IPv6 Addressing (Session 13228)

#SHAREorg

SHARE
Technology · Connections · Results

Kevin Manweiler, CCIE 5269
kmanweil@cisco.com

Junnie Sadler, CCIE 7708
jrsadler@cisco.com

Date of Presentation Tuesday, August 13, 2013: 3:00 PM-4:00 PM

Session Number **(13228)**

SHARE in Boston

Agenda

- IPv6 Introduction
- Addressing
- Summary and Q&A

IPv6 Introduction

How different is IPv6 from IPv4?

- Viewed at a distance, not very different
  - IPv4 with bigger addresses – prefix and host part defined
  - Routing – uses same or similar protocols (BGP4+, RIPng, IS-IS, OSPFv3)
  - Firewalls instead of network address translators in most cases
  - DNS AAAA records instead of A records
  - Address management: Auto-configuration or DHCPv6 as opposed to DHCP
  - QoS procedures identical
  - IPsec identical but carried in an IPv6 extension header instead of separately
- In detail, many differences
  - 128 bits is a *long* address – expressed in hex quads, not decimal
  - Attention to scaling and improved operation
  - Address amplification not required
  - Multicast uses new protocols - MLD

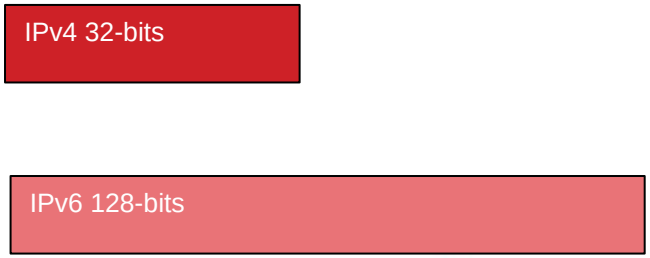

RIRs – Regional Internet Registries

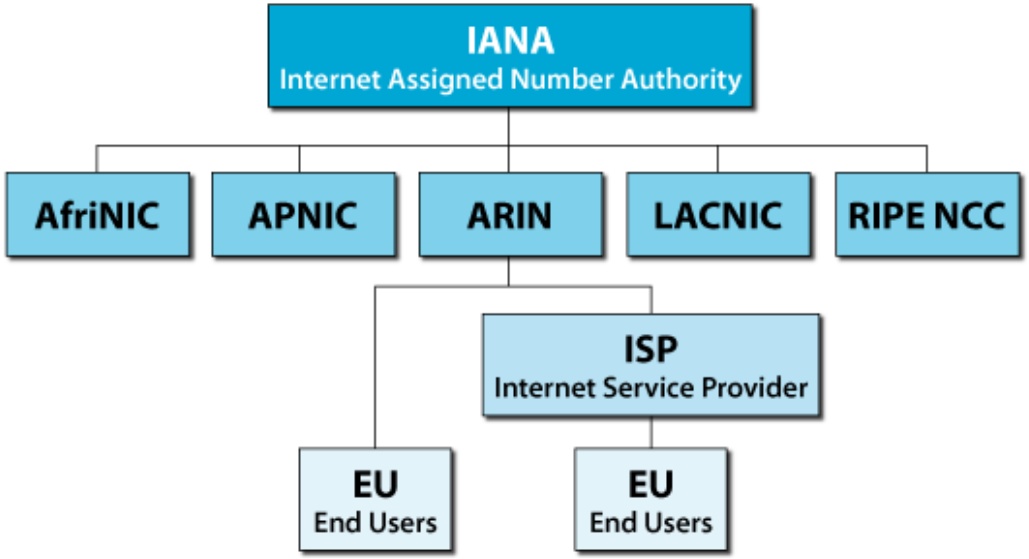

- Apply for Provider Independent address space from appropriate RIR

- www.afrinic.net
- www.arin.net
- www.apnic.net
- www.lacnic.net
- www.ripe.net

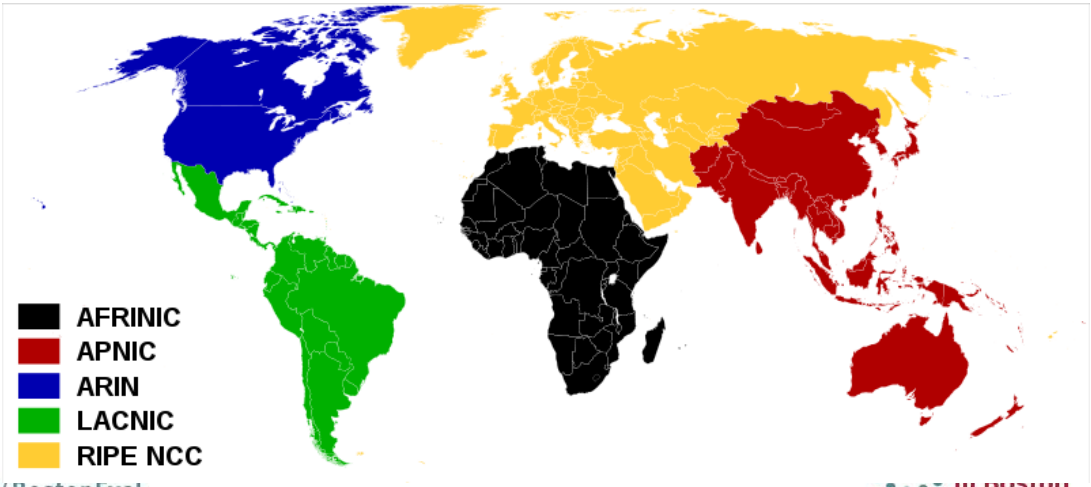

Address Allocation Process

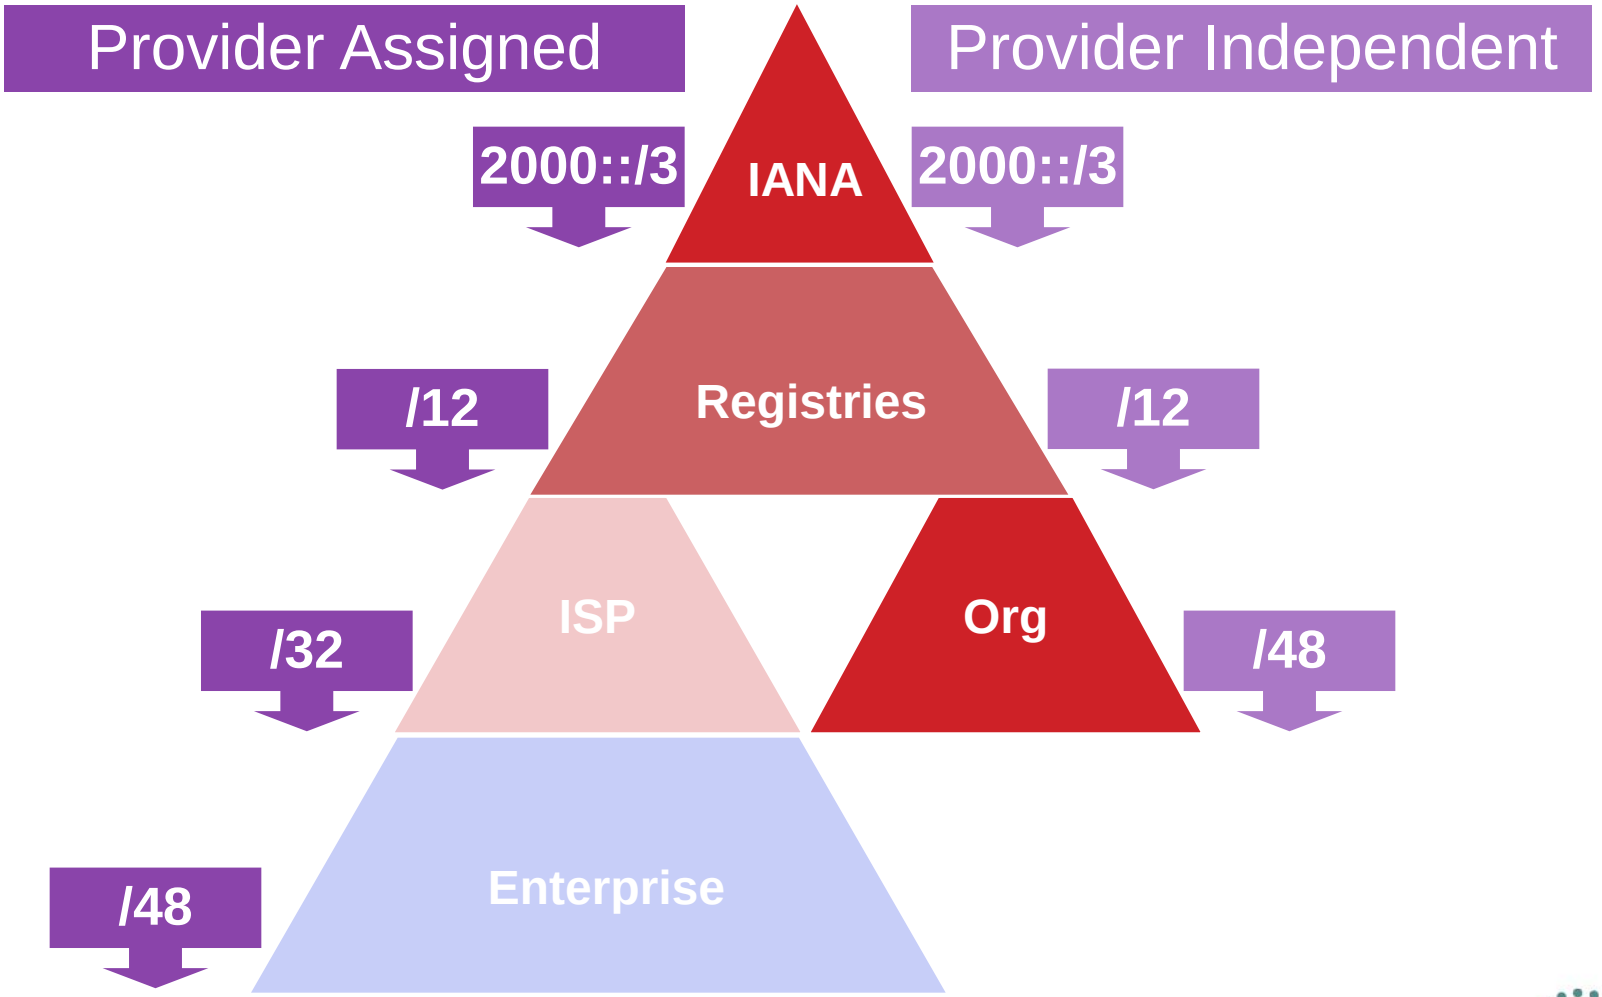

IPv6 Address Allocation Process

Partition of Allocated IPv6 Address Space (example)

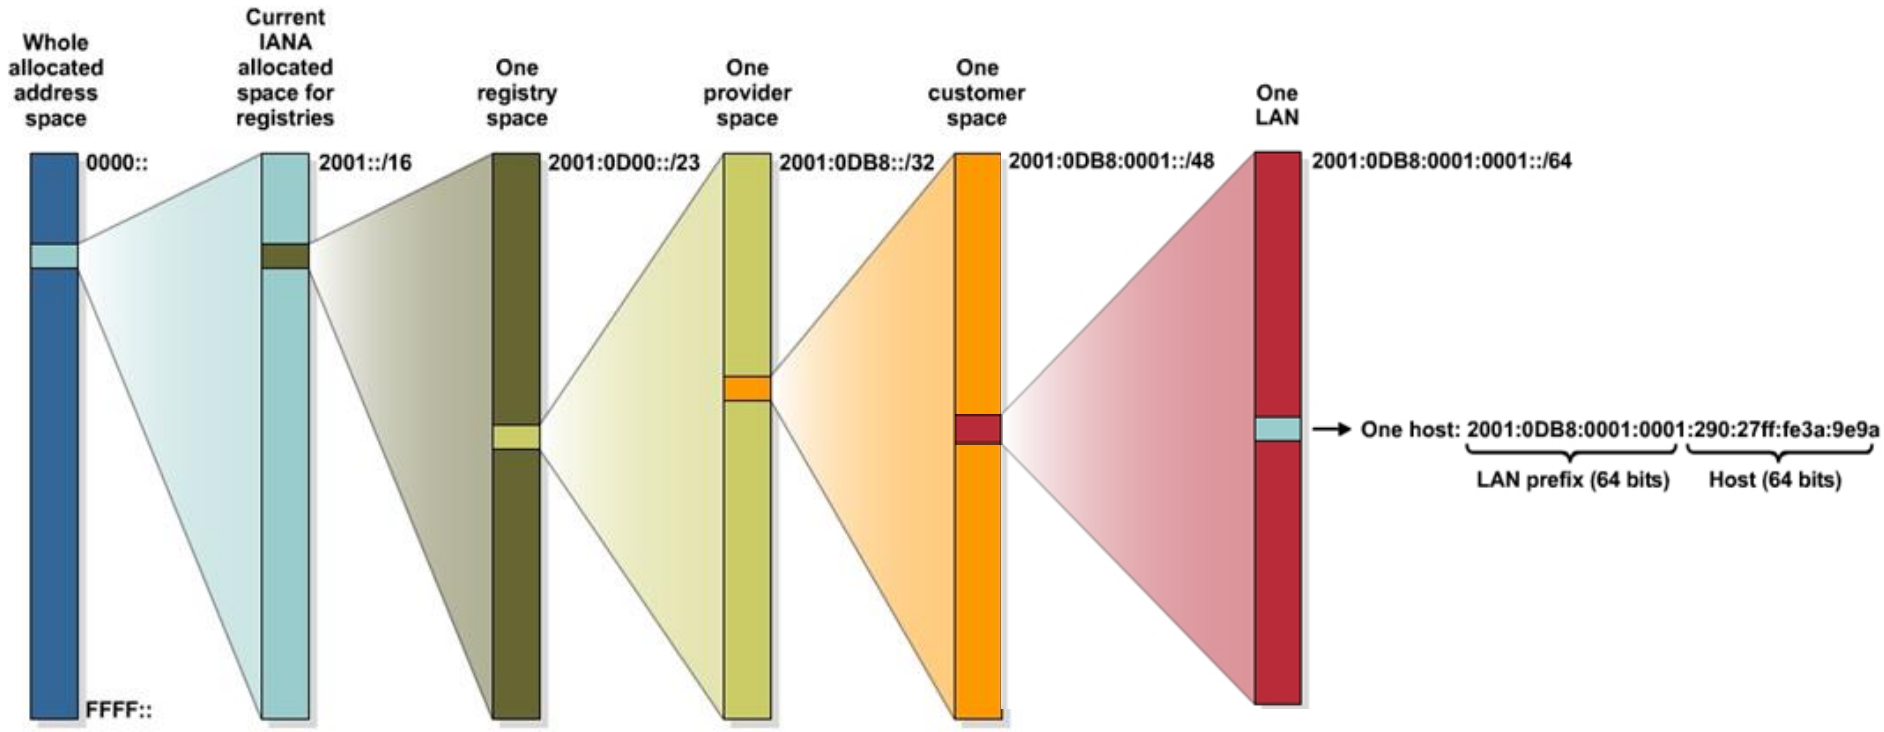

Address Considerations

- Provider Independent or Provider Assigned
- Global, ULA, ULA + Global
- Prefix-length allocation
- Subnet length
- Address allocation method
- IP Address Management (IPAM)

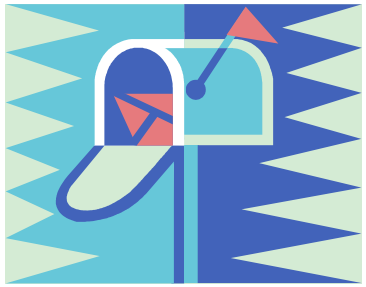

IPv6 Addressing

IPv6 Addressing

IPv4 32-bits

192.168.123.231

IPv6 128-bits

2001:0db8:0000:130F:0000:0000:087C:140B

$2^{32} = 4,294,967,296$

$2^{128} = 340,282,366,920,938,463,463,374,607,431,768,211,456$

$2^{128} = 2^{32} * 2^{96}$

2^{96} = 79,228,162,514,264,337,593,543,950,336 times the number of possible IPv4 Addresses (79 trillion trillion)

Important IPv6 Prefix Notations

* Nalini Elkins

/8	11111111xxxxxxxx.xxxxxxxxxxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx	00XX:: FFXX::
/16	1111111111111111.xxxxxxxxxxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx	00XX:: FFXX::
/32	1111111111111111.1111111111111111. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx	0000:0000:: FFFF:FFFF::
/48	1111111111111111.1111111111111111. 1111111111111111.xxxxxxxxxxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx	0000:0000:0000:: FFFF:FFFF:FFFF::
/56	1111111111111111.1111111111111111. 1111111111111111.11111111xxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx	0000:0000:0000:00xx:: FFFF:FFFF:FFFF:FFxx::
/64	1111111111111111.1111111111111111. 1111111111111111.1111111111111111. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx. xxxxxxxxxxxxxxxx.xxxxxxxxxxxxxxxx	0000:0000:0000:0000:: FFFF:FFFF:FFFF:FFFF::

What's my IPv6 Address?

SHARE
Technology · Connections · Results

2001:db8:1f05:347:213:f7ff:fe53:77af

- Lower 64 bits belong to the host
  - Configured Manually
  - Derived from OUI-48/64 ("MAC") address
  - Randomly selected and changing (RFC4941 privacy addressing)
- Upper 64 bits are the prefix and subnet
  - Prefix learned by hosts using SLAAC (radvd) or DHCPv6
  - DNS servers: manually configured, learned through DHCP (RFC3736)
- See your address from the command line
  - Windows: **ipconfig /all** or **netsh interface ipv6 show addresses**
  - MacOS/Linux: **ifconfig** or **ip -6 address show**

IPv6 Addresses

- IPv6 addresses are 128 bits long
  - Segmented into 8 groups of four HEX characters
  - Separated by a colon (:)
  - 50% for network ID, 50% for interface ID
  - Network portion is allocated by Internet registries 2^64 (1.8×10^{19})
  - Still leaves us with ~ 3 billion network prefixes for each person on earth

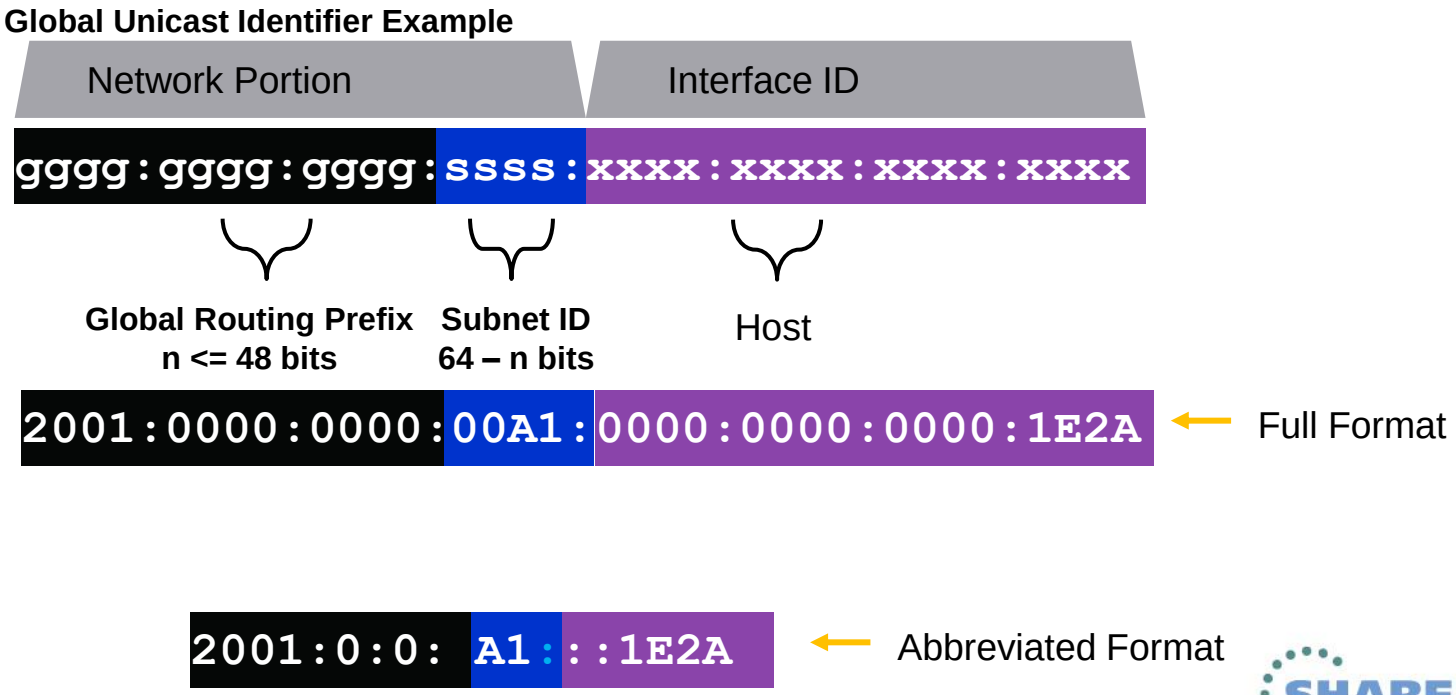

Comparison of Address Type Syntax (CATS)

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

IPv6 Address Types

- Three types of unicast address scopes

Link-Local – Non routable exists on single layer 2 domain (FE80::/64)

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

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

- Multicast addresses begin with FF00::/8

Flags (f) carried in 3rd nibble (4 bits)

Scope (s) incorporated into 4th nibble (4 bits)

Unicast prefixes provide global multicast prefix (RFC 3306)

Global Unicast Addresses

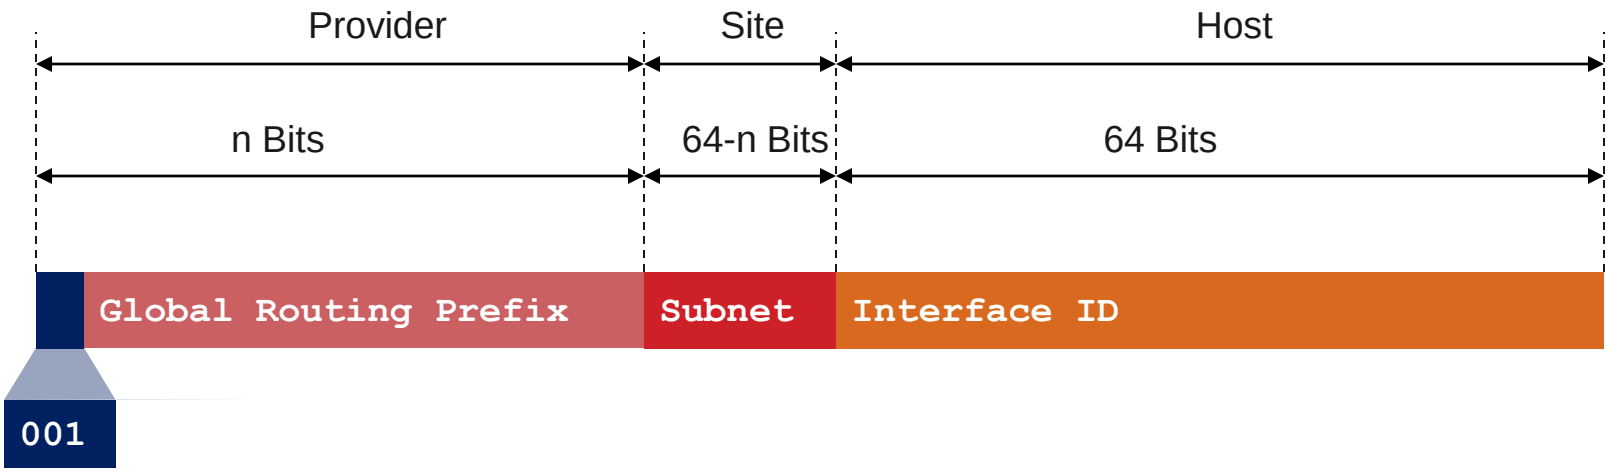

- Addresses for generic use of IPv6
- Structured as a hierarchy to try and keep the aggregation

IPv6 Addresses

Unicast Addresses

- 'Normal' IPv6 addresses
- Aggregatable with prefixes of arbitrary bit-length (similar to classless IPv4 subnetting)
- Several types:
  - global unicast
  - site-local unicast (deprecated)
  - link-local unicast
- Special purpose:
  - IPv6 mapped addresses

Anycast Addresses

- Assigned to more than one device/interface
- Typically on multiple nodes
- Packet routed to the 'nearest' interface
- Any unicast address can be an anycast address
- Explicit configuration required (interface must be told the address is of type anycast)

Multicast Addresses

- The address is an identifier for a group of interfaces
- Typically on multiple nodes
- An interface may have any number of multicast addresses (i.e. belong to any number of multicast groups)
- Hex value 'FF' at start of address identifies it as a multicast address

IPv6 Address Architecture is documented in RFC 4291

Important IPv6 Special Addresses

Unspecified Address -	**::/128**
Loopback/Local Address -	**::1/128**
Default Address -	**:: or ::/0**
Toredo Tunneling -	**2001::/32**
6to4 Tunneling -	**64:FF9B::/96**
Documentation -	**2001:DB8::/32**

End Host View

- Host Application Decides How to Connect
  - The **application** decides what address family or families to use
  - Possible strategies: Prefer IPv6, prefer first answer, “Happy Eyeballs”
- Conduct a DNS lookup
  - Can be a DNS lookup using either IPv4 or IPv6 based DNS server
  - Request an IPv4 record (A), IPv6 (AAAA) record, or both
  - The protocol of the address lookup and subsequent connection can differ
- Receive a DNS reply
  - Some providers (e.g. Google) will vary answers based on DNS source
  - There may be both A and AAAA records from which the application can choose
- Connect
  - Using some strategy, decide how to make the connection (e.g., prefer IPv6)

IPv4 and IPv6 Header Comparison

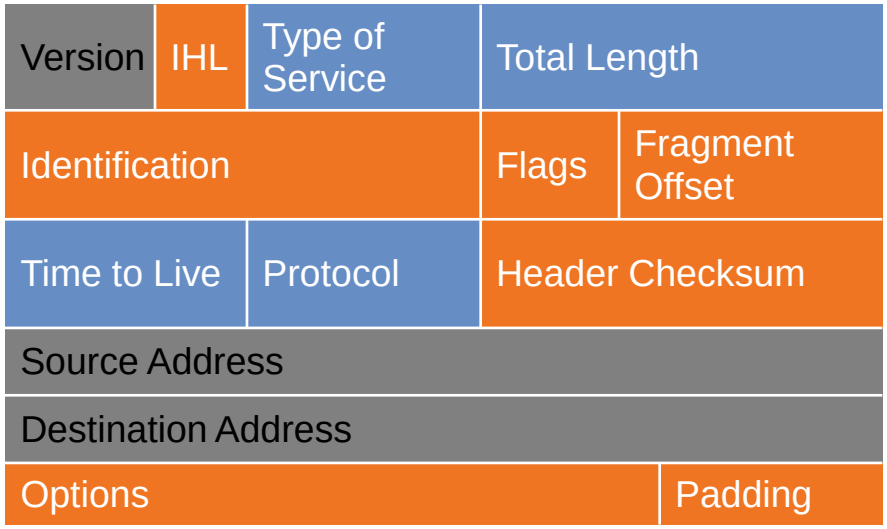

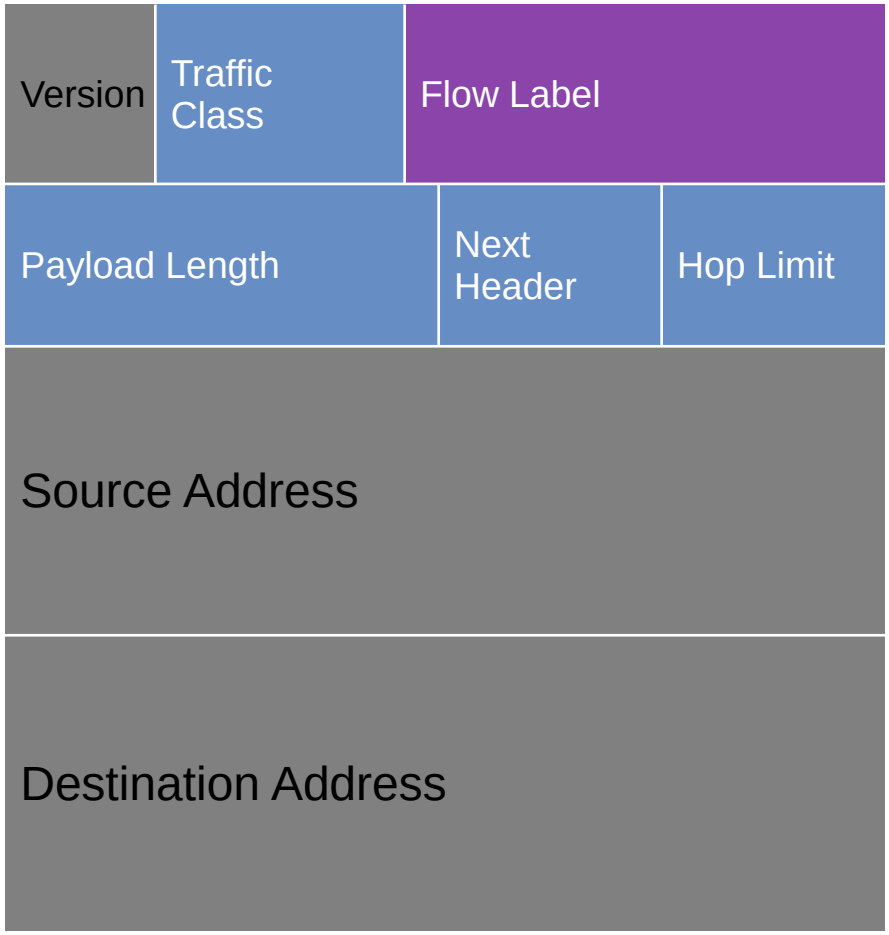

Legend

- **Field's Name Kept from IPv4 to IPv6**
- **Fields Not Kept in IPv6**
- **Name and Position Changed in IPv6**
- **New Field in IPv6**

IPv6 Working's

IPv4 and IPv6 Multicast Comparison

Service	IPv4 Solution	IPv6 Solution
Addressing Range	32-bit, Class D	128-bit (112-bit Group)
Routing	Protocol Independent, All IGPs and MBGP	Protocol Independent, All IGPs and MBGP with v6 mcast SAFI
Forwarding	PIM-DM, PIM-SM, PIM-SSM, PIM-bidir, PIM-BSR	PIM-SM, PIM-SSM, PIM-bidir, PIM-BSR
Group Management	IGMPv1, v2, v3	MLDv1, v2
Domain Control	Boundary, Border	Scope Identifier
Interdomain Solutions	MSDP across Independent PIM Domains	Single RP within Globally Shared Domains

- Static RP, BSR, No Auto-RP
- Embedded RP

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 message
  - Also send to inform change of link layer address

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

IPv6 Address Assignment

Similar to IPv4

Manually configured

Assigned via DHCP

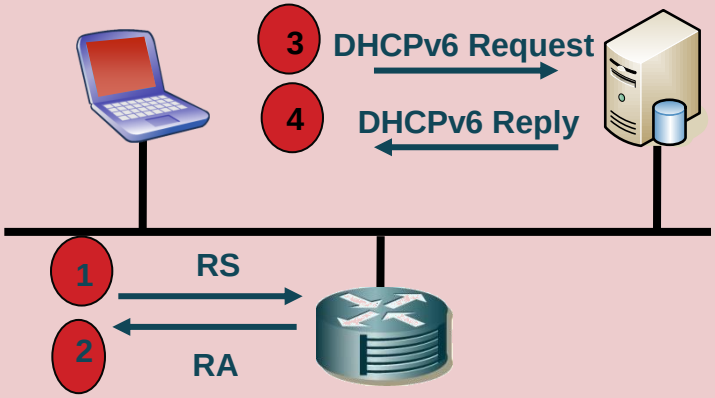

New in IPv6

Stateless configuration

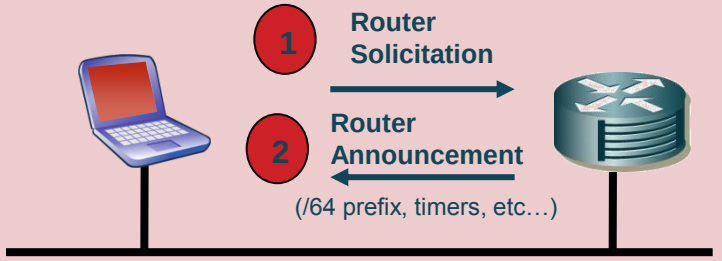

IPv6 Address = /64 prefix + EUI64 (e.g. MAC address)

Auto-generated pseudo-random number (rfc3041)

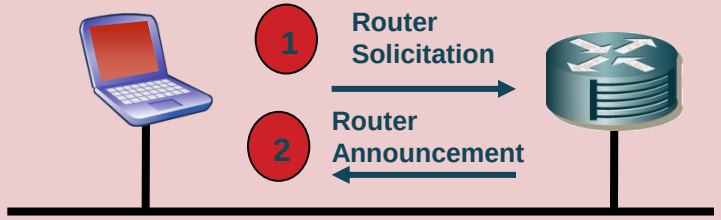

IPv6 Address = /64 prefix + Random 64 bits (rfc3041)

Network Level Considerations

- Global Unique Addresses
  - Commonly referred to as Global IPv6 addresses
  - Assigned by upstream provider
  - A multi-homed site may have one or more available IPv6 address ranges
  - The IPv6 address selection algorithm is key for good operation (RFC3484)

ULA, ULA + Global or Global

- What type of addressing should I deploy internal to my network? It depends:
  - ULA-only—Today, no IPv6 NAT is useable in production so using ULA-only will not work externally to your network
  - ULA + Global allows for the best of both worlds **but** at a price— much more address management with DHCP, DNS, routing and security
  - Global-only—Recommended approach but the old-school security folks that believe topology hiding is essential in security will bark at this option

Q and A


CISCO