



ROCKY MOUNTAIN IPV6 SUMMIT

APRIL 9, 2008


INTRODUCTION TO THE IPV6 PROTOCOL

Scott Hogg

GTRI - Director of Advanced Technology Services

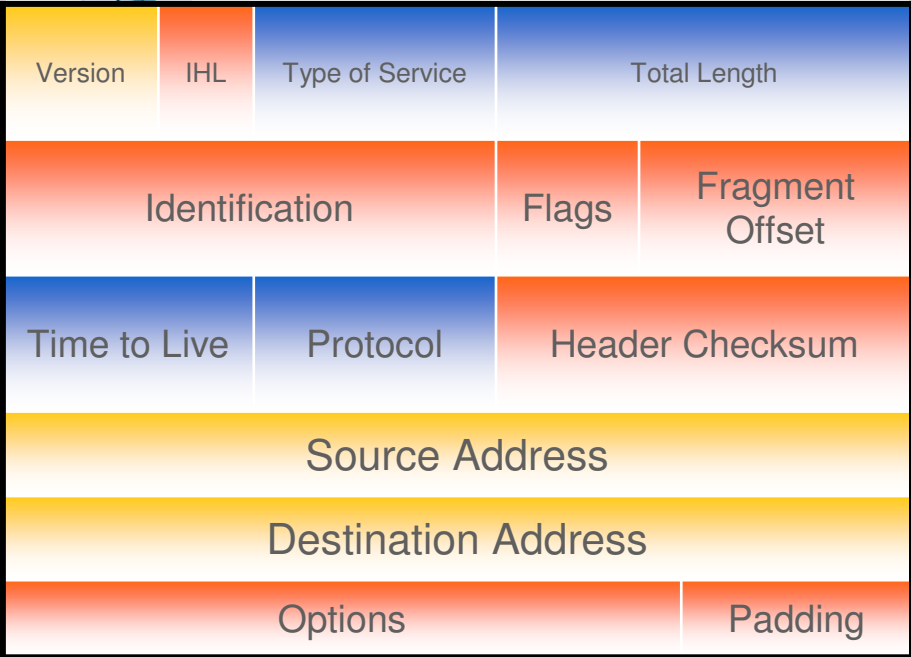
CCIE #5133, CISSP

Rocky Mountain IPv6 Task Force

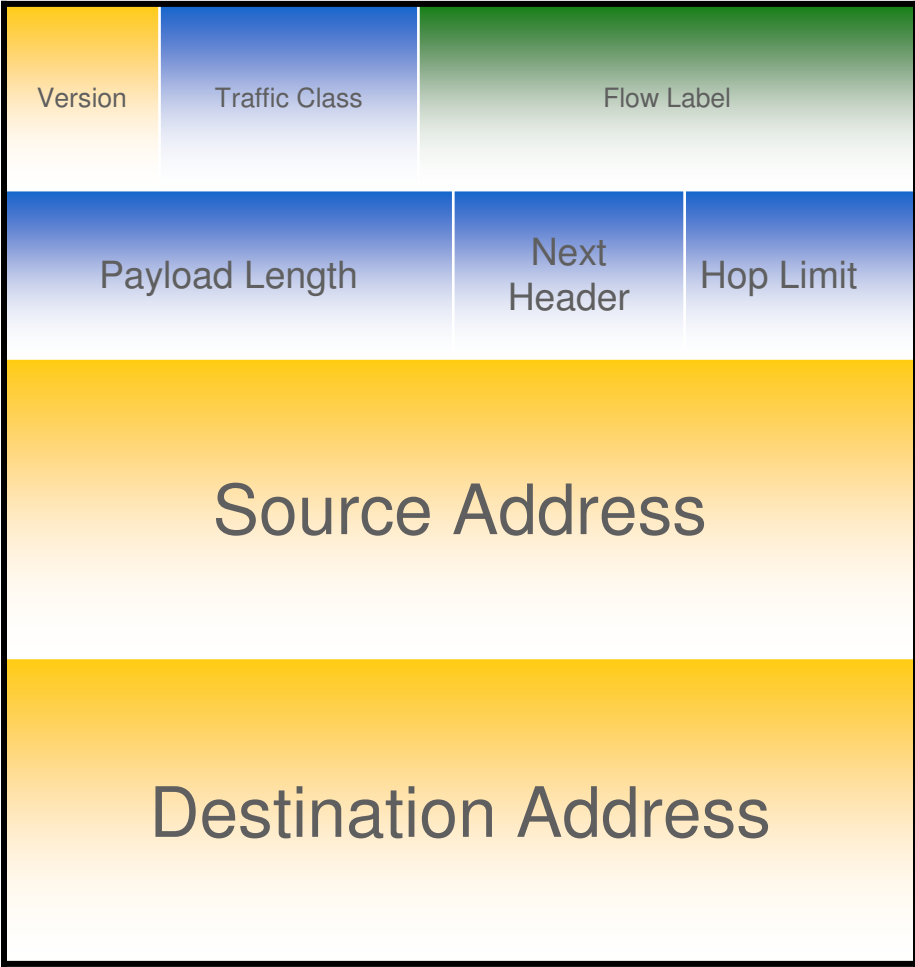


IPv6 HEADER

IPv4 Header 20 bytes






IPv6 Header, 40 bytes fixed



- Legend**
- field's name kept from IPv4 to IPv6
 - fields not kept in IPv6
 - Name & position changed in IPv6
 - New field in IPv6

IPv6 EXTENSION HEADERS



Next Header Field:

0 – Hop-by-Hop Options
60 – Destination Options
(If Routing header is used)

43 – Routing

44 – Fragment

46 – RSVP

51 – AH

50 – ESP

88 – EIGRP

89 – OSPF

6 – TCP

17 – UDP

58 – ICMPv6

135 – Mobility Header

59 – None (no next header)

8-bits	8-bits	
Option Type	Option Data Length	Option Data (Variable Length)

(Next)



INCREASED IPv6 ADDRESSES

- IPv6 Increased Src/Dst Address to 128 bits
- $2^{128} = 340,282,366,920,938,463,463,374,607,431,768,211,456$ addresses (~340 undecillion)
- If each IP address equaled one gram
 - IPv4 would be 1/76th the weight of the Empire State Building
 - IPv6 would be 56.7 billion X the Earth's weight
- 67 billion billion (6.65×10^{23}) addresses per cm² of the Earth's surface
- 1246 IPv6 addresses per square meter of the area of the Milky Way galaxy
- That ought to be enough!

ADDRESSING : FORMAT PREFIX

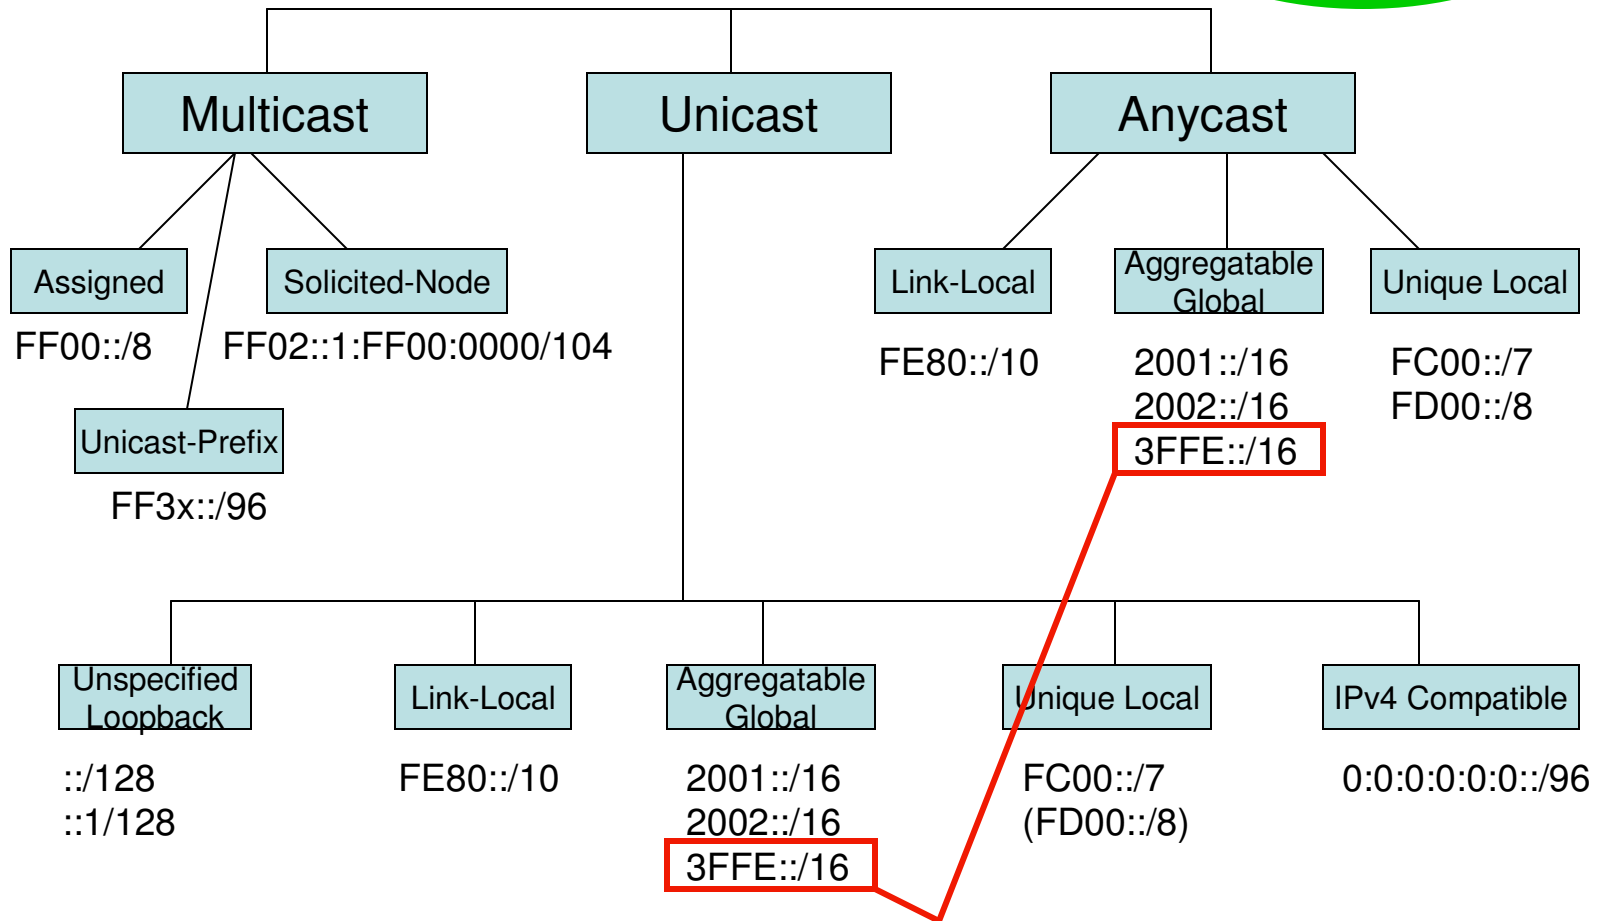
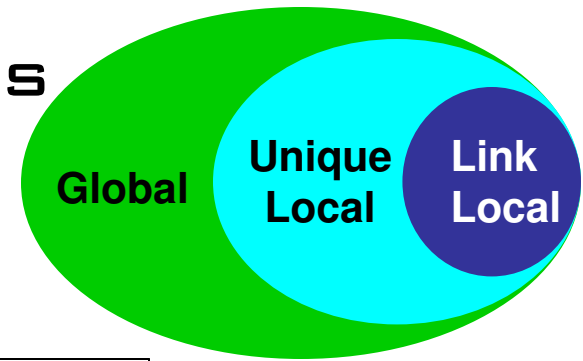
• Reserved (::0/128)	0000 0000
• Unassigned	0000 0001
• Reserved for NSAP Allocation	0000 001
• Reserved for IPX Allocation – later deprecated	0000 010
• Unassigned	0000 011
• Unassigned	0000 1
• Unassigned	0001
• Aggregatable Global Unicast Addresses (2001::/16)	001
• Provider-Based Unicast Address	010
• Unassigned	011
• Rsvd for Neutral-Interconnect-Based Unicast Addr	100
• Unassigned	101
• Unassigned	110
• Unassigned	1110
• Unassigned	1111 0
• Unassigned	1111 10
• Unique Local (FC00::/10)	1111 110
• Unassigned	1111 1110 0
• Link Local Use Addresses (FE80::/10)	1111 1110 10
• Site Local Use Addresses (FEC0::/10) - deprecated	1111 1110 11
• Multicast Addresses (FF00:/8)	1111 1111



IPv6 ADDRESSING NOTATION

- 128 bits get converted into more readable form
 - 0011 1111 1111 1110 1001 0000 1110 0000 0000 0000
0000 0011 0000 0000 0000 0000 / 0000 0000 0000 0000
0000 0000 0101 0000 0000 0000 0000 0000 0000 0000
0000 0000
- Convert bits to hex
 - 3FFE:90E0:0003:0000:0000:0050:0000:0000
- Reduce by removing leading zeros
 - 3FFE:90E0:3:0:0:50:0:0
- Use :: to consolidate multiple zeros – only once
 - 3FFE:90E0:3::50:0:0
 - or
 - 3FFE:90E0:3:0:0:50::
- Prefix format/notation
 - 3FFE:90E0:3::/64

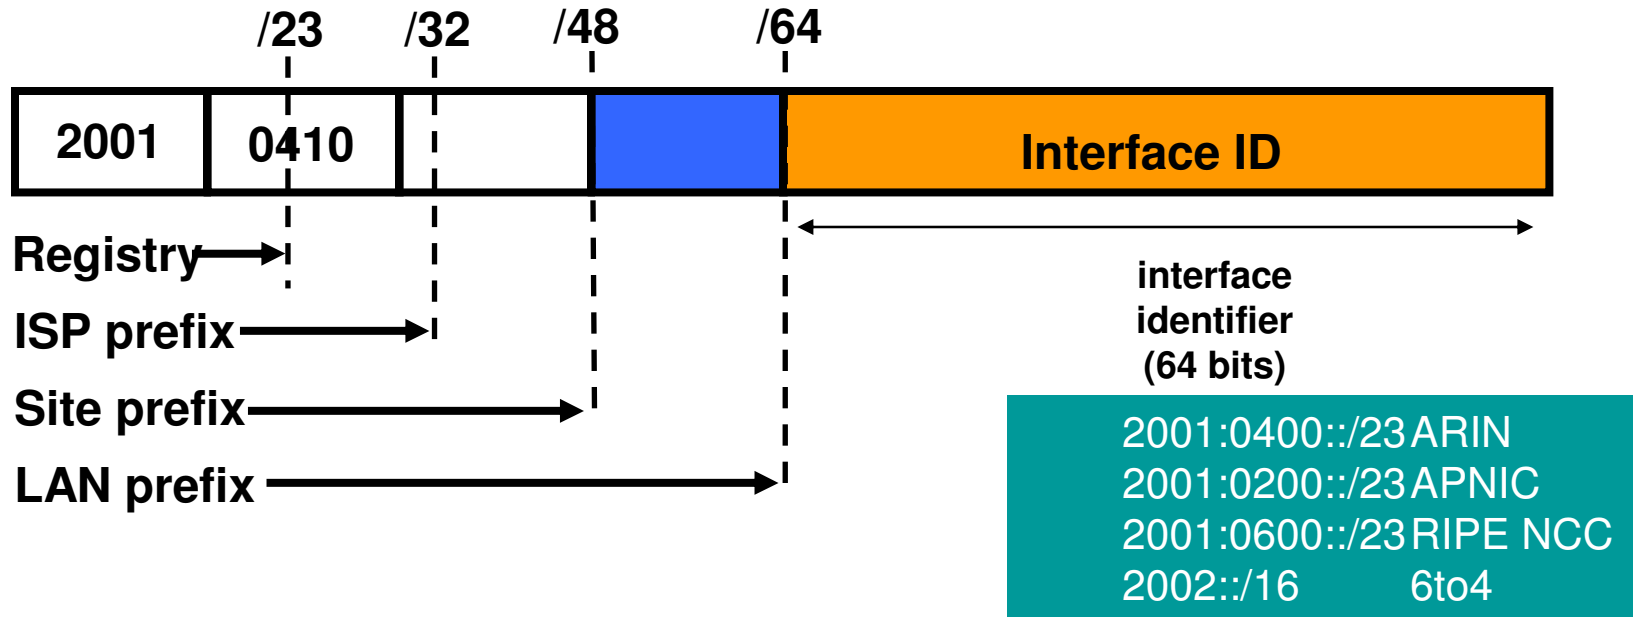
IPv6 ADDRESS TYPES



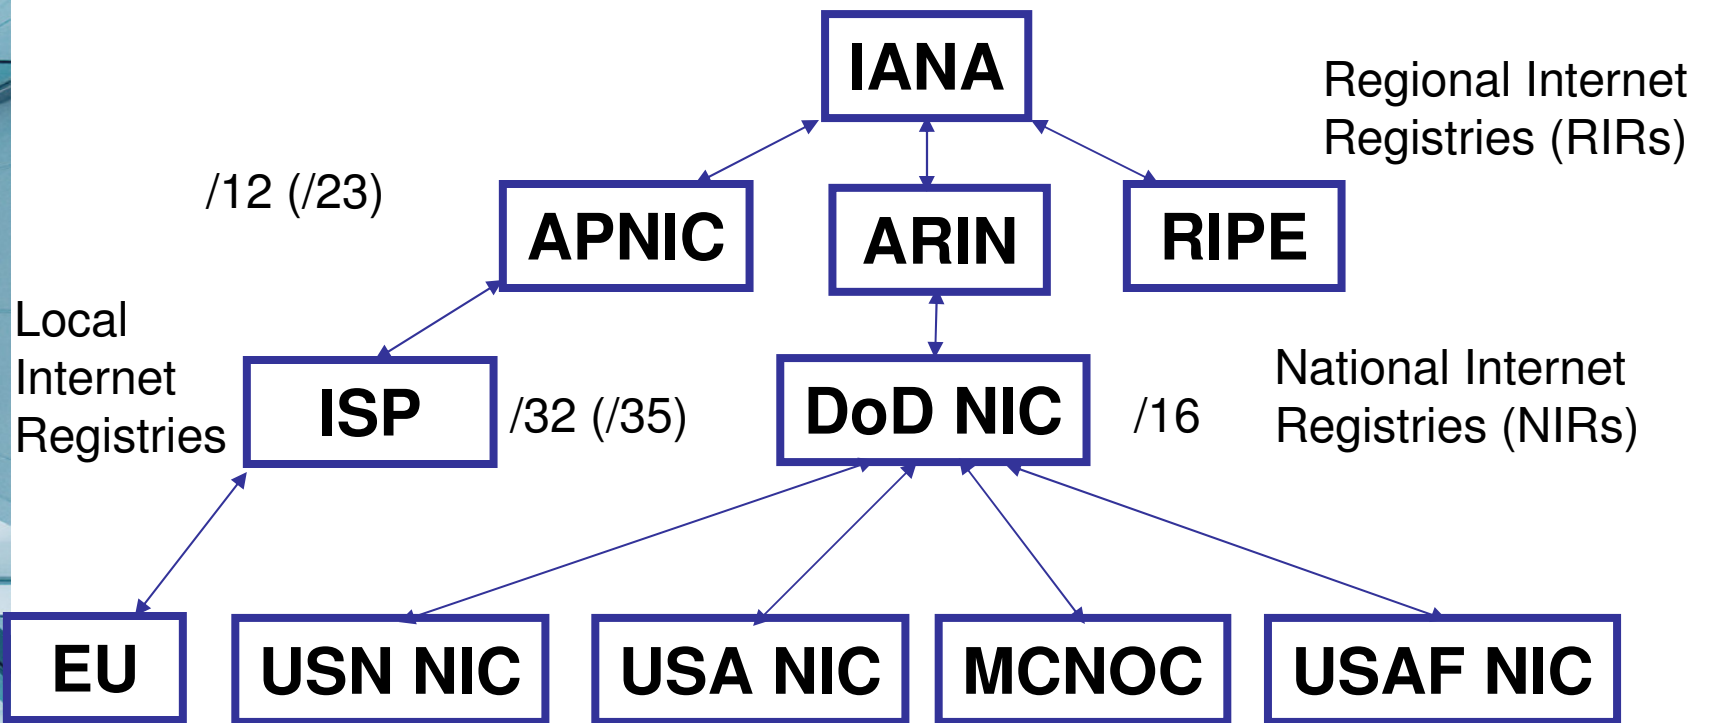
6Bone - Deprecated

ADDRESS ALLOCATION POLICY

- The allocation process is under review by the Registries:
 - IANA allocates 2001::/16 to registries
 - Each registry gets a /23 prefix from IANA
 - Formerly, all ISP were getting a /35
 - With the new policy, Registry allocates a /32 prefix to an IPv6 ISP
 - Then the ISP allocates a /48 prefix to each customer (or potentially /64)



DoD IPv6 ADDRESS MANAGEMENT



IPv6 ADDRESSING PLAN

Subnet Prefix
2001:0DB8:0001:7AB5::/64

Site Prefix
2001:0DB8:0001::/48

ISP Prefix
2001:0DB8::/32

This is where most of your addressing work will be done



WAN/Region Subnet

Site Subnet
8 bits = 256 sites

LAN Subnet
6 bits = 64 LANs

ICMPv6

- More powerful than ICMPv4
- ICMPv6 uses IPv6 extension header # 58 (RFC 2463)

Type	Description
1	Destination Unreachable
2	Packet too Big
3	Time exceeded
4	Parameter problem
128	Echo Request
129	Echo Reply
130	Multicast Listener Query – sent to ff02::1 (all nodes)
131	Multicast Listener Report
132	Multicast Listener Done – sent to ff02::2 (all routers)
133	Router Solicitation (RS) – sent to ff01::2 (all routers)
134	Router Advertisement (RA) – sent to ff01::1 (all nodes)
135	Neighbor Solicitation (NS) – sent to ff02:0:0:0:0:1:ff00:::/104
136	Neighbor Advertisement (NA)
137	Redirect message

Diagram annotations:

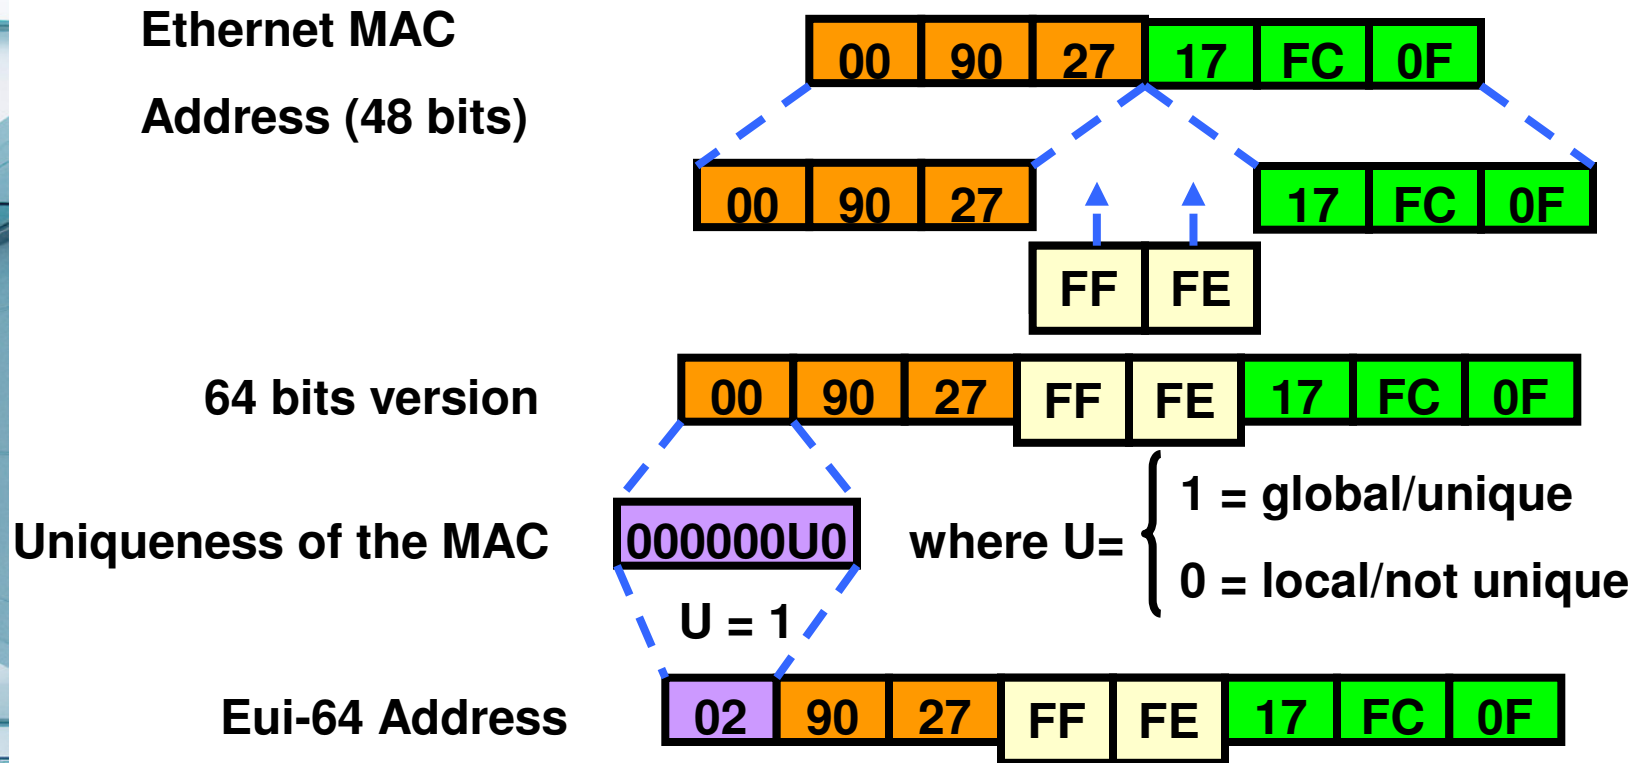
- Teal arrow from 128 to PING
- Grey arrow from 129 to MLD
- Blue arrow from 130 to Prefix Advertisement
- Blue arrow from 133 to Prefix Advertisement
- Red arrow from 135 to ARP Replacement
- Green arrow from 137 to Router Redirection



IPv6 AUTO-CONFIGURATION

- IPv4 Address Configuration
 - Static or DHCP
 - IPv4 Address, Subnet Mask, Default Gateway
 - Domain Name, DNS Resolver
- IPv6 Address Configuration
 - Static addressing
 - DHCPv6 (stateful configuration)
 - Stateless DHCPv6 on router – RFC 3315
 - Neighbor Discovery (stateless configuration)

INTERFACE ID - EUI-64



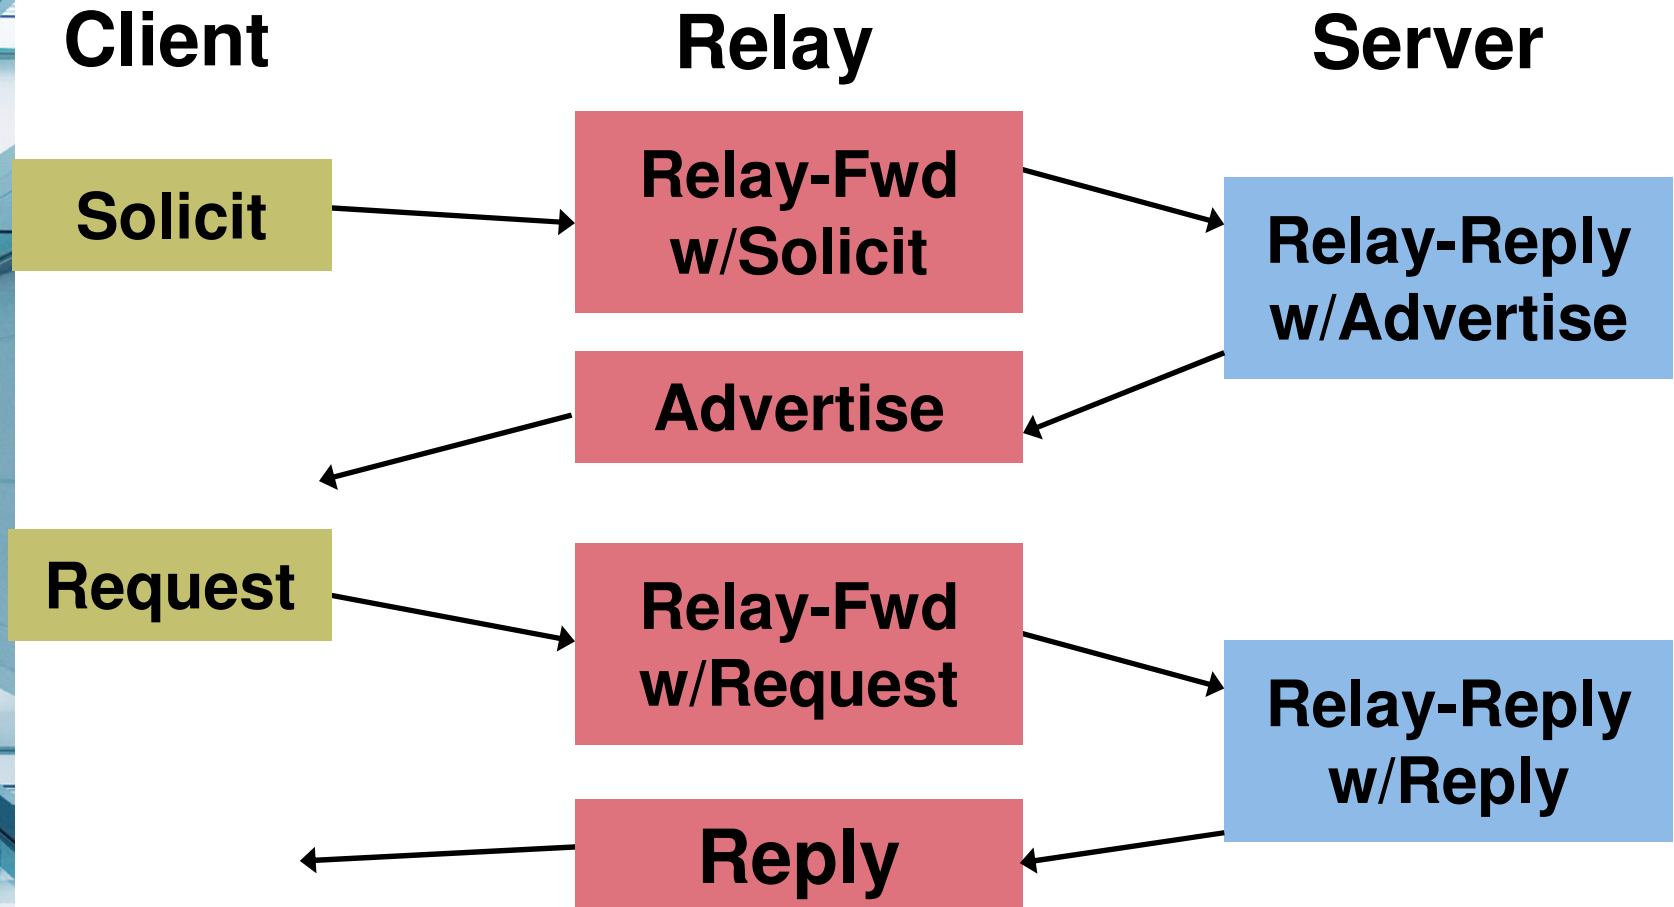
EUI-64 address is formed by inserting "FFFE" and ORing a bit identifying the uniqueness of the MAC address.



DHCPv6

- Similar to DHCPv4, but hosts check RAs first.
- RAs contain A, M, and O bits
 - A-bit causes SLAAC (stateless address autoconfiguration)
 - M-bit and O-bit indicate SAC (stateful address configuration)
 - DHCPv6 can be used concurrently with stateless configuration (M bit in Router Advertisements).
- Can be used for renumbering using Reconfigure-Init message.
- DHCPv6 provides DNS information, stateless autoconfiguration does not, for now.

DHCPV6 OPERATION



- All_DHCP_Relay_Agents_and_Servers (FF02::1:2)
- All_DHCP_Servers (FF05::1:3)
- DHCP Messages: Clients listen UDP port 546. Servers and relay agents listen on UDP port 547

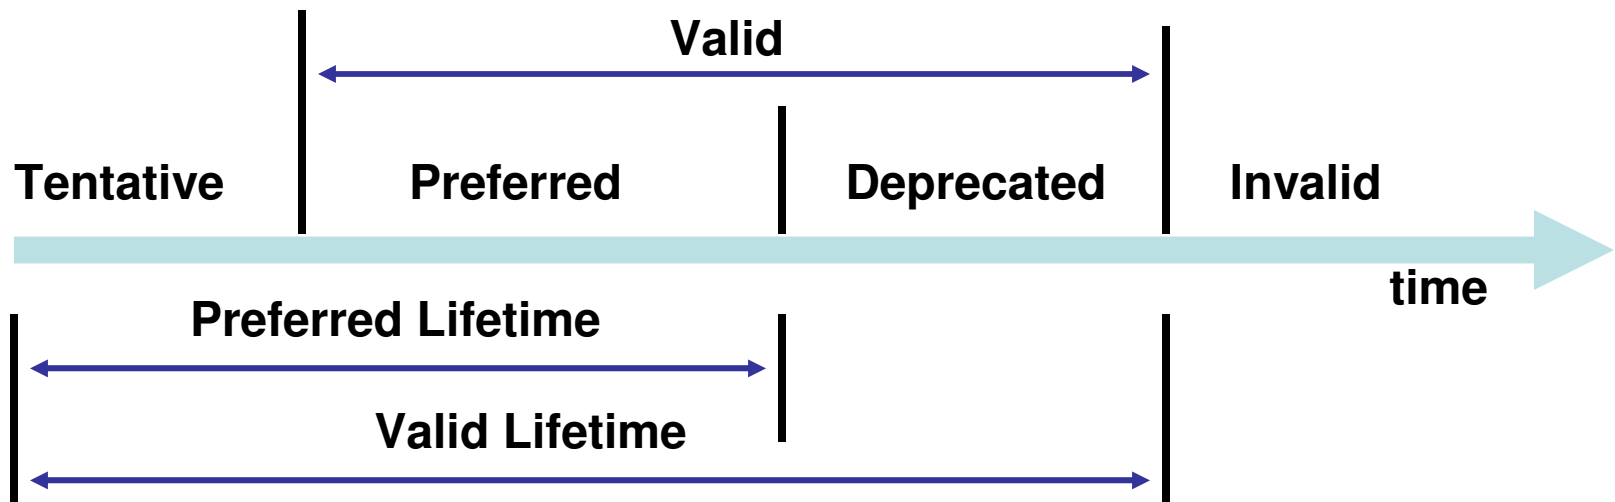


DHCPv6

- DHCPv6 clients listen on UDP port 546
- DHCPv6 servers and relay agents listen on UDP port 547
- Solicit, Advertise, Request, Reply messages
- Managed Address Configuration Flag - M flag
 - When set to 1 it tells hosts to get stateful address through DHCPv6
- Other Stateful Configuration Flag - O flag
 - When set to 1 it tells host to use DHCPv6 to get other config info
 - M=0, O=0
 - End host uses stateless auto-configuration
 - M=0, O=1
 - DHCPv6 stateless, autoconfigured address but DNS and other information is added with DHCPv6 (i.e. from a router)
 - M=1, O=0
 - Not practical, no DNS info
 - M=1, O=1
 - End host uses DHCPv6 for all information

IPv6 ADDRESS LIFETIMES

- Tentative : the address is in the process of being verified as unique
- Preferred : a node can send and receive unicast traffic to and from a preferred address
- Deprecated : the address is still valid, but using it for new communication is discouraged
- Invalid : the address can no longer send unicast traffic to or receive it from a node. An address enters this state after the valid lifetime expires





OTHER IPv6 FEATURES

- IPv6 requires every network link be capable of minimum MTU of 1280 bytes
- IPv6 routers don't fragment packets
 - Hosts perform their own Path MTU Discovery
 - MTU path discovery uses ICMP "packet too big" error messages
- Provider selection (based on policy, performance, cost, ...)
- Host mobility (route to current location)
- Auto-readdressing (route to new address)
 - (Use IPv6's routing extension header)

QUESTION AND ANSWER

Q:

&

A:

SHogg@GTRI.com
Scott@HoggNet.com

Mobile: 303-949-4865