred    6d21h48m47s        21h46m  2001:db8:4646:1:bd86:eac2:f5f1:39c1
Link       Preferred       infinite      infinite  fe80::4f02:8a49:41ad:a136

netsh interface ipv6>show route
Querying active state...
Publish  Type      Met  Prefix                    Idx  Gateway/Interface Name
-------  --------  ---- ------------------------  ---  ---------------------
no       Autoconf    8  2001:db8:4646:1::/64        5  Local Area Connection
no       Autoconf  256  ::/0                        5  fe80::20d:bdff:fe87:f6f9
```

Timer’s & Tables

Address Timers & State

- Tentative – Address in verification process (DAD)
- Preferred – Address can be used for communication
- Valid – Address can be used, may be Preferred or Deprecated
- Deprecated – Address can be used on existing connections
- Invalid – Address is not available for use

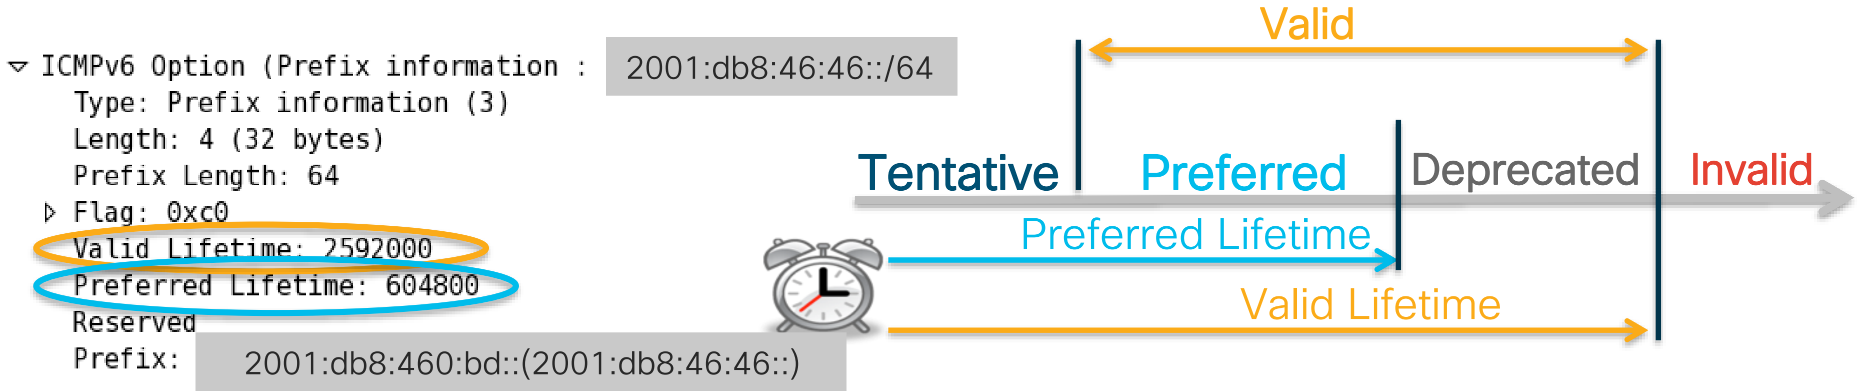

Host Cache Conceptual Data Structures

- Maintained for each interface connected on a host
- Host uses the PL & DRL to work out the destination for outbound packet
- Then it saves the result in the DC, layer 3 mapping to next hop
- Hosts use neighbor discovery to get the link address and update the NC

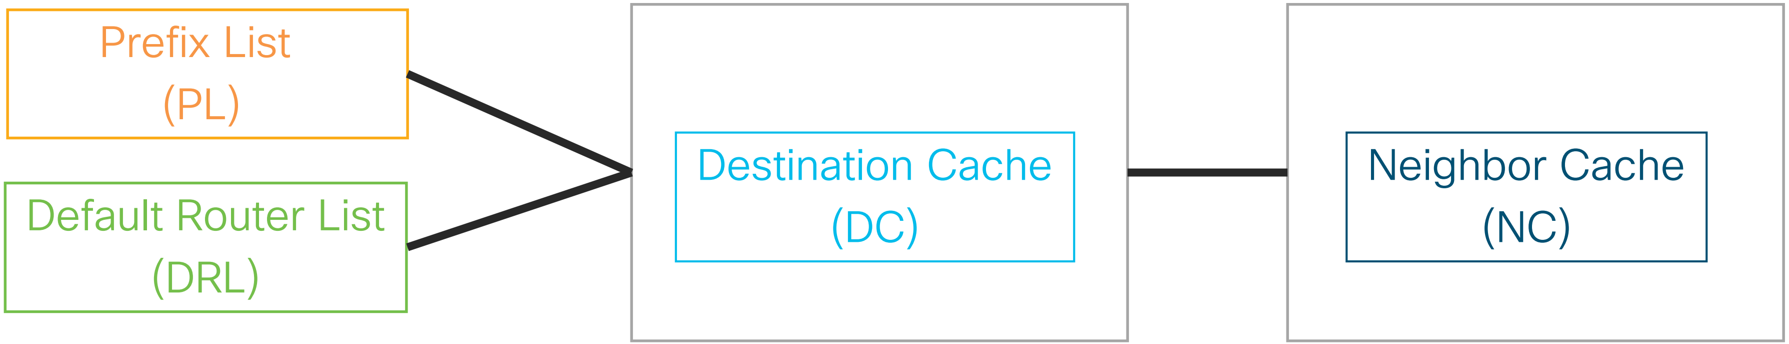

Host Cache Conceptual Model

- Prefix List – contains on link prefixes (L bit) and validation timers
- Default Router List – must be a neighbor usable to the host (Pref bits)
- Destination Cache – resolves next hop IPv6 address and MTU

Prefix List (PL)	Valid Timer
fe80::/10	∞
2001:db8:4646:34::/64	322486

Default Router List (DRL)	Preference
fe80::34:1	H
fe80::34:11	M

Destination Cache (DC)	Neighbor	PMTU
fe80::34:1	fe80::34:1	1500
2001:db8:4646:34::1	2001:db8:4646:34::1	9000
2001:db8:4646:555::22	fe80::34:1	1500
2001:db8:4646:717::98	fe80::34:1	1500
2001:db8:4646:34::3	2001:db8:4646:34::3	9000

Host Neighbor Cache

- Mapping of the neighbors IPv6 address to its link layer address
- Includes the status of the “R” flag in the returned NAs
- Must not create a new entry for “gratuitous” NA
    - Though such an NA can update an existing entry

Neighbor	Link Layer	Is Router	State
fe80::34:1	00-00-0C-83-5C-3E	1	Reachable
2001:db8:4646:34::1	00-00-0C-83-5C-3E	0	Stale
2001:db8:4646:34::3	04-48-9A-16-37-FB	0	Stale
ff02::1	33-33-00-00-00-01	0	~

Neighbor Cache State Machine

- Incomplete – Pending address resolution, NS message outstanding
- Reachable – Recently used mapping, Can be refreshed by ULP
- Stale – Not currently communicating, waiting for next queued packet
- Delay –Using stale binding, awaiting (ULP) return traffic
- Probe – Sending Unicast NS to node (after Delay timer, 3x1 sec)

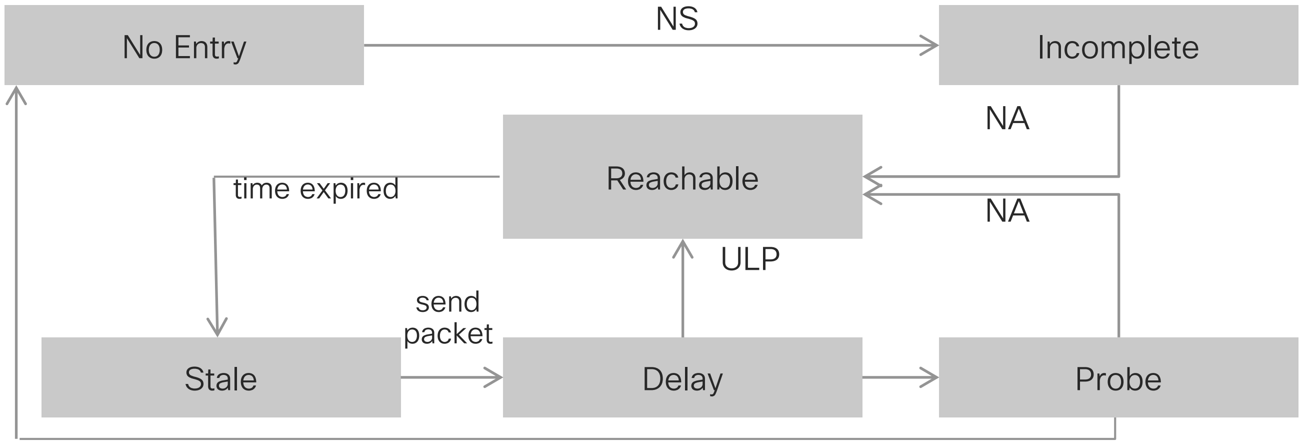

RFC 6724 – Source Address Selection

- IPv6 over IPv4, but which IPv6 first..
- Scope = Local, ULA, Global
- State = Preferred over Deprecated
- Interface = Assigned vs. another
- Type = Temporary over Public

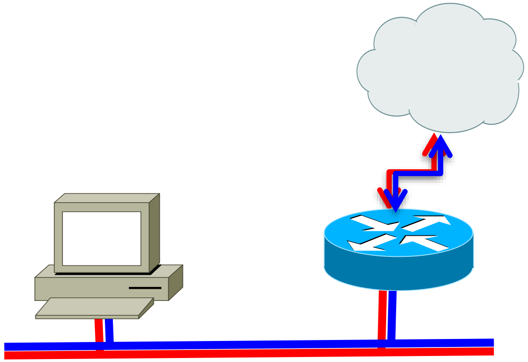

↓	↓	↓	↓
Public	Preferred	2001:db8:4646:1:4f02:8a34:bead:a136	Idx
Temporary	Preferred	2001:db8:4646:1:bd86:ea49:41f1:39c1	Idx
Link	Preferred	fe80::4f02:8a34:bead:a136	Idx

Dual Stack OS Considerations

- In a dual stack case, an application can:
  - Query DNS for IPv4 and/or IPv6 records
  - Parallel or serial connection request
  - Set resolution delay to 50ms

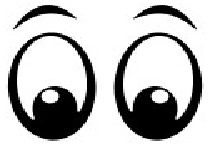

RFC 8305
Happy Eyeballs

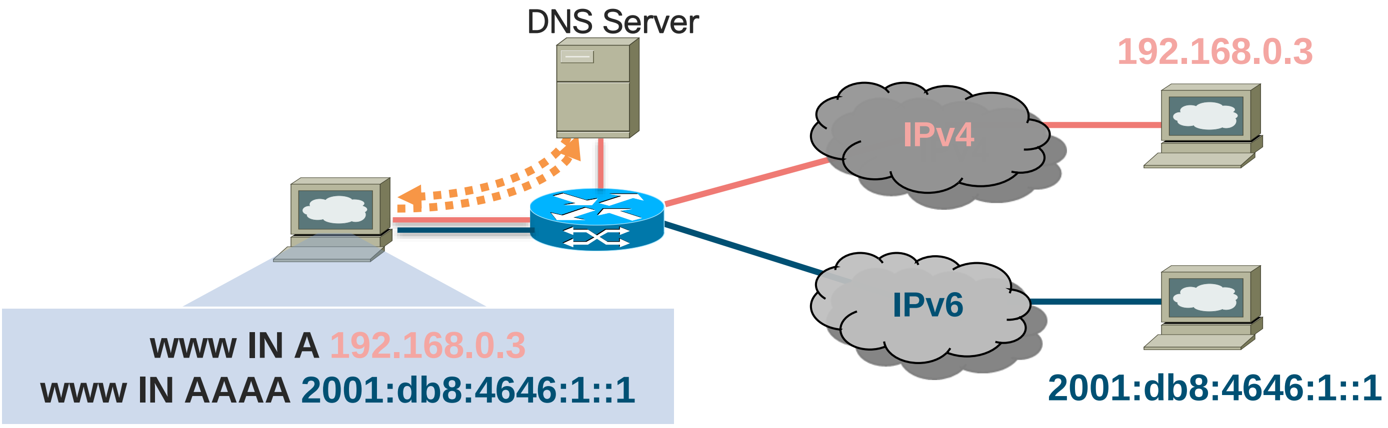

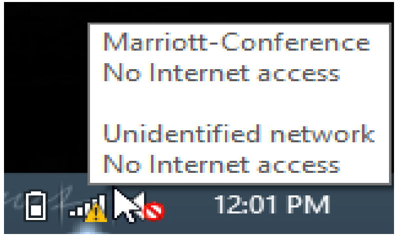

cisco

IPv6

Address planning & some thoughts

Nicole Wajer
Chief Stroopwafel Officer

TECRST-1991

Follow
@vlinder_nl

Cisco Live DISTINGUISHED Speaker

cisco Live!
Barcelona | January 27-31, 2020

Global Address Assignment

- Provider Allocated (PA)
  - From your ISP, single homed
  - /48 - /60
- Provider Independent (PI)
  - Multi home, Multi provider
  - /32 - /48
- Local Internet Registry (LIR)
  - Regional registry member
  - Acquire & manage space
  - /29 - /48

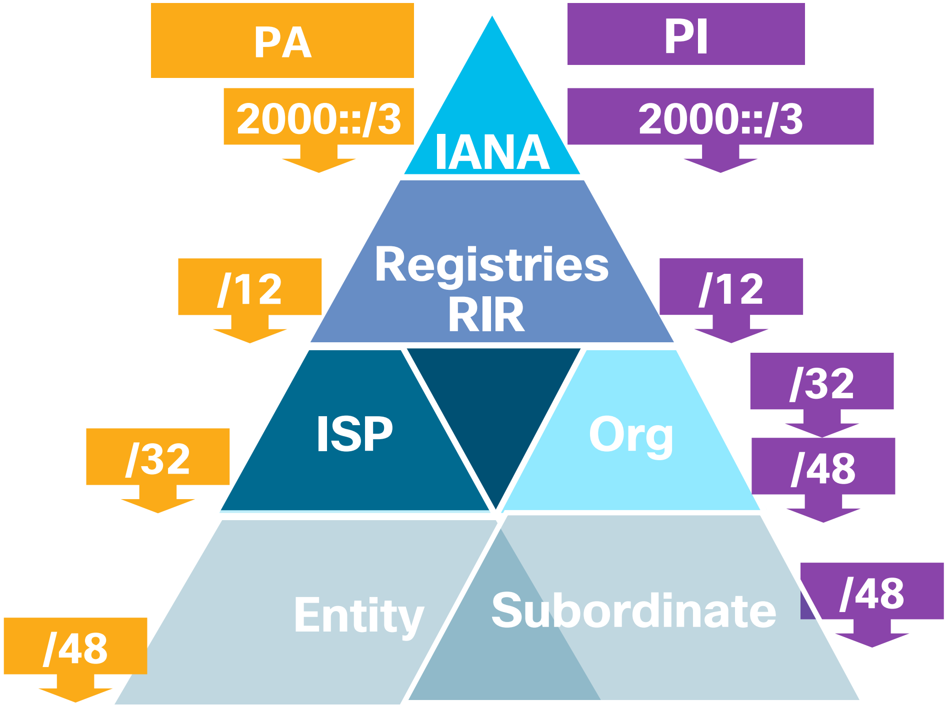

Multi-national Model

- PA or PI from each region you operate in
- Coordination of advertised space within each RIR, policy will vary
- Most run PI from primary region

Unique Local Address (ULA)

- Prefix Generation (RFC 4193) non sequential /48, M&A challenges
- To be avoided in most cases, v6ops-ula-usage-recommendations
- Caution with older OS's (RFC 3484) using ULA & IPv4
- Multiple policies to maintain (ACL, QoS, Routing, etc..)

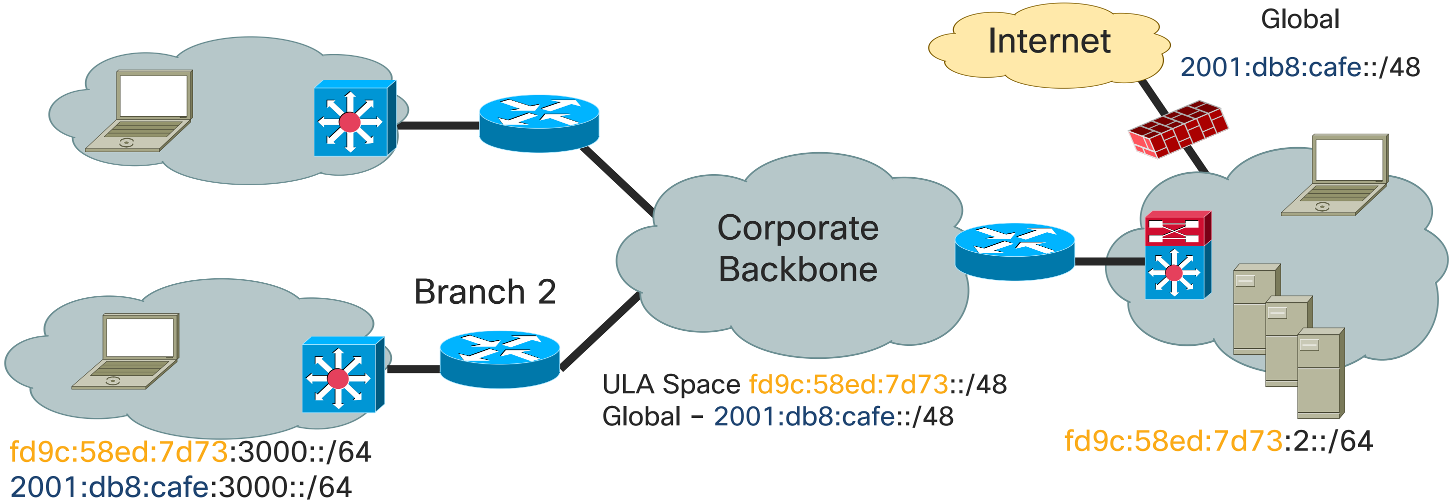

Infrastructure using Link Local Addressing

- Topology hiding, Interfaces cannot be seen by off link devices
- Reduces routing table prefix count, less configuration
- Need to use ULA or GUA for generating ICMPv6 messages
- What about DNS?, Traceroute, WAN Connections, etc..
- RFC7404 – Details pros and cons

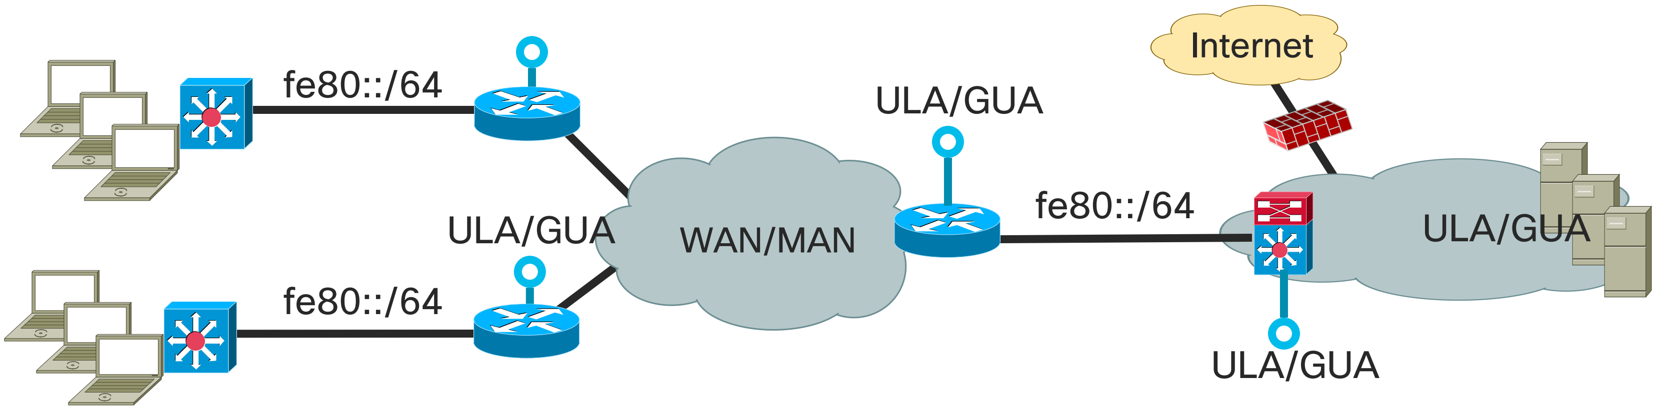

Prefix Length Considerations

- Anywhere a host exists /64
  - RFC 7421, rational for /64
- Point to Point /127
  - Should not use all 0's or 1's interface
  - Nodes 1 & 2, not in the same subnet
  - Reserve the /64,
    configure /127 with :a & :b
  - RFC 6164: /127 cache exhaust
  - Reserve a /64, configure a /127
- Loopback or Anycast /128

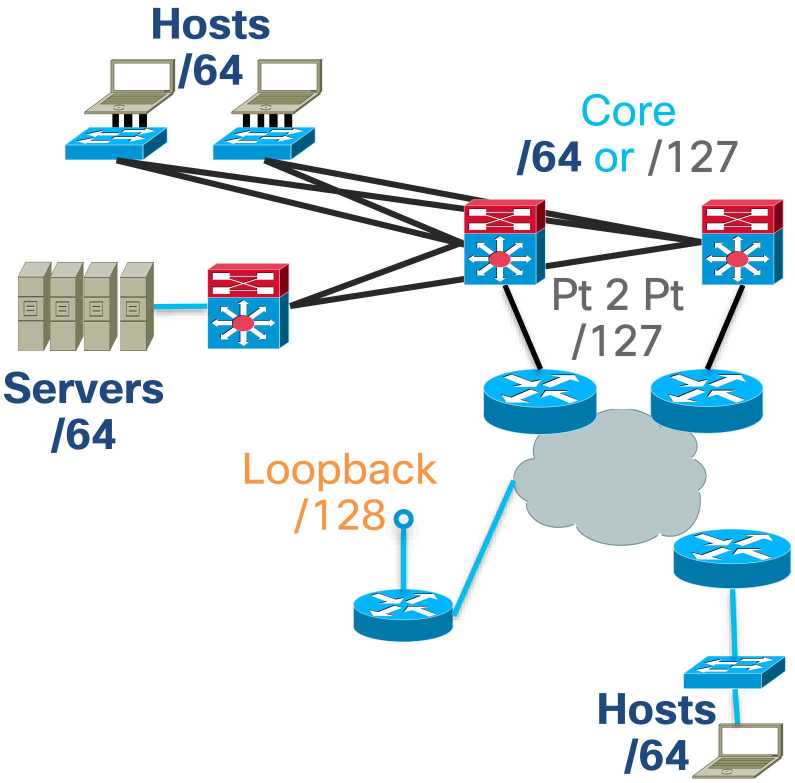

Why is IPv6 Addressing Plan Important?

- Helps you to get your head around IPv6 deployment
  - Structure - Services
  - Hierarchy - Security
- Puts structure in place
  - Like belts and braces
- First real "toe in the cold water"
  - In terms of deployment
  - Makes the "jump" easier
- Well thought-through addressing plan
  - Think network operations and troubleshooting

cisco Live!

Let it Go!

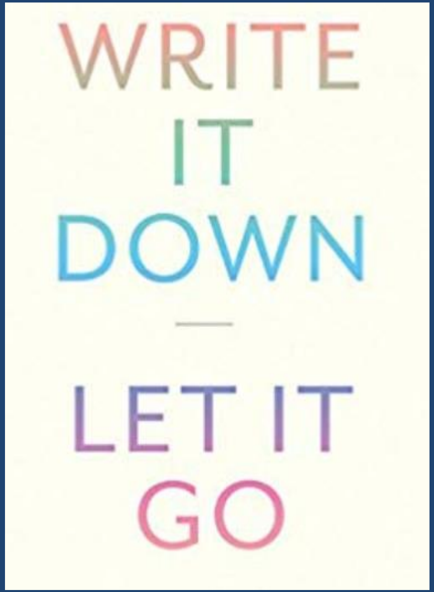

IPv6 provides more options than IPv4

- Links can have multiple prefixes:
  - Every link (point-to-point, VLAN) has link-local addresses
  - Links can have global addresses
  - Links can have ULA addresses
- Multiple global and ULA prefixes per link is possible
  - Can be useful for renumbering to have multiple prefixes active at the same time
  - But do you really want that during normal operation???
- Keep it simple!

IPv6 provides more options than IPv4

- Links can have multiple prefixes:
  - Every link (point-to-point, VLAN) has link-local addresses
  - Links can have global addresses
  - Links can have ULA addresses
- Multiple global and ULA prefixes per link is possible
  - Can be useful for renumbering to have multiple prefixes active at the same time
  - But do you really want that during normal operation???
- Keep it simple!

```
en0: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
	ether 00:24:36:ef:1d:89
	inet6 fe80::224:36ff:feef:1d89%en0 prefixlen 64 scopeid 0x4
	inet6 2a02:a213:a300:9300:224:36ff:feef:1d89 prefixlen 64 autoconf
	inet6 2a02:a213:a300:9300:1015:5159:6b61:8f5f prefixlen 64 autoconf
	inet6 2001:9e0:803:aa2:224:36ff:feef:1d89 prefixlen 64 autoconf
	inet6 2001:9e0:803:aa2::80 prefixlen 64 autoconf
	inet 192.168.178.10 netmask 0xffffff00 broadcast 192.168.178.255
```


IPv4 Address Assessment

- Assess how the existing IPv4 address space is used
- Useful information for
  - IPv6 integration
  - IPv4 address consolidation
  - Reclaiming unused address space
- Use existing tools
  - IPAM
  - ARP tables
  - DHCP logs

→ Better visibility into how the existing

Address space is used

→ Can better answer when IPv6 is critical

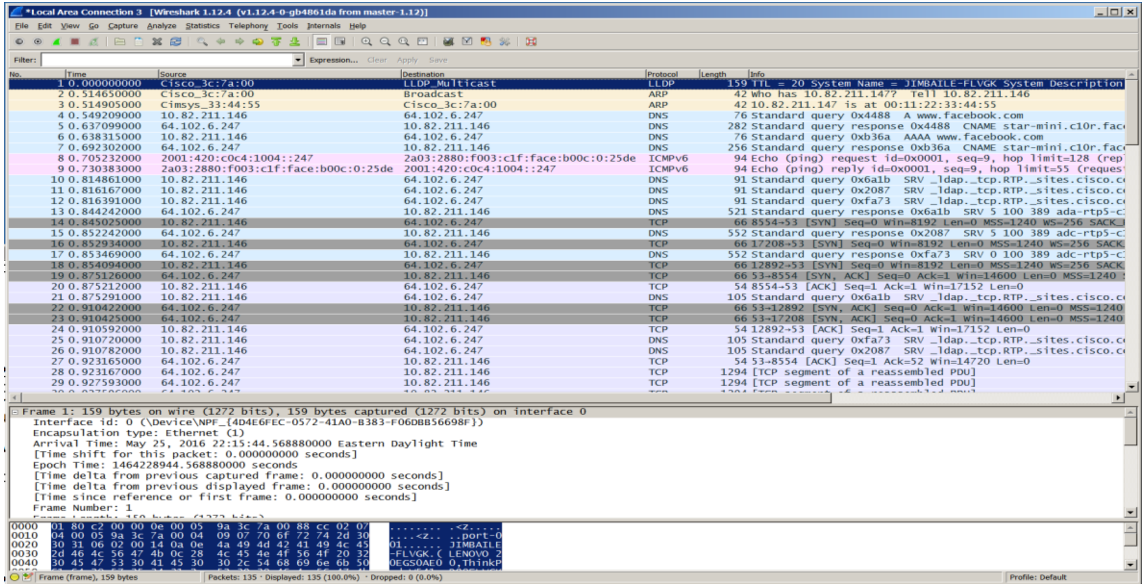

Addressing Plan Requirements and Considerations

- Requirements
  - Length of prefix and bits to work with
    - Enterprises usually multiple /48
    - Highly dependent on RIR policy
    - SPs should get /29 (≥ 35 bits)
  - Avoid breaking the nibble boundary
  - Think of # of prefixes at each level
  - Templates will be your friends
  - IP Address Management Tools

- Considerations
  - Clear addressing for different parts of the network
    - WAN/Core, Campus, branch, DC, Internet Edge etc.
  - Different Locations/Services
  - Encoding of information
  - Ease of aggregation
  - Leaving space for growth

Numbering your infrastructure

- Use only Link-Local on links between routers?
  - Easy to configure, smaller attack surface, harder to debug
- Use global addresses for systems and ULA for infrastructure?
  - Smaller attack surface, complex addressing, confusing to ops, ICMP problems
- Use both global addresses and ULA side-by-side for everything?
  - Complex to configure and maintain, unpredictable behaviour, (almost) no benefits
- Use global addresses everywhere?
  - Use IPv6 as designed: This is the best choice in most cases!

Host Address Assignment

- Depends on the existing processes and the adaptability of that process
- Methods are not mutually exclusive - all three can be used
- Must control the stateless address assignment of addresses

	Manual	Stateless	Stateful DHCPv6
Pros	Address is stable Controlled assignment Well understood process	Scales well Time to deploy Widely implemented	Well understood process Controlled assignment Time to deploy
Cons	Does not scale Time to deploy	No control on assignment process Not well understood Lack of management Unpredictable (privacy extensions)	Implementation in OS Must design for HA

Methodology – Avoid Zero Compression

- Operations may not understand what you do
- Avoid the zero compression "gotcha"
  - `2001:db8::/32`
  - `2001:db8:0000::/48`
  - `2001:db8:0000::/52`
  - `2001:db8::`
- Same goes for host based subnets
  - `2001:db8:1020:0000::/64` could look like `2001:db8:1020::`

Example - How Many Subnets in a Location?

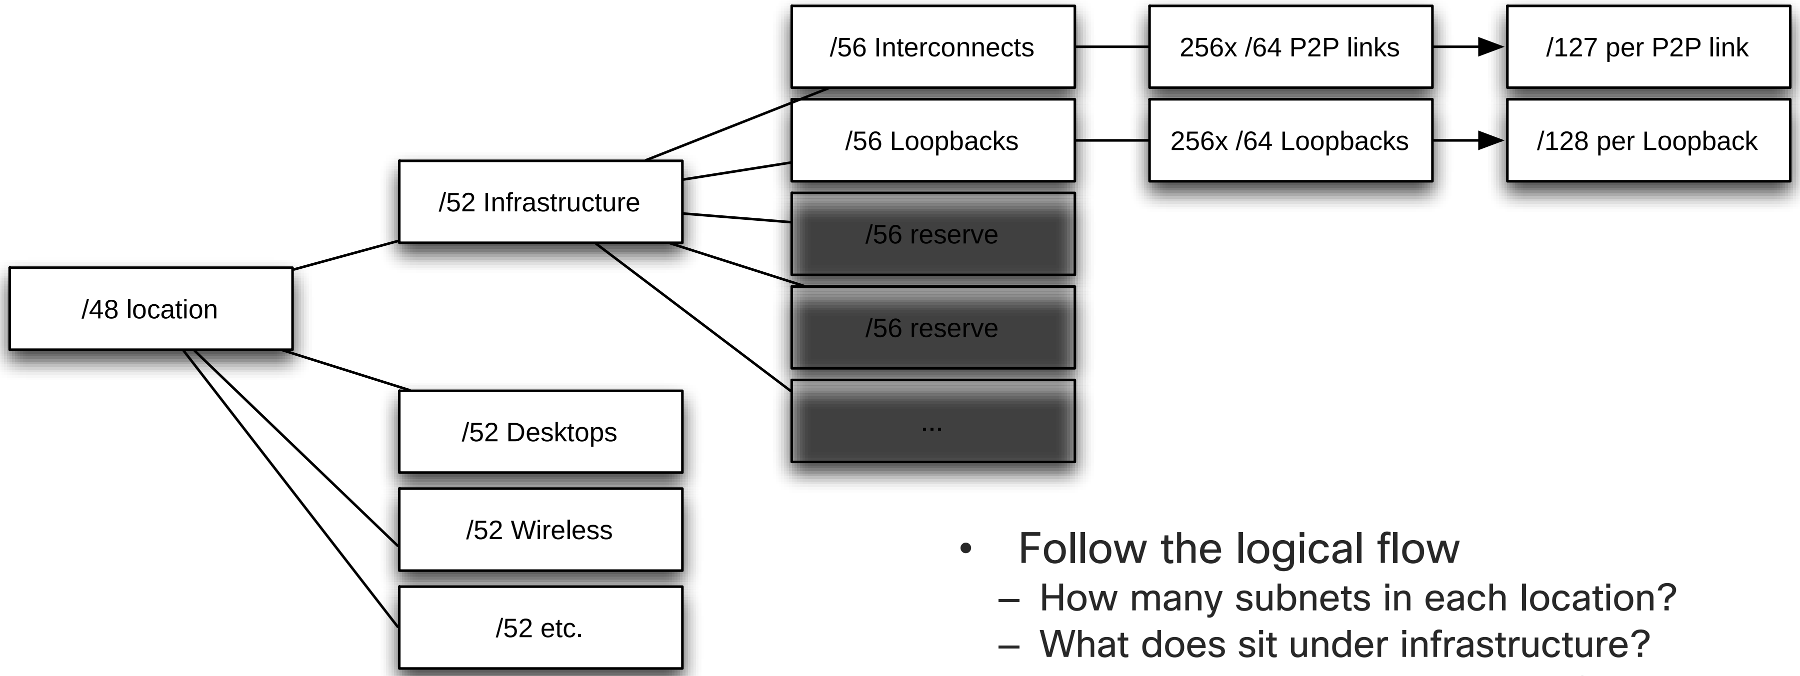

- Follow the logical flow
  - How many subnets in each location?
  - What does sit under infrastructure?
  - How many point-to-point links?
  - Where is the reserve?

Ideas for grouping prefixes

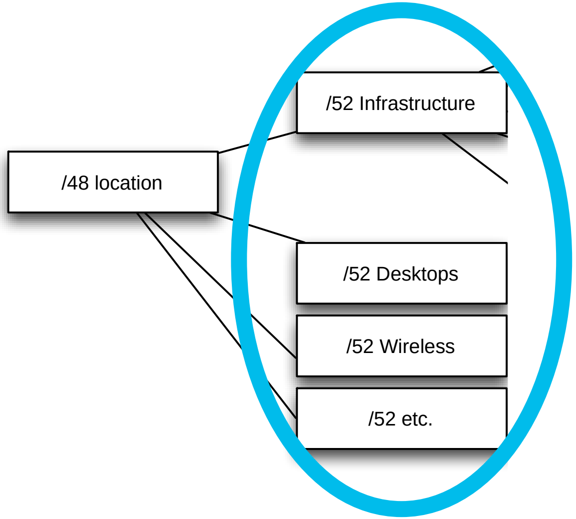

- Always make one group for infrastructure
  - Easier to protect and monitor
  - Using group number 0 gives shorter addresses
- Other grouping ideas:
  - Organizational division (group by responsibility)
  - Type of use (group by expected traffic)
  - Security boundary (group by firewall policy)
  - Whatever makes sense in your network!

- Further reading
  - https://www.ripe.net/support/training/material/IPv6-for-LIRs-Training-Course/Preparing-an-IPv6-Addressing-Plan.pdf

4 Rules

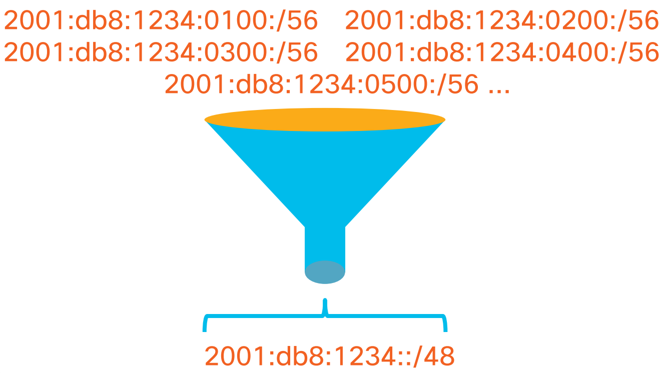

1. Keep it SIMPLE

   You don't want to spend weeks explaining it!

2. Embed information to help operations

   To help troubleshooting and operation of the network

   Examples: location, country, firewall zone, VLAN, IPv4 addresses (be careful)

3. Plan for expansion (build in reserve)

   Cater for future growth, mergers & acquisitions, new locations

   Reserved vs. assigned

4. Exploit hierarchy / aggregation

   Good aggregation is essential, just one address block (per location)

   Ensures scalability and stability

Address Planning Exercise

Idea Enterprise

- Furniture Company
- Has presence throughout Europe
  - Currently set up in w/ 19 regional offices
  - Plans are to expand to 37 regional offices
- They also have a sister company (Idea ISP) which is providing telecommunications services throughout Europe

Idea Enterprise Current State of the Network

- Idea has grown organically through a policy of acquisitions and mergers over the past few years.
- Use of private (RFC 1918) and/or illegal IPv4 address blocks, NAT is widely used. This is negatively impacting the behaviour of some enterprise applications.
- Idea has decided to strategically deploy IPv6 within the IdeA enterprise network. This will enable applications and services to be moved from IPv4 to IPv6 on a case-by-case basis
- For its WAN connectivity, Idea enterprise uses the MPLS VPN service offered by Idea ISP.

Idea Enterprise IPv6 High Level Requirements

- Idea ISP is a RIPE member and have been allocated a /19 IPv6 address block. Idea Enterprise has been provided 2001:db8::/32 from its ISP. Idea ISP will be interconnecting all the IPv6 locations of the Idea enterprise network.
- The most important requirements for the IPv6 addressing design are for it to be **highly hierarchical, uniform and scalable**. This will greatly simplify the design, operation and troubleshooting of the network.

Idea Enterprise Detail Requirements

- As a general rule, Idea would like to use byte (8-bit)-boundaries between the different hierarchies of the IPv6 addressing. HINT!!!
- At the first level, the addressing scheme needs to support at least 37 regional offices (HINT!!!). Also some address blocks should be reserved for future growth in the larger countries.
- At the second level (within each region), there are a number of campus locations. It is at this level that connectivity into the Idea ISP network is provided. The largest region has about 40 campus locations (HINT!!!).

Idea Enterprise Detail Requirements (Cont.)

- At the third level (within each campus location), the number of buildings within each campus (4-6 maximum). Therefore, allocating these blocks on a byte boundary is deemed as overkill. A nibble (4-bit) boundary will suffice here. HINT!!!
- A separate "virtual building" address block needs to be set aside for network infrastructure addressing within that campus location
- At the forth level (within each building), individual IPv6 subnets need to be assigned to individual VLANs
- An additional requirement is to divide up the network infrastructure block in ranges for loopback, link and network services addressing

Building An Address Plan For Idea Enterprise (Tasks)

- Design an IPv6 address plan for Idea enterprise applying with what you have learned in this session and the mentioned HINTS
- Work top-down through the address plan
- Focus first on the end-system addressing
- Think about the network infrastructure addressing
- There are multiple acceptable solutions, it's more important to think about the problem and apply the methodology

Idea Enterprise IPv6 Addressing Exercise

Global Prefix: 2001:db8::/32

Adress Scope	Prefix Length	# of address block allocation w/in higher level	First address block allocation w/in higher level	Last address block w/in higher level
Region				
Campus Location				
Building				
End-System				
Network Infrastructure				
Loopbacks				
Links				
Network Services				

Solution

Idea Enterprise IPv6 Addressing Exercise

For Your Reference

Global Prefix: 2001:db8::/32

Adress Scope	Prefix Length	# of address block allocation w/in higher level	First address block allocation w/in higher level	Last address block w/in higher level
Region	2001:db8::/40	256	2001:db8::/40	2001:db8:ff00::/40
Campus Location				
Building				
End-System				
Network Infrastructure				
Loopbacks				
Links				
Network Services				

Region 1 2001:db8:1000::/40 -> Reserved 2001:db8:1100::/40 - 2001:db8:1f00::/40

Region 2 2001:db8:2000::/40 -> Reserved 2001:db8:2100::/40 - 2001:db8:2f00::/40

Idea Enterprise IPv6 Addressing Exercise

For Your Reference

Adress Scope	Prefix Length	# of address block allocation w/in higher level	First address block allocation w/in higher level	Last address block w/in higher level
Region	2001:db8::/40	256	2001:db8::/40	2001:db8:ff00::/40
Campus Location	2001:db8::/48	256	2001:db8::/48	2001:db8:ffff::/48
Building				
End-System				
Network Infrastructure				
Loopbacks				
Links				
Network Services				

Region 1 Campus 1

2001:db8:1010::/48 -> Reserved 2001:db8:1011::/48 thru 2001:db8:101f::/48

Idea Enterprise IPv6 Addressing Exercise

For Your Reference

Adress Scope	Prefix Length	# of address block allocation w/in higher level	First address block allocation w/in higher level	Last address block w/in higher level
Region	2001:db8::/40	256	2001:db8::/40	2001:db8:ff00::/40
Campus Location	2001:db8::/48	256	2001:db8::/48	2001:db8:ffff::/48
Building	2001:db8::/52	16	2001:db8::/52	2001:db8:ffff:f000::/48
End-System				
Network Infrastructure				
Loopbacks				
Links				
Network Services				

Region 1 Campus 1 Building 1

2001:db8:1010:1000::/52

Region 1 Campus 1 Building 2

2001:db8:1010:2000::/52

Reserved

2001:db8:1010:7000::/52 - 2001:db8:1010:f000::/52

Idea Enterprise IPv6 Addressing Exercise

For Your Reference

Adress Scope	Prefix Length	# of address block allocation w/in higher level	First address block allocation w/in higher level	Last address block w/in higher level
Region	2001:db8::/40	256	2001:db8::/40	2001:db8:ff00::/40
Campus Location	2001:db8::/48	256	2001:db8::/48	2001:db8:ffff::/48
Building	2001:db8::/52	16	2001:db8::/52	2001:db8:ffff:f000::/48
End-System	2001:db8::/64	4096	2001:db8::/64	2001:db8:ffff:ffff::/64
Network Infrastructure				
Loopbacks				
Links				
Network Services				

Region 1 Campus 1 Building 1 VLAN 100 2001:db8:1010:1100::/64

Region 1 Campus 1 Building 1 VLAN 200 2001:db8:1010:1200::/64

Idea Enterprise IPv6 Addressing Exercise

For Your Reference

Adress Scope	Prefix Length	# of address block allocation w/in higher level	First address block allocation w/in higher level	Last address block w/in higher level
Region	2001:db8::/40	256	2001:db8::/40	2001:db8:ff00::/40
Campus Location	2001:db8::/48	256	2001:db8::/48	2001:db8:ffff::/48
Building	2001:db8::/52	16	2001:db8::/52	2001:db8:ffff:f000::/48
End-System	2001:db8::/64	4096	2001:db8::/64	2001:db8:ffff:ffff::/64
Network Infrastructure				
Loopbacks	/128	2^128	Specific block chosen in the building prefix	
Links	/127	2^64	Specific block chosen in the building prefix	
Network Services	/56	256	Specific block chosen at a higher level	

Region 1 Campus 1 Building 1 Loopbacks

2001:db8:1010:11ff::/64

Region 1 Campus 1 Building 1 Links

2001:db8:1010:11fe::/64

Network Services (DNS,DHCP)

2001:db8:ffff:ff00::/56

Design Considerations

The Scope of IPv6 Deployment

- Planning and coordination is required
  - Network engineers & operators
  - Security engineers
  - Application developers
  - Desktop / Server engineers
  - Web hosting / content developers
  - Business development managers
  - …

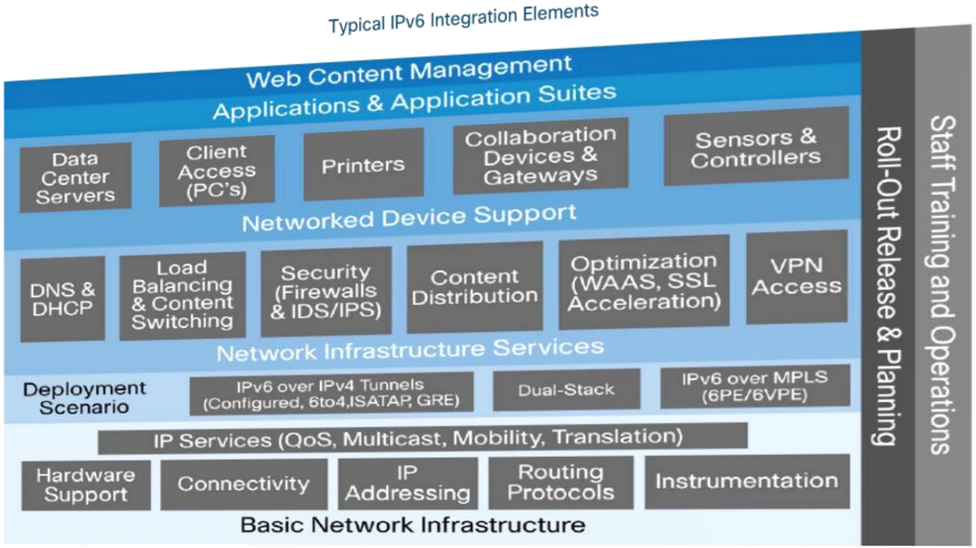

Dual Stack Mode

- Preferred Method, Versatile, Scalable and Highest Performance
- No Dependency on IPv4, runs in parallel on the same HW
- No tunnelling, MTU, NAT or performance degrading technologies
- Does require IPv6 support on all devices

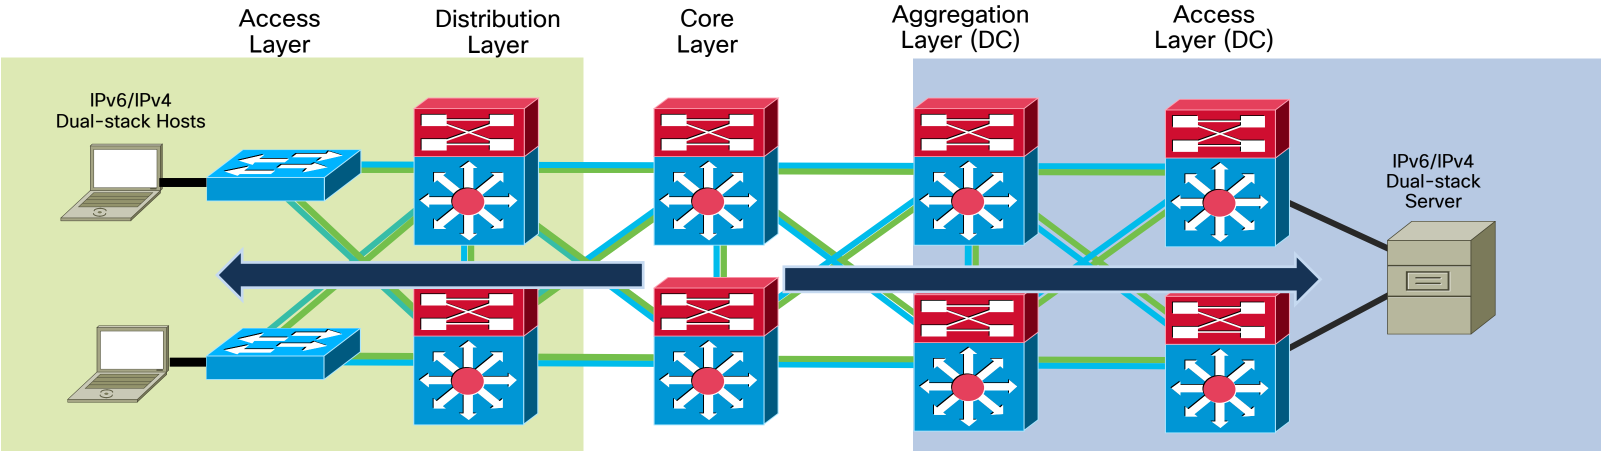

IPv4 & IPv6 Combined

TECRST-2001

- Should we use both on the same link at Layer 3?
- Routing protocols OSPFv3, EIGRP combined or separate?
- Fate sharing between the data and control planes per protocol

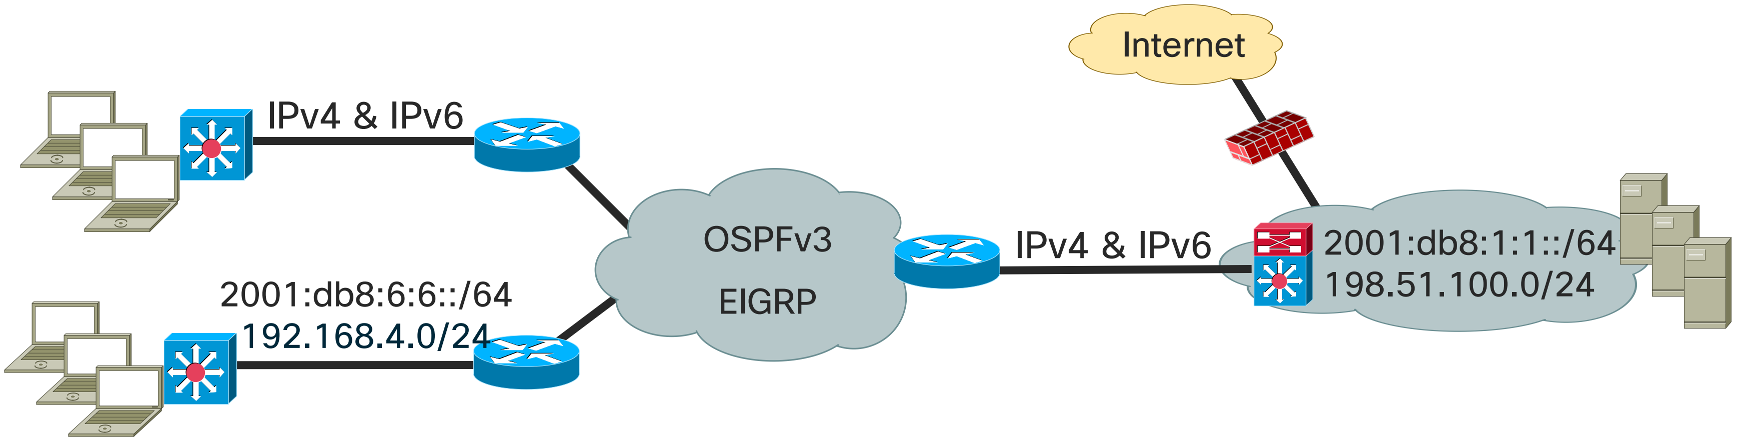

IPv6 /64 Deployments

- What if you have a server that needs more than 64k connections?
  - Use multiple addresses
- What if multiple apps want the same (well-known) port?
  - Give each application/container a separate IPv6 address, DNS name
- Possible solution: Assign a /64 to each server

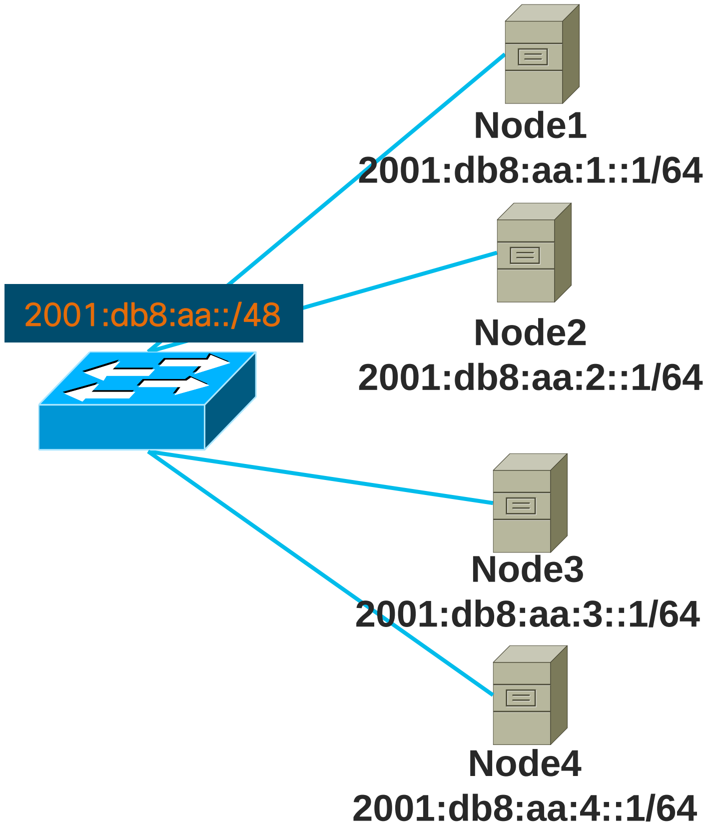

Client Provisioning DHCPv6 & SLAAC

- SLAAC address tracking
  - Radius accounting, CAM table scrapes
  - Microsoft now supports RDNSS in RA's

username=joe@example.org Acct-Session-Id=xyz Acct-Status-Type=Start Framed-IP-Address=192.0.2.1 Framed-IPv6-Address=fe80::d00d

username=joe@example.org Acct-Session-Id=xyz Acct-Status-Type=Alive Framed-IP-Address=192.0.2.1 Framed-IPv6-Address=fe80::d00d Framed-IPv6-Address=2001:db8::d00d Framed-IPv6-Address=2001:db8::d00d

DHCPv6 challenges

- MAC address for reservations, inventory, tracking
- Android does not support DHCPv6

- Understand the implications of switching methods
  - Inconsistent amongst the OS's

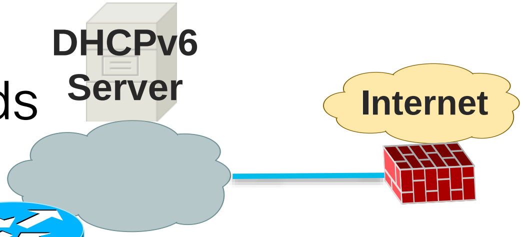

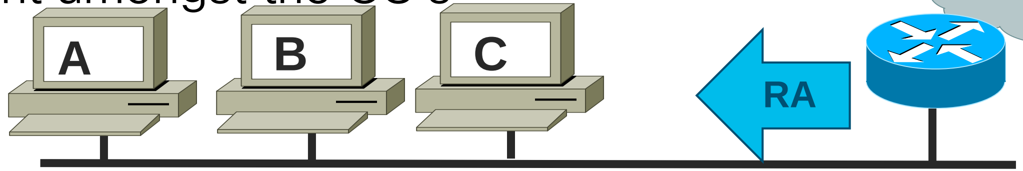

Disabling SLAAC/Privacy Addresses

- Enable DHCPv6 via the M flag
- Disable prefix auto configuration
- Enable router preference to high
- Enable DHCPv6 relay destination

```
interface fastEthernet 0/0
  ipv6 address 2001:db8:4646:acc1::1/64
  ipv6 nd managed-config-flag
  ipv6 nd prefix default no-autoconfig
  ipv6 nd router-preference high
  ipv6 dhcp relay destination 2001:db8:4646::café
```

IPv6 First Hop Redundancy Protocols

- FHRPs provide resilient default gateway
  - First hop address to end-stations
- IPv6 has a “built in” FHRP mechanism
  - Neighbor Unreachability Detection (NUD)
- HSRP, GLBP, and VRRP alternatives
  - Millisecond timers for fast convergence
- Preempt timers need to be tuned
  - To avoid black-holed traffic

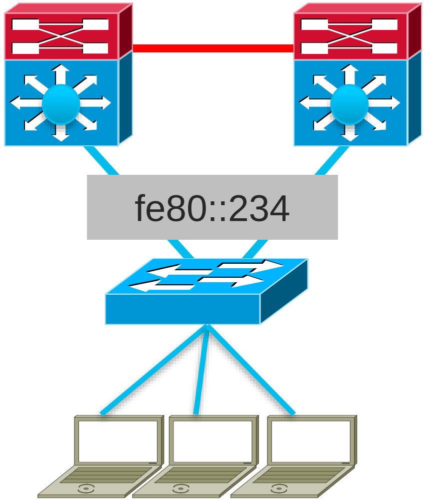

QoS CLI

- Class maps can match both IPv4 and IPv6 traffic
  - Can be broken into “ip” and “ipv6” matching
- Design principles still the same
  - Mark at the edge
  - Trust boundaries still apply
  - Queue sizing

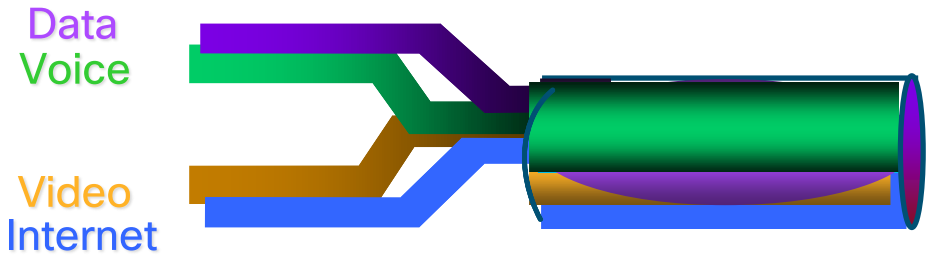


```
class-map match-any Critical_Data
  match dscp af21
class-map match-any Voice
  match dscp ef
class-map match-all Scavenger
  match dscp cs1
class-map match-any Bulk_Data
  match dscp af11
!
policy-map DISTRIBUTION
  class Voice
    priority percent 10
  class Critical_Data
    bandwidth percent 25
    random-detect dscp-based
  class Bulk_Data
    bandwidth percent 4
    random-detect dscp-based
  class Scavenger
    bandwidth percent 1
  class class-default
    bandwidth percent 25
    random-detect
```

Wireless LAN Controller BCP's

- WLC version 8.x increases support of IPv6
- CAPWAP, SNMP, NTP, Radius, Syslog, CDP, WebAuth
- Interface groups, same SSID over multiple VLAN's

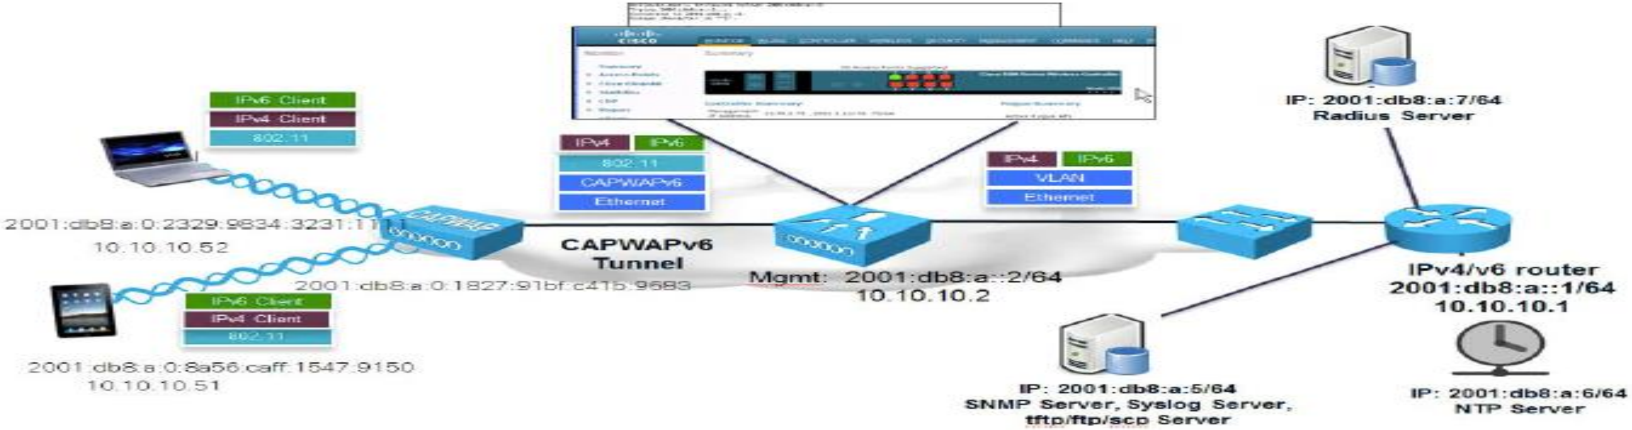

Wi-Fi Multicast Background

BRKRST-3304

- Radio is a shared media
  - Hosts must “awaken” to see if Multicast is for them
  - Multicast packets are not acknowledged or retransmitted
  - AP transmits bcast/mcast frames at the lowest possible rate

cisco Live!

Neighbour Discovery Multicast Suppression

- Scaling 802.11 multicast reliability issues
- NDP process is multicast "chatty", Unicasting reduces the effect
- Caching allows the Controller to "proxy" the NA, based on gleaning

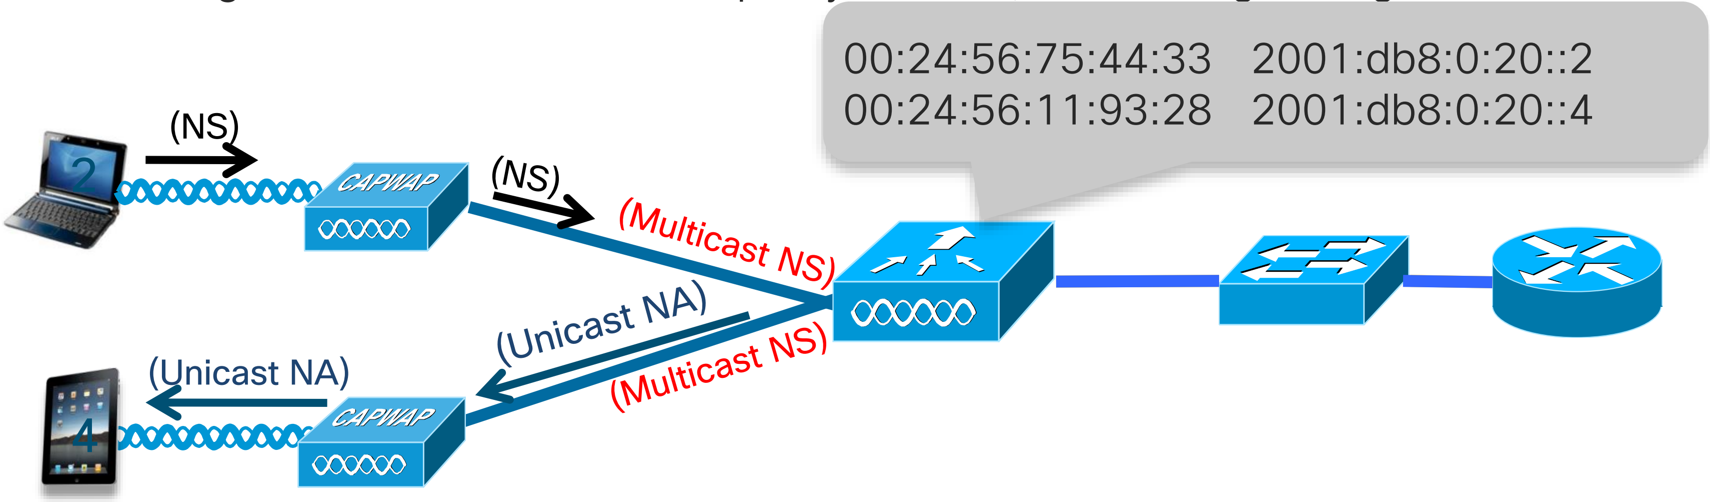

Router Advertisement Throttler

BRKRST-3304

- Rate limit RA's from the legitimate router
- Inspect the RS, convert the responding RA to L2 Unicast

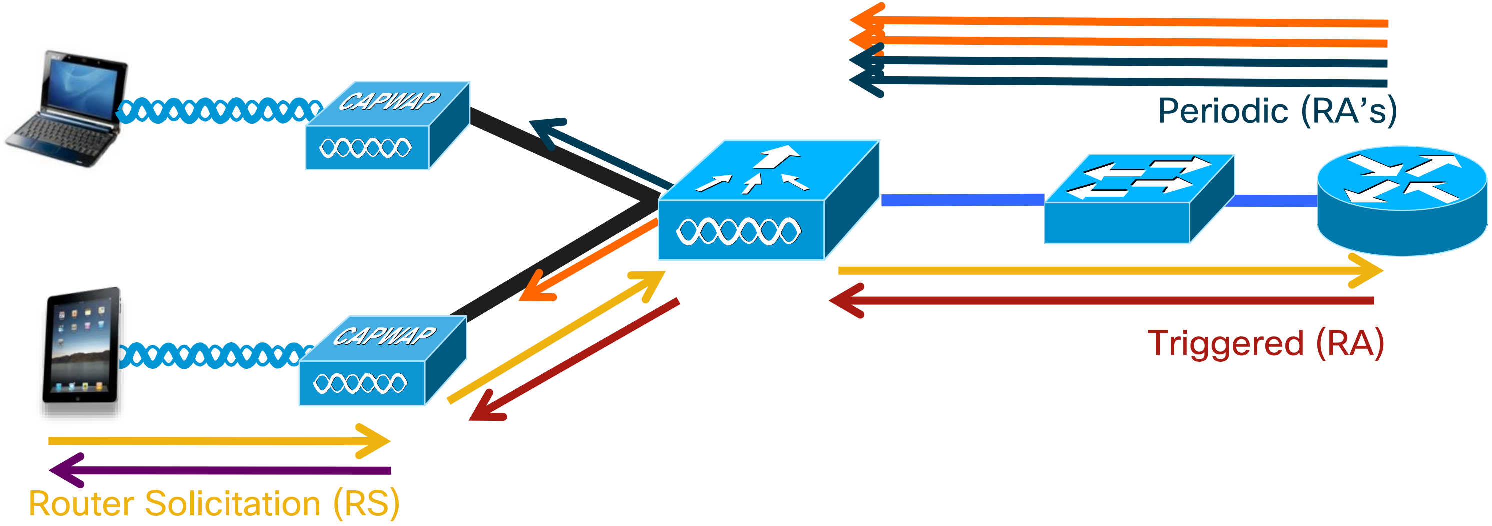

TECRST-2001

IPv6 FHS

RA Guard

Protection:
- Rogue or malicious RA
- MiM attacks

DHCPv6 Guard

Protection:
- Invalid DHCP Offers
- DoS attacks
- MiM attacks

Source/Prefix Guard

Protection:
- Invalid source address
- Invalid prefix
- Source address spoofing

Destination Guard

Protection:
- DoS attacks
- Scanning
- Invalid destination address

RA Throttle

Facilitates:
- Scale converting multicast traffic to unicast

ND Multicast Suppress

Reduces:
- Control traffic necessary for proper link operations to improve performance

Core Features **Advance Features** **Scalability & Performance**

IPv6 Snooping

IPv6 Operational Procedures

IPv6 and DNS

- Add an IPv6 address to a host, create AAAA record in DNS zone
- Only global or unique local, do not use link local addresses
- Uses PTR records in “ip6.arpa” for reverse mapping
- External servers/services should have PTR records configured
- Clients recieve DNS info from manual configuration, DHCPv6, RAs

Function	IPv4	IPv6
Hostname to IP Address	A Record `www.example.com. A 192.168.30.1`	AAAA Record (Quad A) `www.example.com AAAA 2001:db8:c18:1::2`
IP Address To Hostname	PTR Record `1.30.168.192.in-addr.arpa. PTR www.example.com.`	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 PTR www.example.com.`

DNS Records Over IP

- IPv6 only networks require IPv6 DNS servers
  - Each DNS server should have its own /64 (2001:db8:4646::53)
- Dual stack networks can use either protocol as transport
  - IPv6 results can be carried over IPv4

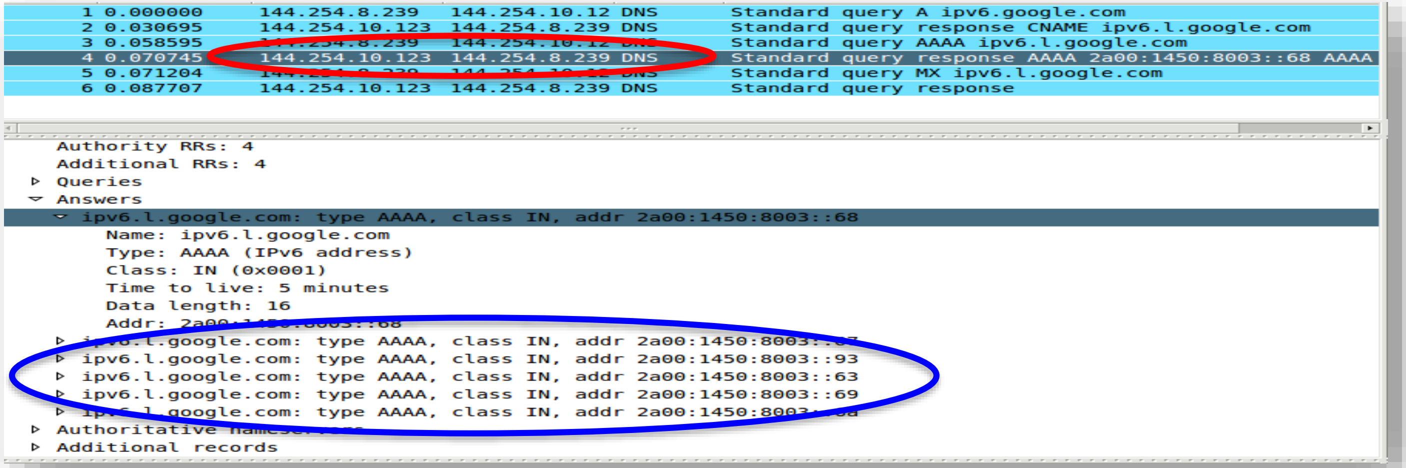

cisco Live!

IPv6 DNS Operation

- Successful DNS lookups over IPv6 require:
  - Each zone provide at least one name server with an IPv6 address
- Clients accessing your site often do so via recursive servers
  - RS request your DNS info on behalf of the client
- Glue records bind the name server (NS1.example.com)
  - To the parent zone (example.com)

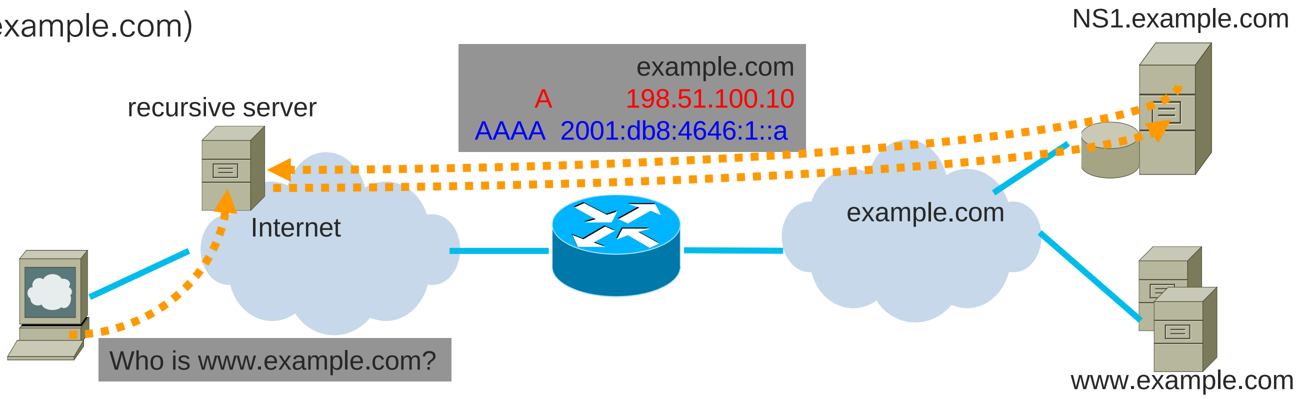

Conclusion

Key Take Away

- Gain Operational Experience now
- Control **IPv6 traffic** as you would IPv4
- **Lead** your OT/LOB's into the Internet

Recommended Reading

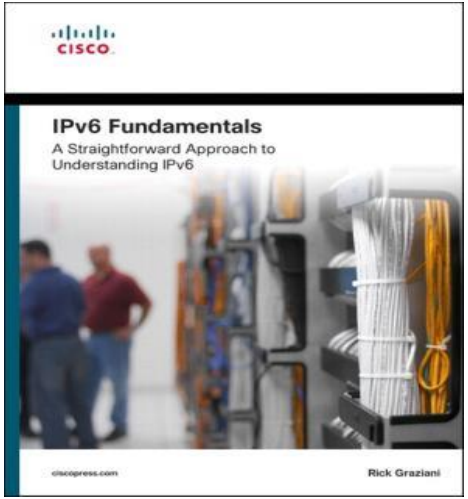

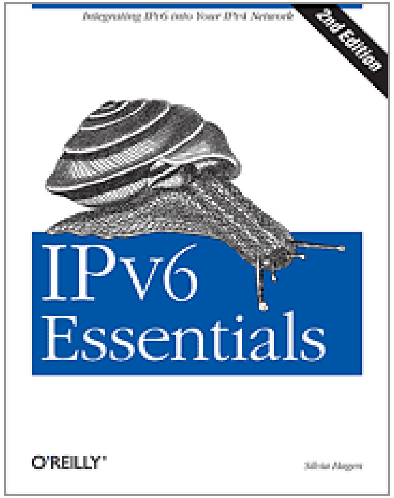

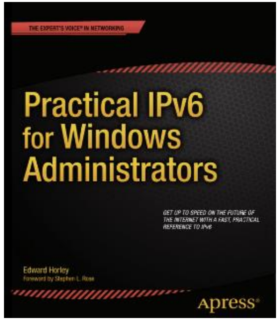

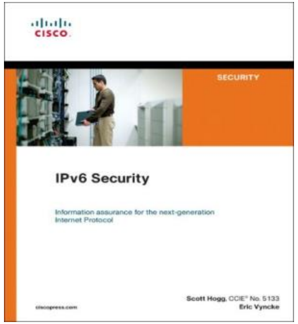

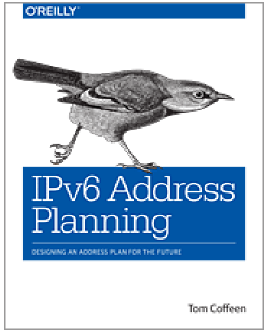

http://www.ciscopress.com/store/ipv6-design-and-deployment-livelessons-9780134655512

Code: MARTIN60
Offer: 60% off
ISBN: 9780134655512

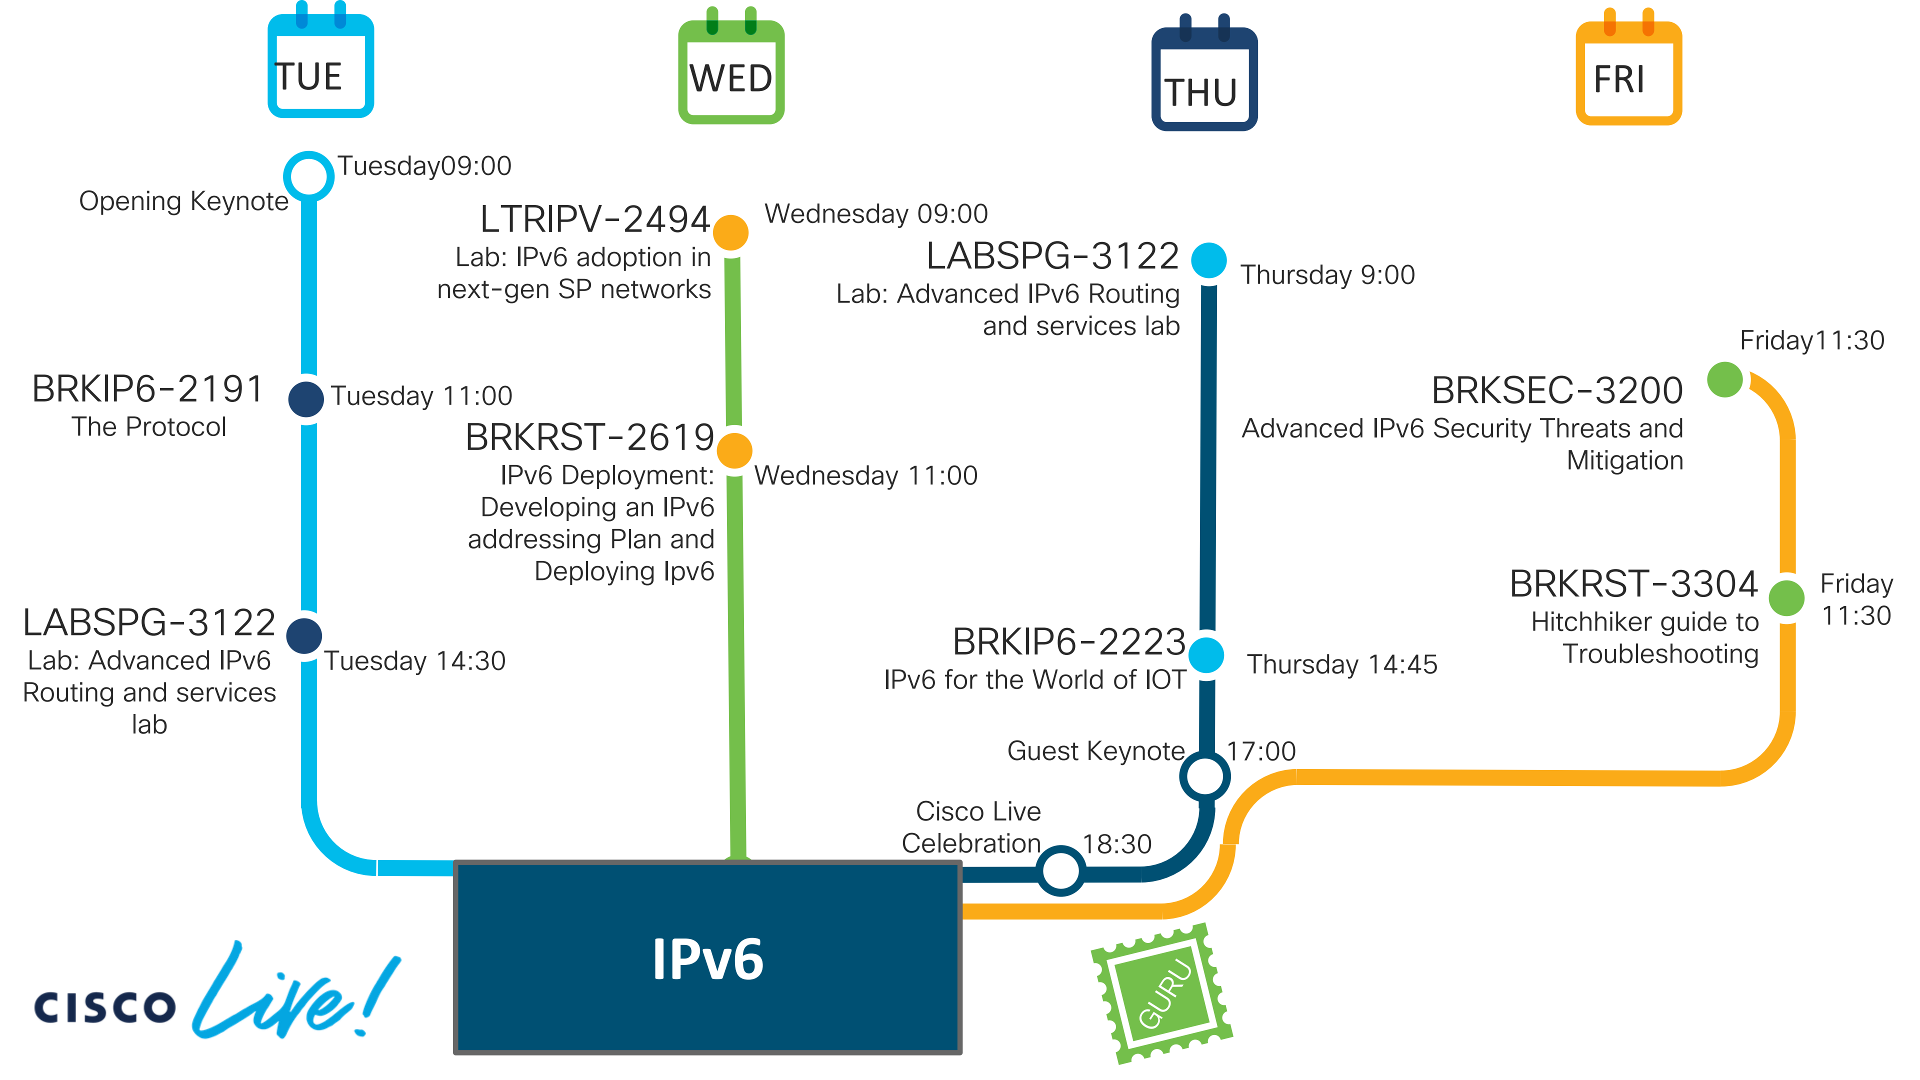
TUE
WED
THU
FRI
Tuesday09:00
Opening Keynote
LTRIPV-2494
Lab: IPv6 adoption in next-gen SP networks
Wednesday 09:00
LABSPG-3122
Lab: Advanced IPv6 Routing and services lab
Thursday 9:00
Friday11:30
BRKSEC-3200
Advanced IPv6 Security Threats and Mitigation
BRKIP6-2191
The Protocol
Tuesday 11:00
BRKRST-2619
IPv6 Deployment: Developing an IPv6 addressing Plan and Deploying Ipv6
Wednesday 11:00
BRKRST-3304
Hitchhiker guide to Troubleshooting
Friday 11:30
LABSPG-3122
Lab: Advanced IPv6 Routing and services lab
Tuesday 14:30
BRKIP6-2223
IPv6 for the World of IOT
Thursday 14:45
Guest Keynote
17:00
Cisco Live Celebration
18:30
IPv6
GURU
cisco Live!

Complete your online session survey

- Please complete your session survey after each session. Your feedback is very important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (starting on Thursday) to receive your Cisco Live t-shirt.
- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Content Catalog on ciscolive.com/emea.

Cisco Live sessions will be available for viewing on demand after the event at ciscolive.com.

Continue your education

Demos in the Cisco campus

Walk-in labs

Meet the engineer 1:1 meetings

Related sessions

Thank you

You make **possible**