The IPv6 Protocol

APNIC Staff Conference 2016

July 12-14

Mt Tamborine

Last updated 29th January 2016

So what has really changed?

- IPv6 does not interoperate with IPv4
  - Separate protocol working independently of IPv4
  - Deliberate design intention
  - Simplify IP headers to remove unused or unnecessary fields
  - Fixed length headers to “make it easier for chip designers and software engineers”

What else has changed?

- Expanded address space
  - Address length quadrupled to 16 bytes
- Header Format Simplification
  - Fixed length, optional headers are daisy-chained
  - IPv6 header is twice as long (40 bytes) as IPv4 header without options (20 bytes)
- No checksum at the IP network layer
- No hop-by-hop fragmentation
  - Path MTU discovery
- 64 bits aligned
- Authentication and Privacy Capabilities
  - IPsec is integrated
- No more broadcast

IPv4 and IPv6 Header Comparison

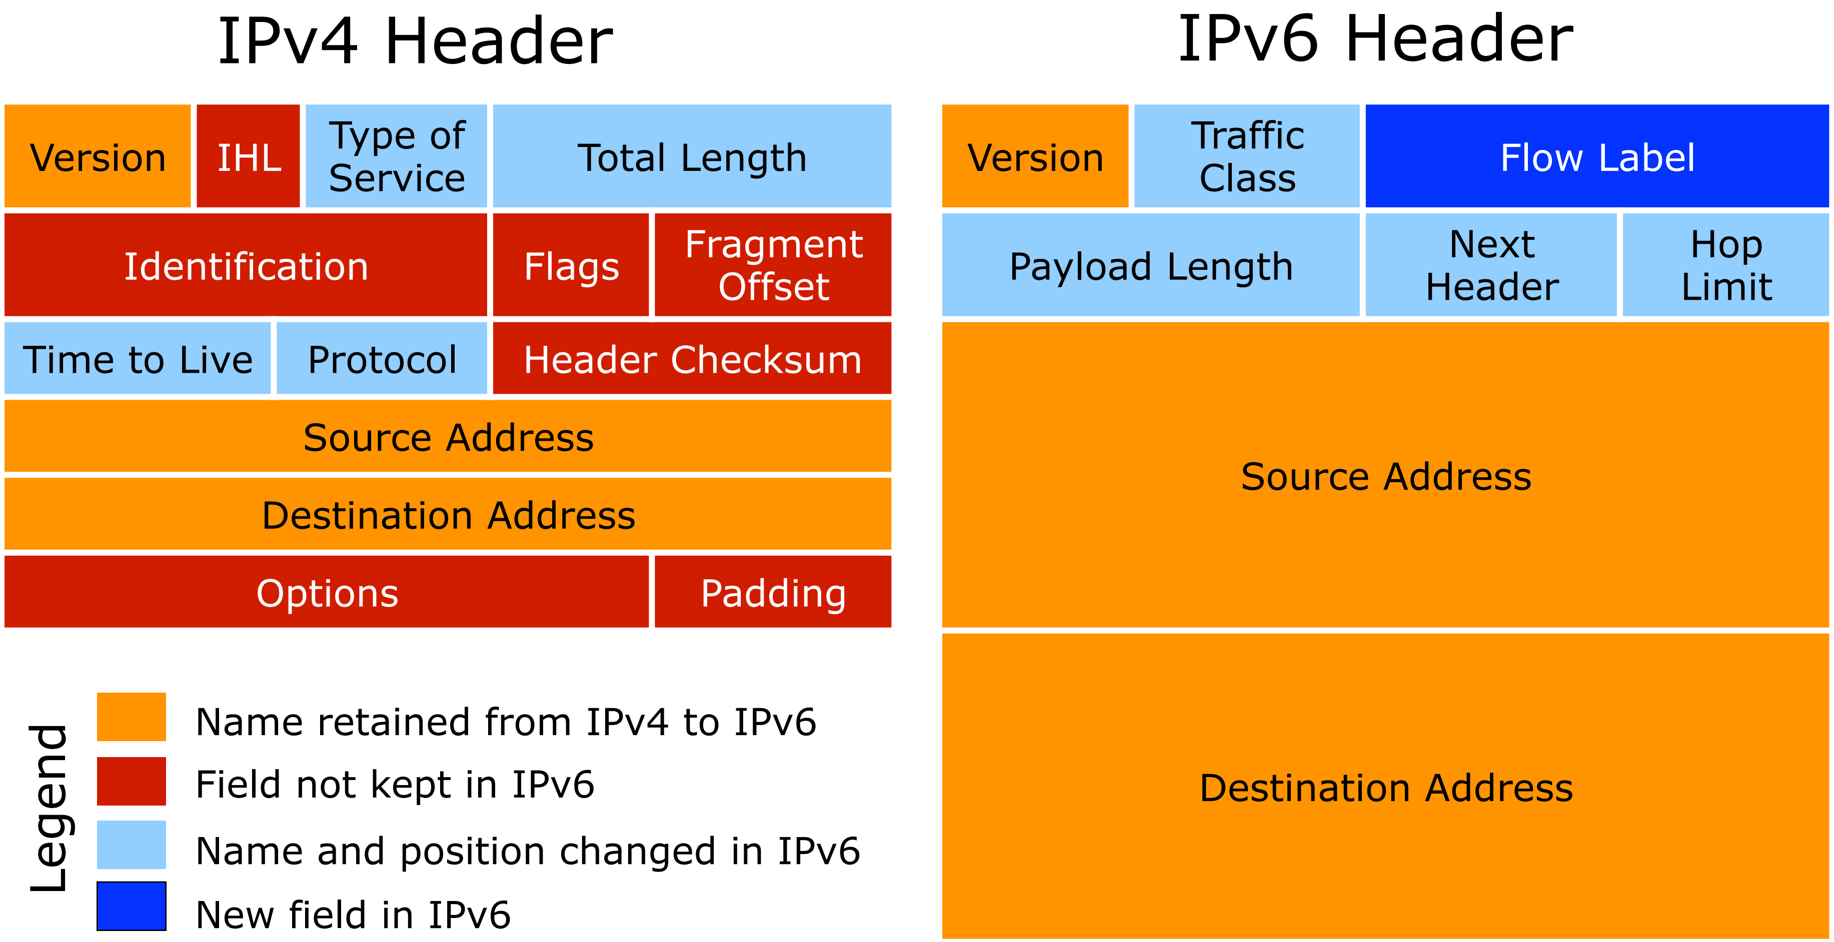

Header Format – Extension Headers

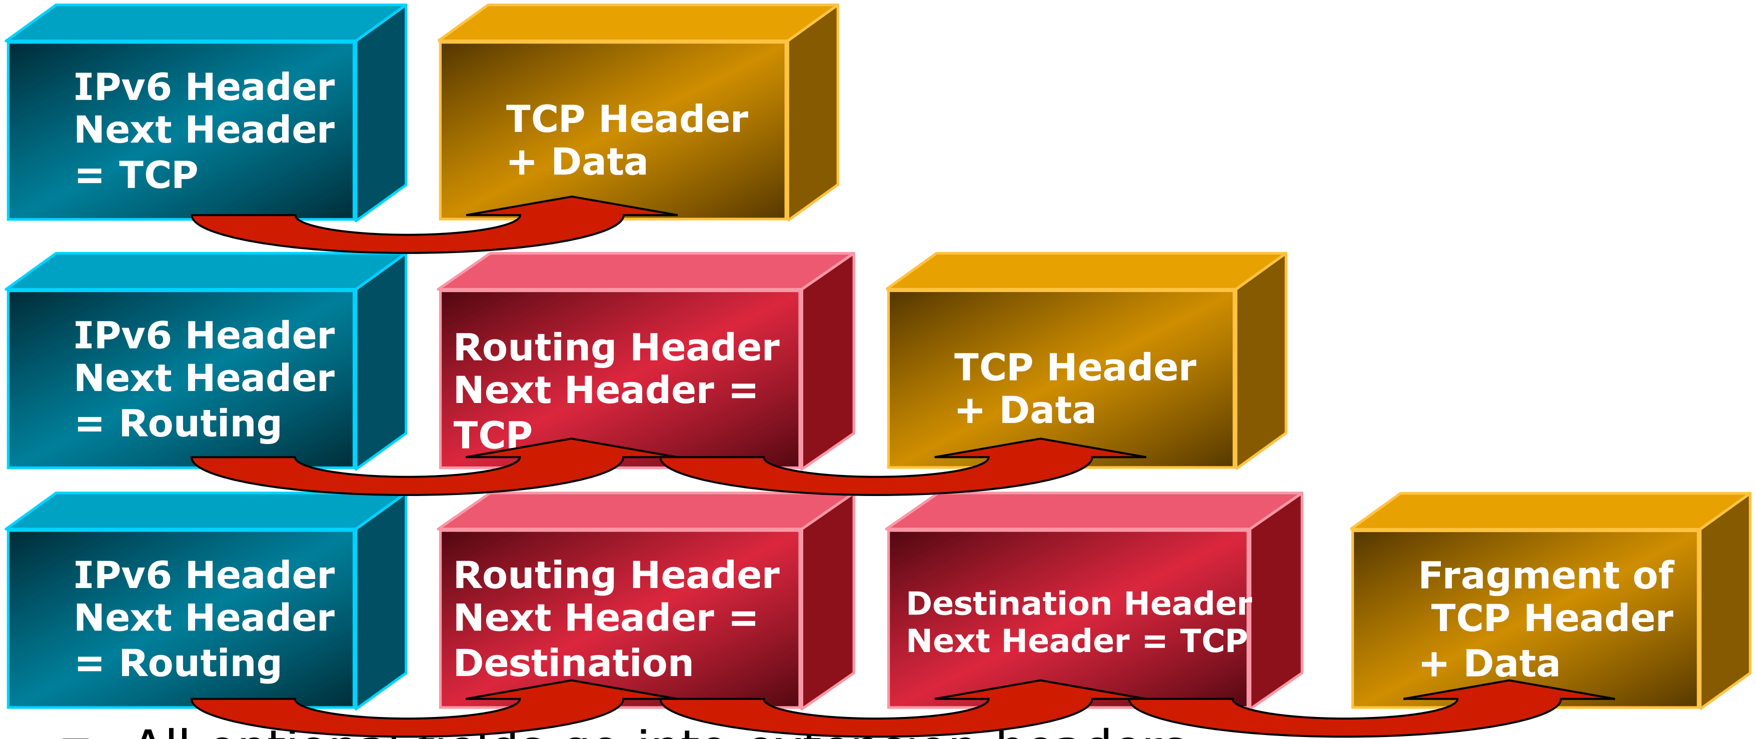

- All optional fields go into extension headers
- These are daisy chained behind the main header
  - The last 'extension' header is usually the ICMP, TCP or UDP header
- Makes it simple to add new features in IPv6 protocol without major re-engineering of devices
- Number of extension headers is not fixed / limited

Larger Address Space

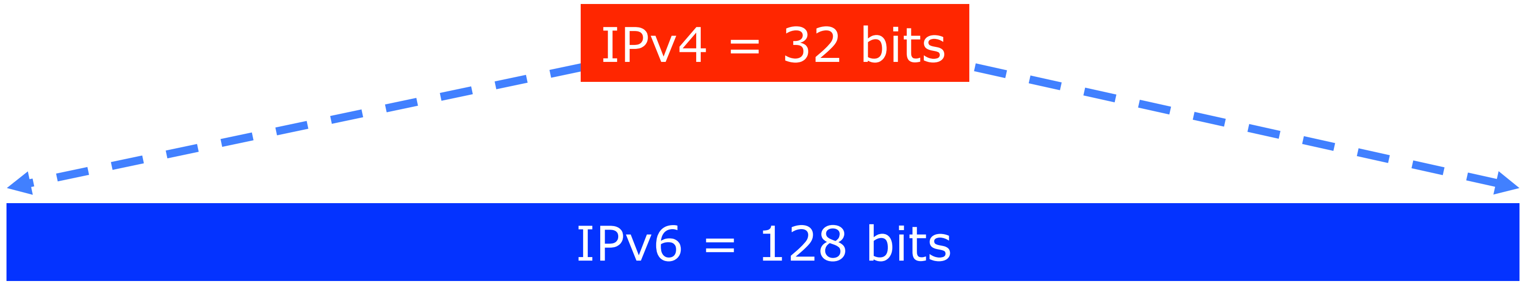

- IPv4
  - 32 bits
  - = 4,294,967,296 possible addressable devices
- IPv6
  - 128 bits: 4 times the size in bits
  - = 3.4 x 10^{38} possible addressable devices
  - = 340,282,366,920,938,463,463,374,607,431,768,211,456
  - = 4.6 x 10^{28} addresses per person on the planet

How was the IPv6 Address Size Chosen?

- Some wanted fixed-length, 64-bit addresses
  - Easily good for 10^{12} sites, 10^{15} nodes, at .0001 allocation efficiency
    - (3 orders of magnitude more than IPv6 requirement)
  - Minimizes growth of per-packet header overhead
  - Efficient for software processing
- Some wanted variable-length, up to 160 bits
  - Compatible with OSI NSAP addressing plans
  - Big enough for auto-configuration using IEEE 802 addresses
  - Could start with addresses shorter than 64 bits & grow later
- Settled on fixed-length, 128-bit addresses

IPv6 Address Representation (1)

- 16 bit fields in case insensitive colon hexadecimal representation
  - 2031:0000:130F:0000:0000:09C0:876A:130B
- Leading zeros in a field are optional:
  - 2031:0:130F:0:0:9C0:876A:130B
- Successive fields of 0 represented as ::, but only once in an address:
  - 2031:0:130F::9C0:876A:130B is ok
  - 2031::130F::9C0:876A:130B is **NOT** ok
  - 0:0:0:0:0:0:0:1 → ::1 (loopback address)
  - 0:0:0:0:0:0:0:0 → :: (unspecified address)

IPv6 Address Representation (2)

- :: representation
  - RFC5952 recommends that the rightmost set of :0: be replaced with :: for consistency
    - 2001:db8:0:2f::5 rather than 2001:db8::2f:0:0:0:5
- IPv4-compatible (not used any more)
  - 0:0:0:0:0:0:192.168.30.1
  - = ::192.168.30.1
  - = ::C0A8:1E01
- In a URL, it is enclosed in brackets (RFC3986)
  - http://[2001:db8:4f3a::206:ae14]:8080/index.html
  - Cumbersome for users, mostly for diagnostic purposes
  - Use fully qualified domain names (FQDN)
  - ⇒ The DNS has to work!!

IPv6 Address Representation (3)

- Prefix Representation
  - Representation of prefix is just like IPv4 CIDR
  - In this representation you attach the prefix length
  - Like IPv4 address:
    - 198.10.0.0/16
  - IPv6 address is represented in the same way:
    - 2001:db8:12::/40

Global Unicast Addresses

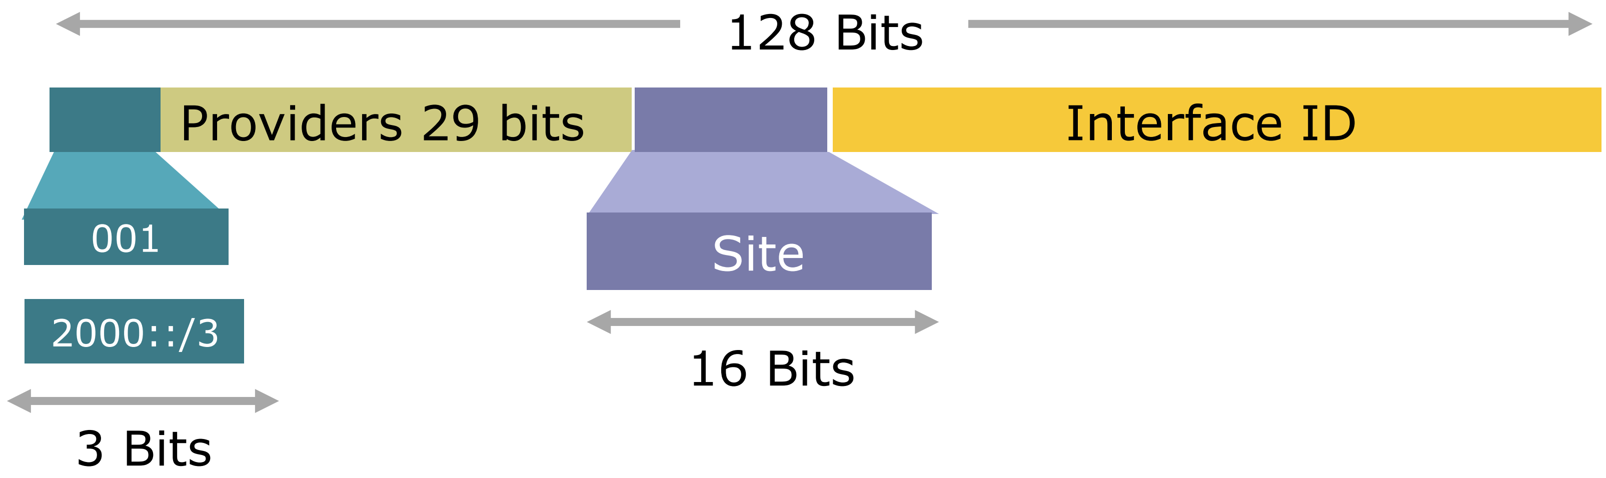

- Address block delegated by IETF to IANA
- For distribution to the RIRs and on to the users of the public Internet
- Global Unicast Address block is 2000::/3
  - This is 1/8th of the entire available IPv6 address space

Unique-Local Addresses

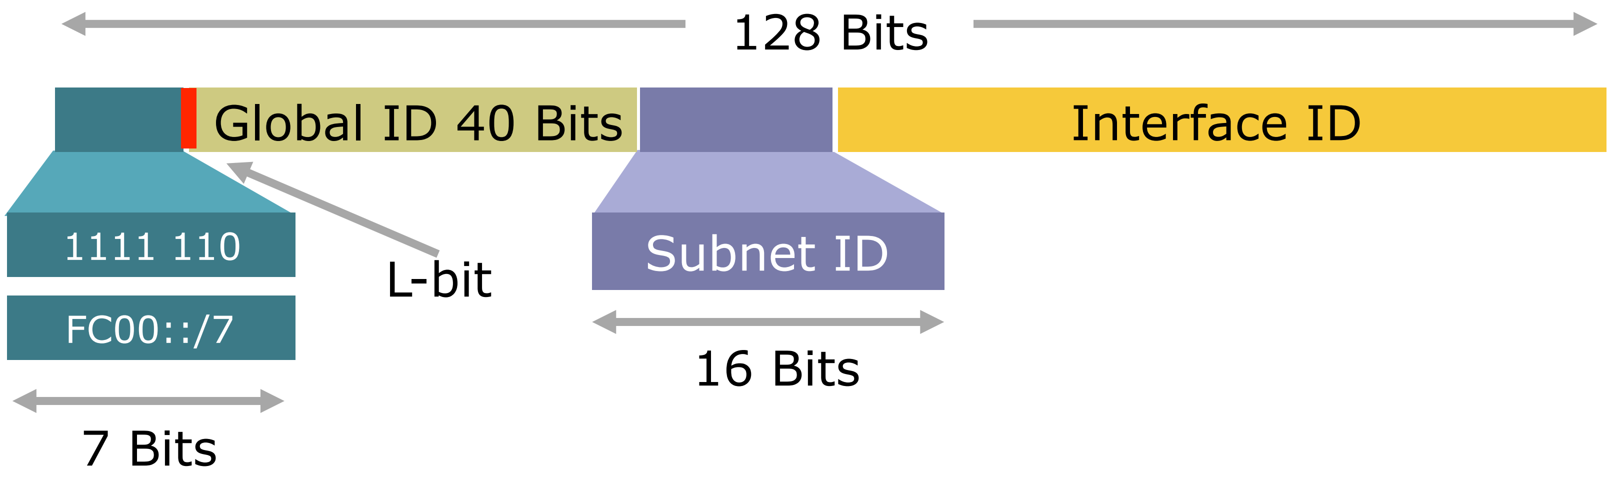

- Unique-Local Addresses (ULAs) are NOT routable on the Internet
  - L-bit set to 1 – which means the address is locally assigned
- ULAs are used for:
  - Isolated networks
  - Local communications & inter-site VPNs
  - (see now expired https://datatracker.ietf.org/doc/draft-ietf-v6ops-ula-usage-recommendations/)

Link-Local Addresses

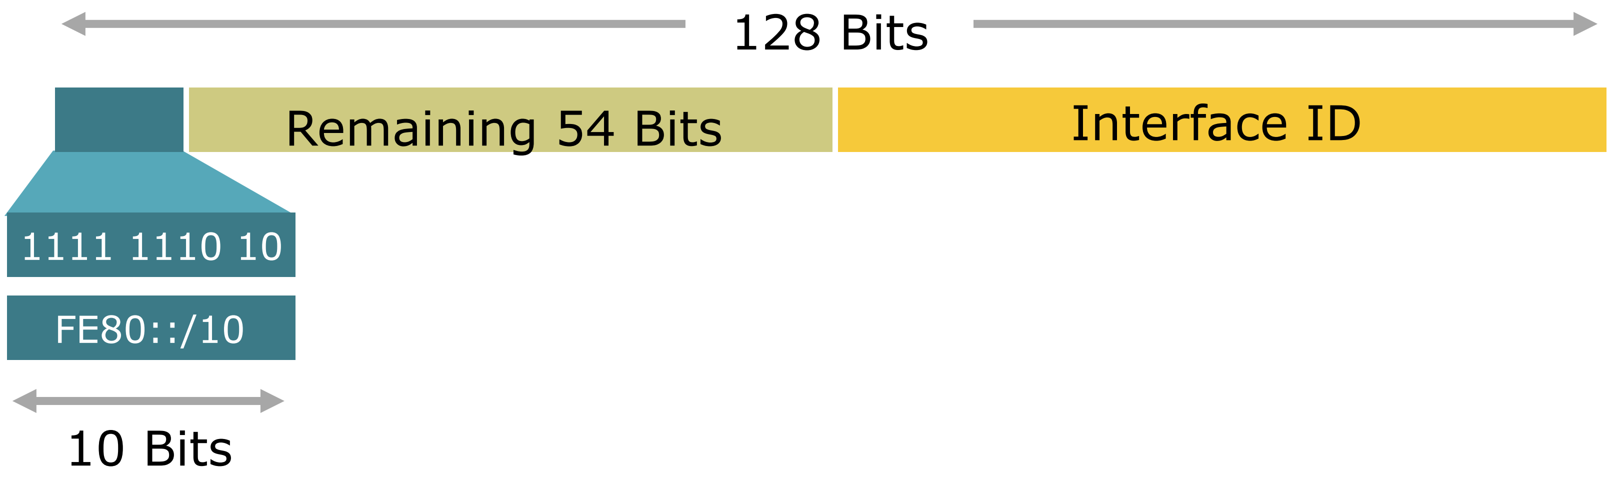

- Link-Local Addresses Used For:
  - Communication between two IPv6 device (like ARP but at Layer 3)
  - Next-Hop calculation in Routing Protocols
- Automatically assigned by Router as soon as IPv6 is enabled
  - **Mandatory Address**
- Only Link Specific scope
- Remaining 54 bits could be Zero or any manual configured value

Multicast Addresses

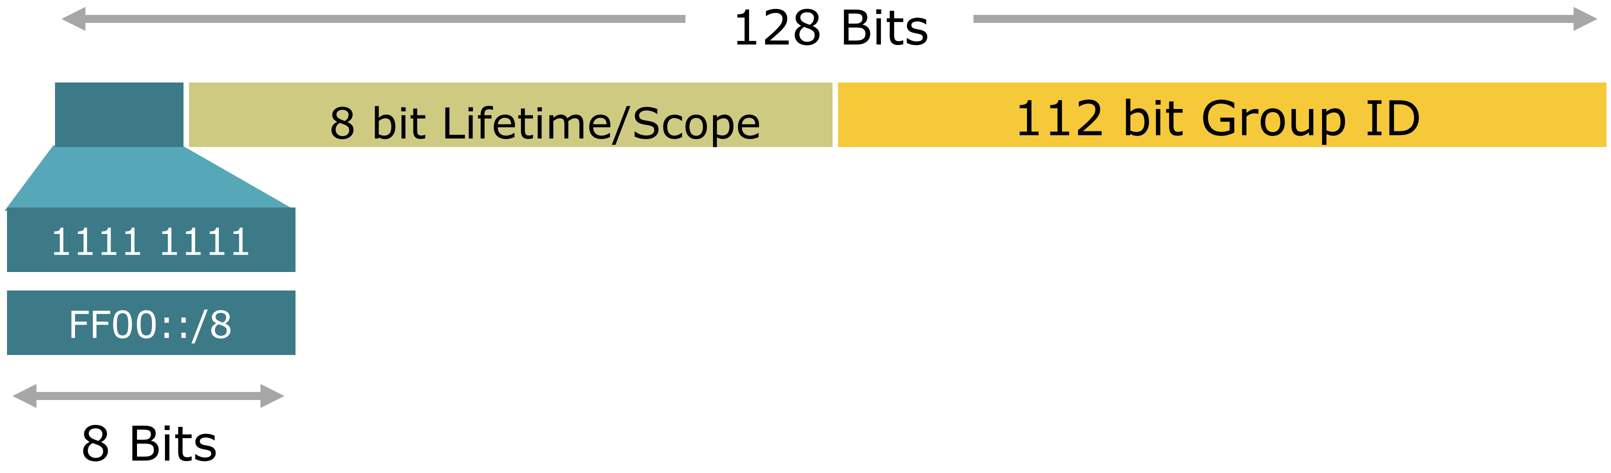

- Multicast Addresses Used For:
  - One to many communication
- 2nd octet reserved for Lifetime and Scope
- Remainder of address represents the Group ID
- (Substantially larger range than for IPv4 which only had 224.0.0.0/4 for Multicast)

Global Unicast IPv6 Address Allocation

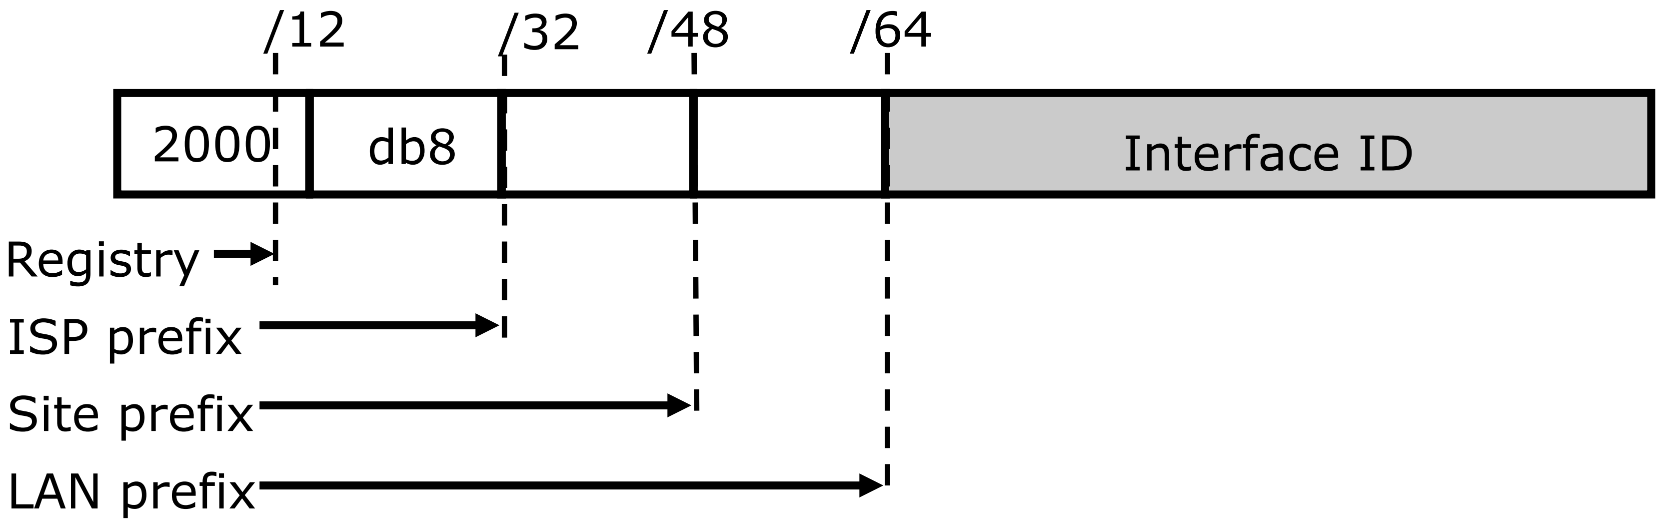

- The allocation process is:
  - The IANA is allocating out of 2000::/3 for initial IPv6 unicast use
  - Each registry gets a /12 prefix from the IANA
  - Registry allocates a /32 prefix (or larger) to an IPv6 ISP
  - Policy is that an ISP allocates a /48 prefix to each end customer

IPv6 Addressing Scope

- 64 bits reserved for the interface ID
  - Possibility of 2^{64} hosts on one network LAN
  - In theory 18,446,744,073,709,551,616 hosts
  - Arrangement to accommodate MAC addresses within the IPv6 address
- 16 bits reserved for the end site
  - Possibility of 2^{16} networks at each end-site
  - 65536 subnets equivalent to a /12 in IPv4 (assuming a /28 or 16 hosts per IPv4 subnet)

IPv6 Addressing Scope

- 16 bits reserved for each service provider
  - Possibility of 2^{16} end-sites per service provider
  - 65536 possible customers: equivalent to each service provider receiving a /8 in IPv4 (assuming a /24 address block per customer)
- 29 bits reserved for all service providers
  - Possibility of 2^{29} service providers
  - i.e. 536,870,912 discrete service provider networks
    - Although some service providers already are justifying more than a /32

Interface IDs

- Lowest order 64-bit field of unicast address may be assigned in several different ways:
  - Auto-configured from a 64-bit EUI-64, or expanded from a 48-bit MAC address (e.g., Ethernet address)
  - Auto-generated pseudo-random number (to address privacy concerns)
  - Assigned via DHCP
  - Manually configured

EUI-64

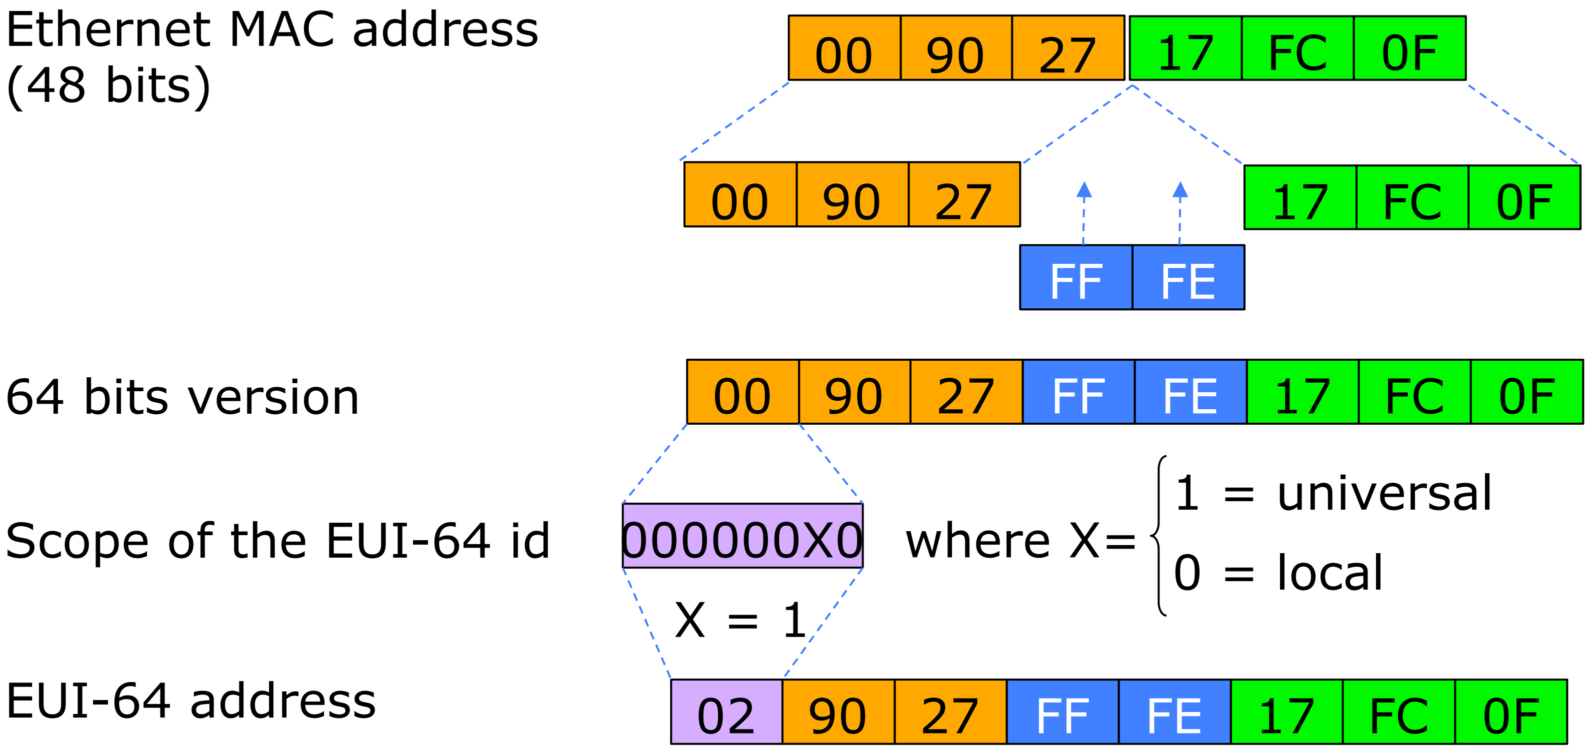

- EUI-64 address is formed by inserting FFFE between the company-id and the manufacturer extension, and setting the “u” bit to indicate scope
  - Global scope: for IEEE 48-bit MAC
  - Local scope: when no IEEE 48-bit MAC is available (eg serials, tunnels)

EUI-64

- Device MAC address is used to create:
  - Final 64 bits of global unicast address e.g.
    - 2001:db8:0:1:290:27ff:fe17:fc0f
  - Final 64 bits of link local address e.g.
    - fe80::290:27ff:fe17:fc0f
  - Final 24 bits of solicited node multicast address e.g.
    - ff02::1:ff17:fc0f
- Note that both global unicast and link local addresses can also be configured manually

IPv6 Addressing Examples

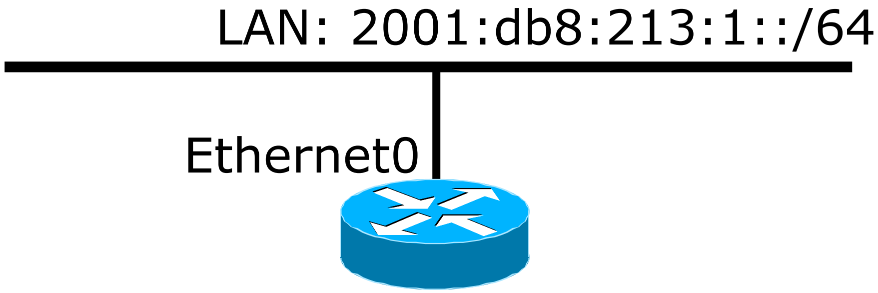


```
interface Ethernet0
 ipv6 address 2001:db8:213:1::/64 eui-64
```

MAC address: 0060.3e47.1530

```
router# show ipv6 interface Ethernet0
Ethernet0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::260:3EFF:FE47:1530
Global unicast address(es):
    2001:db8:213:1:260:3EFF:FE47:1530, subnet is 2001:db8:213:1::/64
  Joined group address(es):
    FF02::1:FF47:1530
    FF02::1
    FF02::2
  MTU is 1500 bytes
```

IPv6 Address Privacy (RFC 4941)

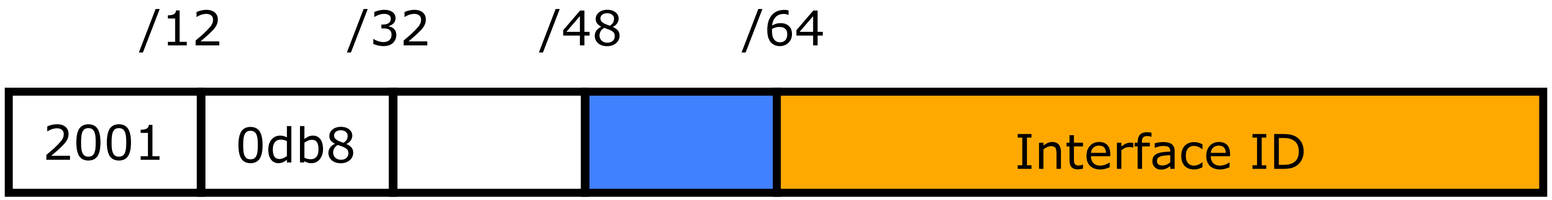

- Temporary addresses for IPv6 host client application, e.g. Web browser
- Intended to inhibit device/user tracking but is also a potential issue
  - More difficult to scan all IP addresses on a subnet
  - But port scan is identical when an address is known
- Random 64 bit interface ID, run DAD before using it
- Rate of change based on local policy
- Implemented on Microsoft Windows XP/Vista/7 and Apple MacOS 10.7 onwards
  - Can be activated on FreeBSD/Linux with a system call

Host IPv6 Addressing Options

- Stateless (RFC4862)
  - SLAAC – Stateless Address AutoConfiguration
  - Booting node sends a “router solicitation” to request “router advertisement” to get information to configure its interface
  - Booting node configures its own Link-Local address
- Stateful
  - DHCPv6 – required by most enterprises
  - Manual – like IPv4 pre-DHCP
    - Useful for servers and router infrastructure
    - Doesn’t scale for typical end user devices

Auto-configuration

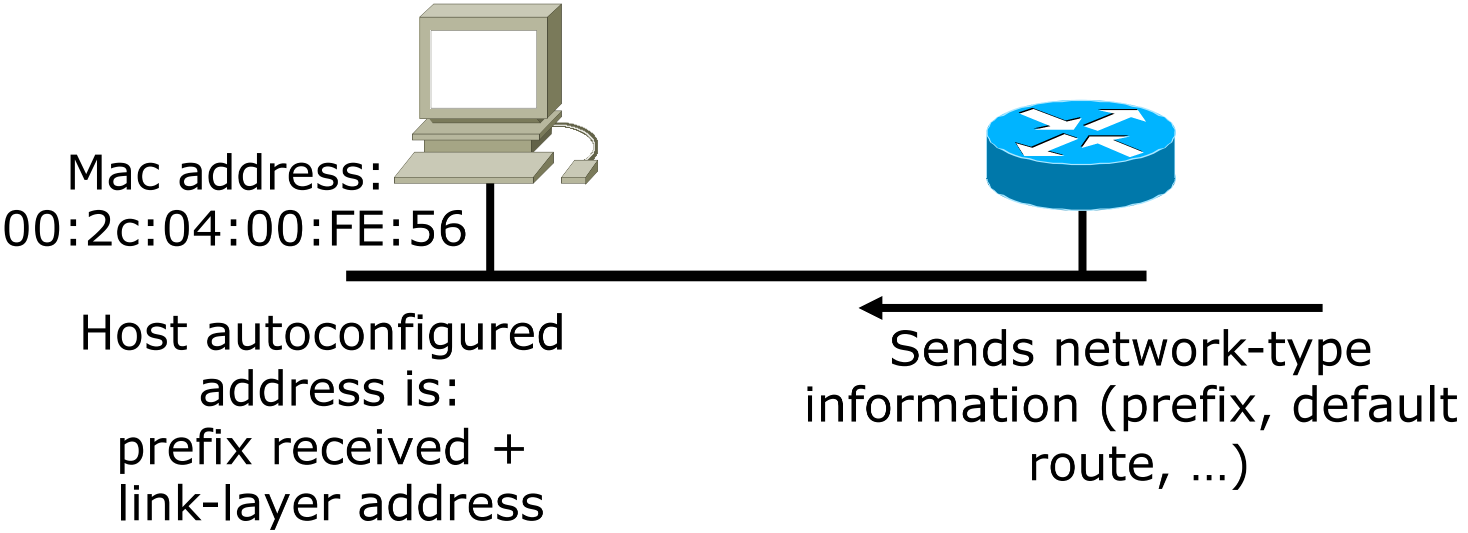

- PC sends router solicitation (RS) message
- Router responds with router advertisement (RA)
  - This includes prefix and default route
  - RFC6106 adds DNS server option
- PC configures its IPv6 address by concatenating prefix received with its EUI-64 address

IPv6 and DNS

- Hostname to IP address:

IPv4	www.abc.test.	A	192.168.30.1
IPv6	www.abc.test.	AAAA	2001:db8:c18:1::2

IPv6 and DNS

- IP address to Hostname:

IPv4	1.30.168.192.in-addr.arpa.	PTR	www.abc.test.
IPv6	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.abc.test.

IPv6 Technology Scope

IP Service	IPv4 Solution	IPv6 Solution
Addressing Range	32-bit, Network Address Translation	128-bit, Multiple Scopes
Autoconfiguration	DHCP	DHCP, Serverless, Reconfiguration
Security	IPsec	IPsec works End-to-End
Mobility	Mobile IP	Mobile IP with Direct Routing
Quality of Service	Differentiated Service, Integrated Service	Differentiated Service, Integrated Service
Multicast	IGMP, PIM, Multicast BGP	MLD, PIM, Multicast BGP, Scope Identifier

What does IPv6 do for:

- Security
  - Nothing IPv4 doesn't already support – IPSec runs in both
- QoS
  - Nothing IPv4 doesn't already support – Differentiated and Integrated Services run in both
  - So far, Flow label has no real use

IPv6 Status – Standardisation

- Several key components on standards track…

Specification (RFC2460)	Neighbour Discovery (RFC4861)
ICMPv6 (RFC4443)	IPv6 Addresses (RFC4291 & 3587)
RIP (RFC2080)	BGP (RFC2545)
IGMPv6 (RFC2710)	OSPF (RFC5340)
Router Alert (RFC2711)	Jumbograms (RFC2675)
Autoconfiguration (RFC4862)	Radius (RFC3162)
DHCPv6 (RFC3315 & 4361)	Flow Label (RFC6436/7/8)
IPv6 Mobility (RFC3775)	Mobile IPv6 MIB (RFC4295)
GRE Tunnelling (RFC2473)	Unique Local IPv6 Addresses (RFC4193)
DAD for IPv6 (RFC4429)	Teredo (RFC4380)
ISIS for IPv6 (RFC5308)	VRRP (RFC5798)

- IPv6 available over:

PPP (RFC5072)	Ethernet (RFC2464)
FDDI (RFC2467)	Token Ring (RFC2470)
NBMA (RFC2491)	ATM (RFC2492)
Frame Relay (RFC2590)	ARCnet (RFC2497)
IEEE1394 (RFC3146)	FibreChannel (RFC4338)
Facebook (RFC5514)	

Recent IPv6 Hot Topics

- IPv4 depletion debate
  - IANA IPv4 pool ran out on 3rd February 2011
    - http://www.potaroo.net/tools/ipv4/
- IPv6 Transition "assistance"
  - CGN, 6rd, NAT64, IVI, DS-Lite, 6to4, A+P…
- Mobile IPv6
- Multihoming
  - SHIM6 "dead", Multihoming in IPv6 same as in IPv4
- IPv6 Security
  - Security industry & experts taking much closer look

Conclusion

- Protocol is “ready to go”
- The core components have already seen several years field experience

The IPv6 Protocol