

IPv6 Protocols & Standards



IPv6 Migration Strategies for
Telecom Service Providers

24th-27th July 2012

Bangkok

Introduction

- The **only** compelling reason for IPv6:
 - IPv6 offers vast address space (unlike IPv4)
- What has changed
 - How the IPv6 protocol compares with IPv4
- The standards status
 - Core standards
 - Transition technologies

So what has really changed?

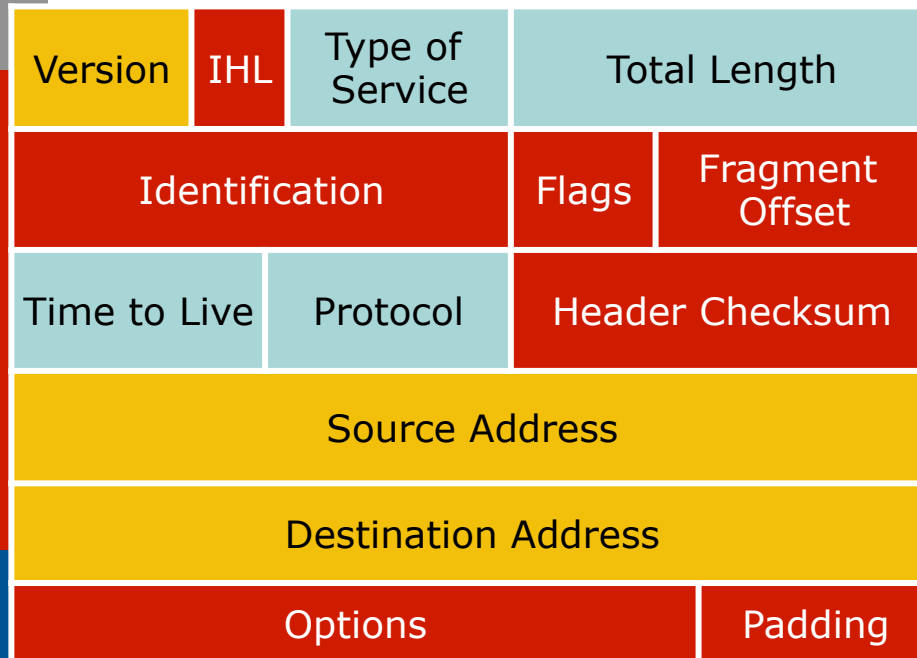
- Expanded address space
 - Address **length** quadrupled to 16 bytes
 - IPv6 address space is 2^{96} times (about 8×10^{28}) bigger than IPv4
- Header Format Simplification
 - Fixed length, optional headers are daisy-chained
 - IPv6 header is twice as long (40 bytes) as IPv4 header without options (20 bytes)

So what has really changed?

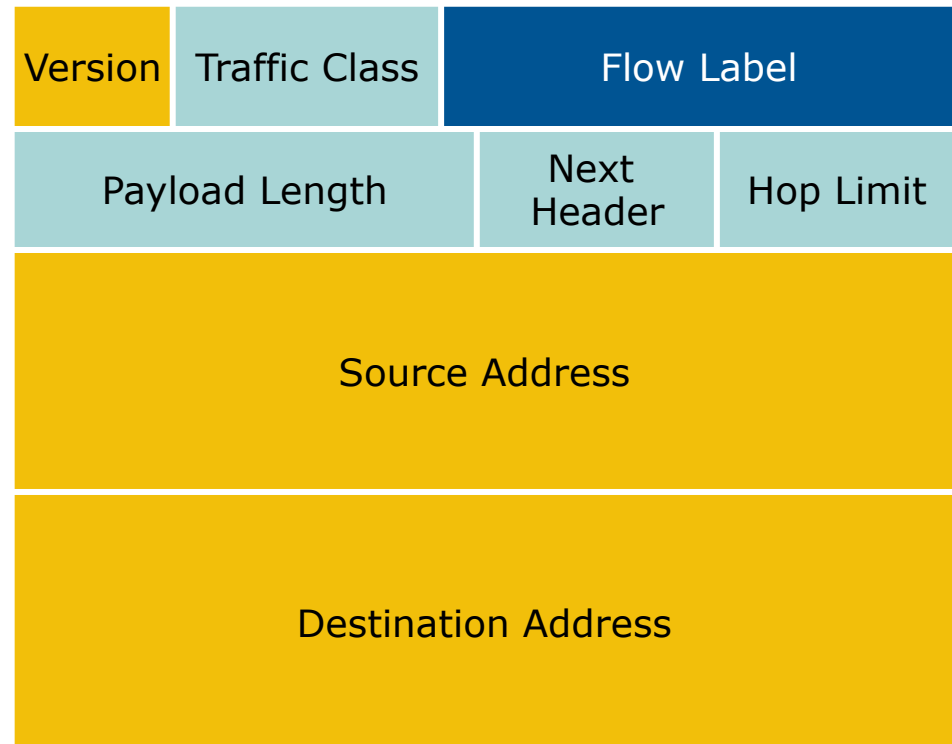
- No checksum at the IP network layer
 - Checksum already done higher up the protocol stack
- No hop-by-hop segmentation
 - Path MTU discovery
- 64 bits aligned
- Authentication and Privacy Capabilities
 - IPsec is mandated
- No more broadcast

IPv4 and IPv6 Header Comparison

IPv4 Header



IPv6 Header



Legend

- Field's name kept from IPv4 to IPv6
- Fields not kept in IPv6
- Name and position changed in IPv6
- New field in IPv6

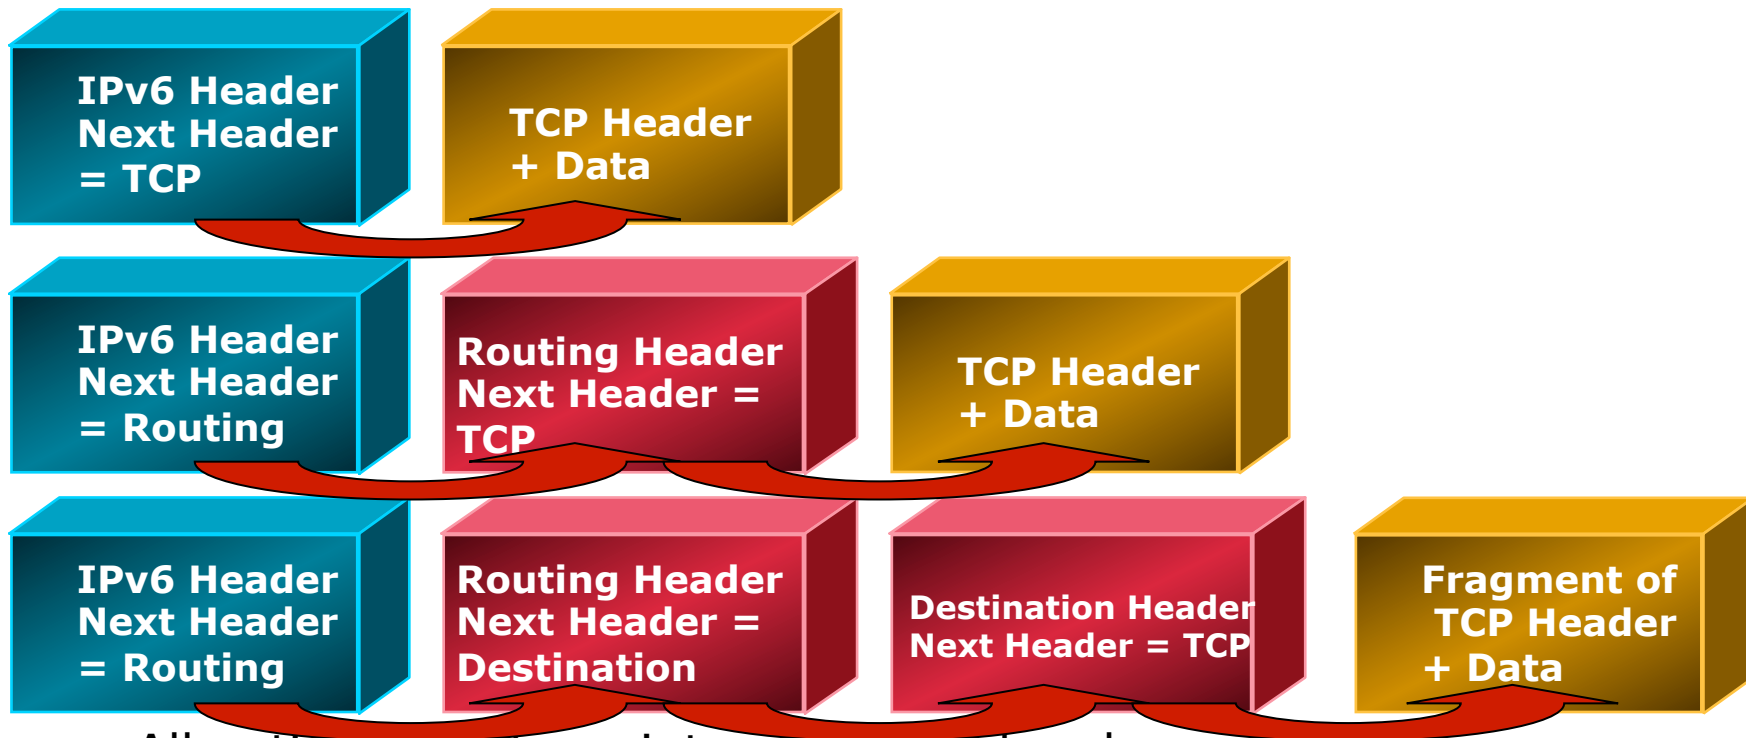
IPv6 Header

- ❑ Version = 4-bit value set to 6
- ❑ Traffic Class = 8-bit value
 - Replaces IPv4 TOS field
- ❑ Flow Label = 20-bit value
- ❑ Payload Length = 16-bit value
 - The size of the rest of the IPv6 packet following the header – replaces IPv4 Total Length
- ❑ Next Header = 8-bit value
 - Replaces IPv4 Protocol, and indicates type of next header
- ❑ Hop Limit = 8-bit value
 - Decreased by one every IPv6 hop (IPv4 TTL counter)
- ❑ Source address = 128-bit value
- ❑ Destination address = 128-bit value

Header Format Simplification

- Fixed length
 - Optional headers are daisy-chained
- 64 bits aligned
- IPv6 header is twice as long (40 bytes) as IPv4 header without options (20 bytes)
- IPv4 contains 10 basic header fields
- IPv6 contains 6 basic header fields
 - No checksum at the IP network layer
 - No hop-by-hop fragmentation

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

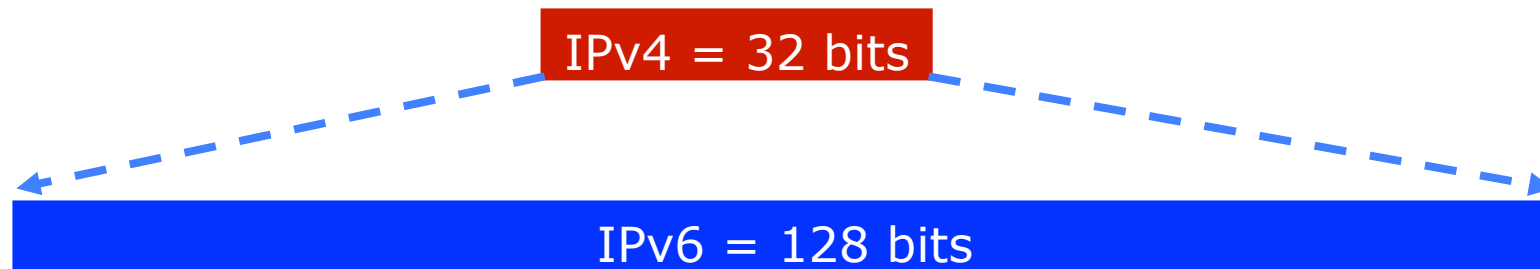
Header Format – Common Headers

- Common values of Next Header field:
 - 0 Hop-by-hop option (extension)
 - 2 ICMP (payload)
 - 6 TCP (payload)
 - 17 UDP (payload)
 - 43 Source routing (extension)
 - 44 Fragmentation (extension)
 - 50 Encrypted security payload (extension, IPSec)
 - 51 Authentication (extension, IPSec)
 - 59 Null (No next header)
 - 60 Destination option (extension)

Header Format – Ordering of Headers

- Order is important because:
 - Hop-by-hop header has to be processed by every intermediate node
 - Routing header needs to be processed by intermediate routers
 - At the destination fragmentation has to be processed before other headers
- This makes header processing easier to implement in hardware

Larger Address Space




- IPv4
 - 32 bits
 - = 4,294,967,296 possible addressable devices
- IPv6
 - 128 bits: 4 times the size in bits
 - = 3.4×10^{38} possible addressable devices
 - = 340,282,366,920,938,463,463,374,607,431,768,211,456
 - $\sim 5 \times 10^{28}$ addresses per person on the planet

How was the IPv6 Address Size Chosen?

- Some wanted fixed-length, 64-bit addresses
 - 64 bits = 1.8×10^{19} possible addresses
 - Easily good for 10^{12} sites or 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
- Agreement 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](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

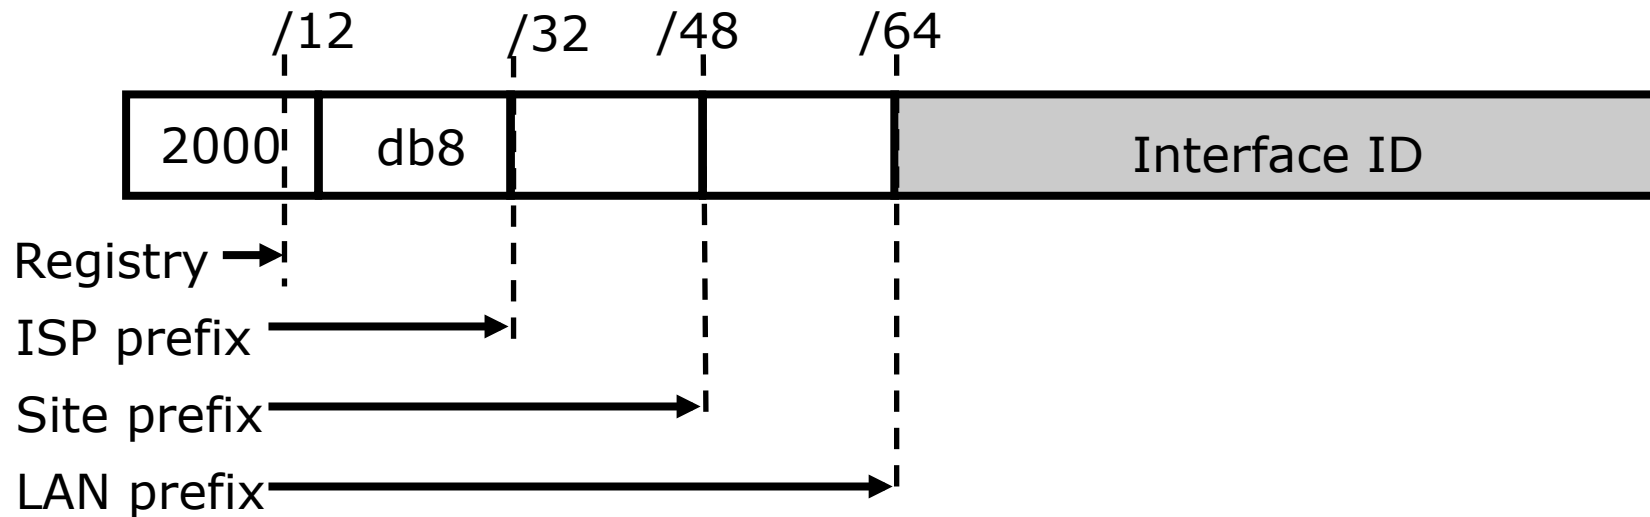
IPv6 Addressing

- ❑ IPv6 Addressing rules are covered by multiple RFCs
 - Architecture defined by RFC 4291
- ❑ Address Types are :
 - Unicast : One to One (Global, Unique Local, Link local)
 - Anycast : One to Nearest (Allocated from Unicast)
 - Multicast : One to Many
- ❑ A single interface may be assigned multiple IPv6 addresses of any type (unicast, anycast, multicast)
 - No Broadcast Address → Use Multicast

IPv6 Addressing

Type	Binary	Hex
Unspecified	000...0	::/128
Loopback	000...1	::1/128
Global Unicast Address	0010	2000::/3
Link Local Unicast Address	1111 1110 10	FE80::/10
Unique Local Unicast Address	1111 1100 1111 1101	FC00::/7
Multicast Address	1111 1111	FF00::/8

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

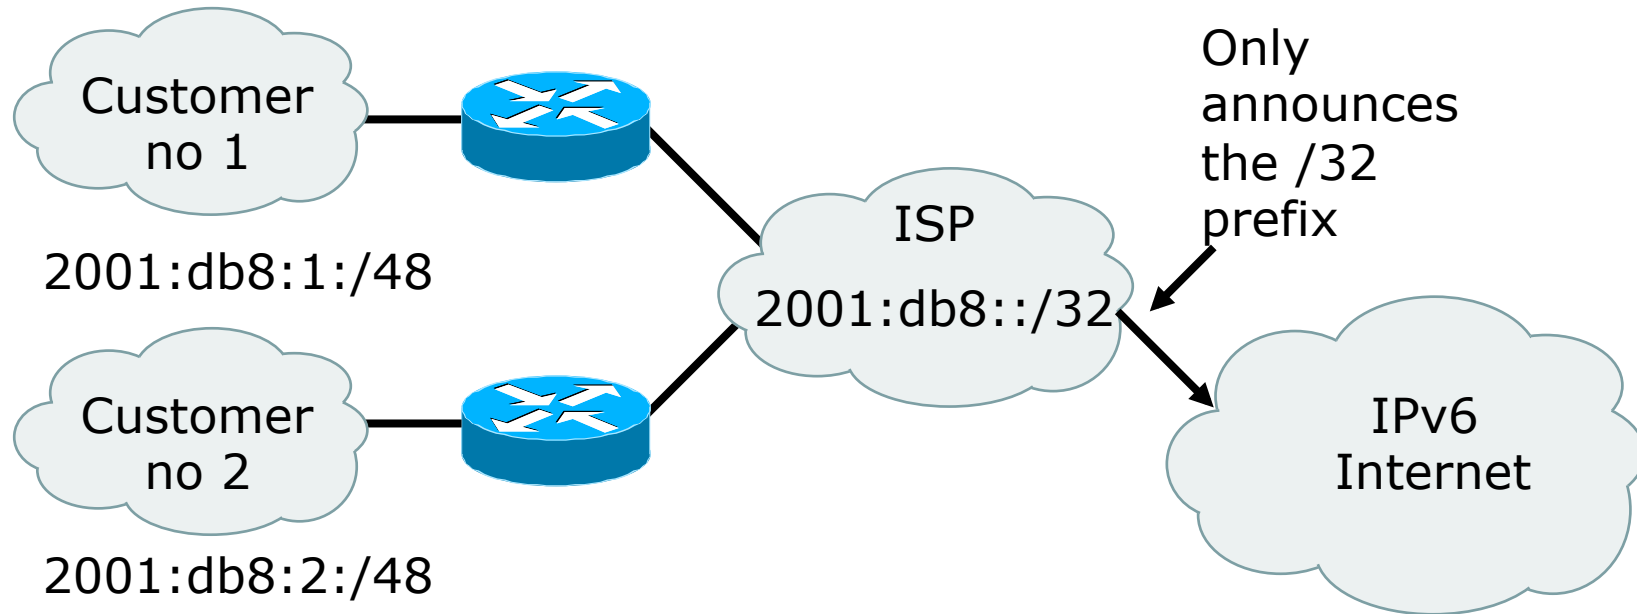
IPv6 Addressing Scope

- 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

Aggregation hopes



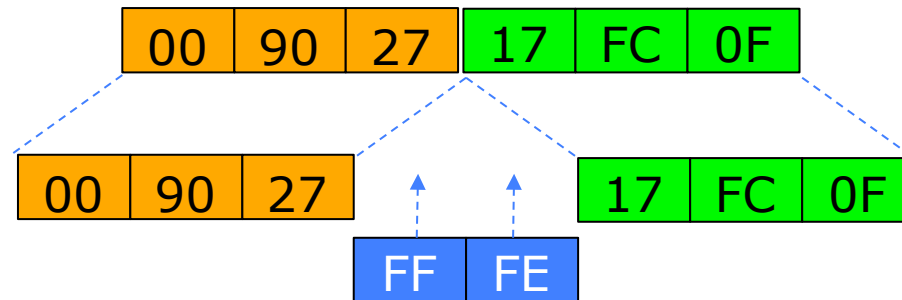
- ❑ Larger address space enables aggregation of prefixes announced in the global routing table
- ❑ Idea was to allow efficient and scalable routing
- ❑ **But current Internet multihoming solution breaks this model**

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

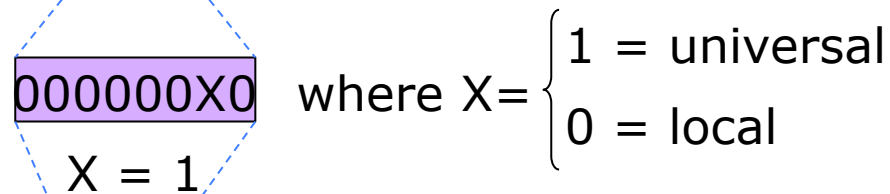
Ethernet MAC address
(48 bits)



64 bits version



Scope of the EUI-64 id



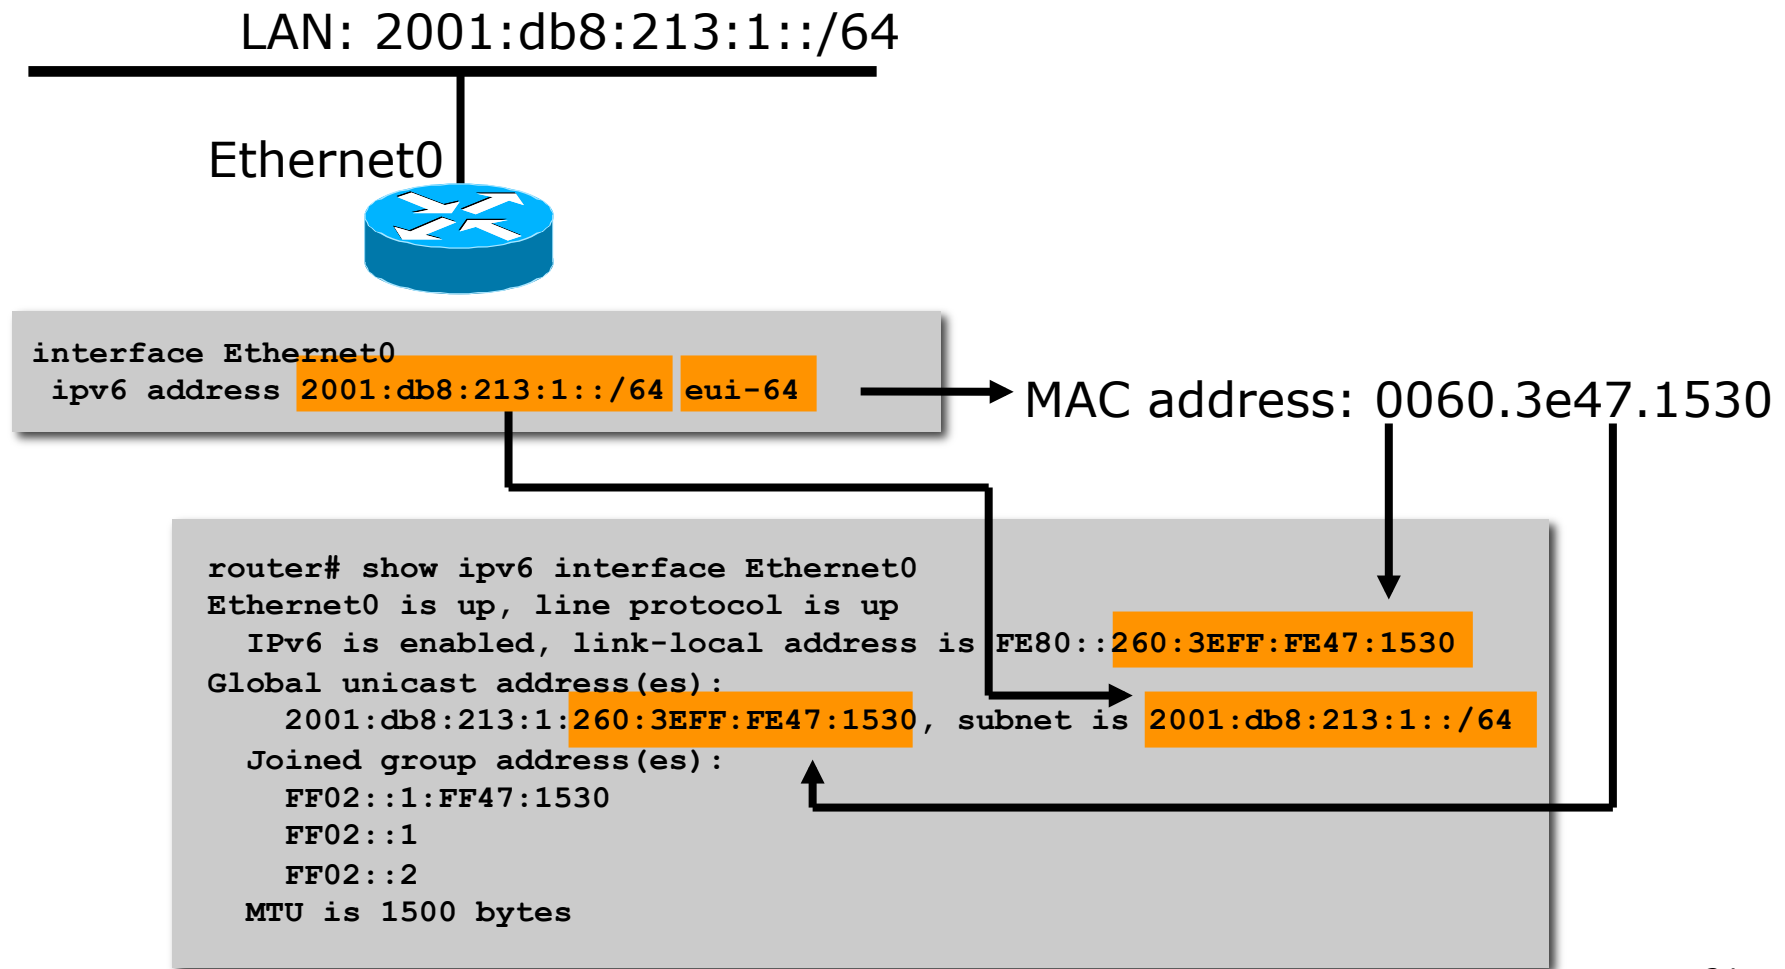
EUI-64 address



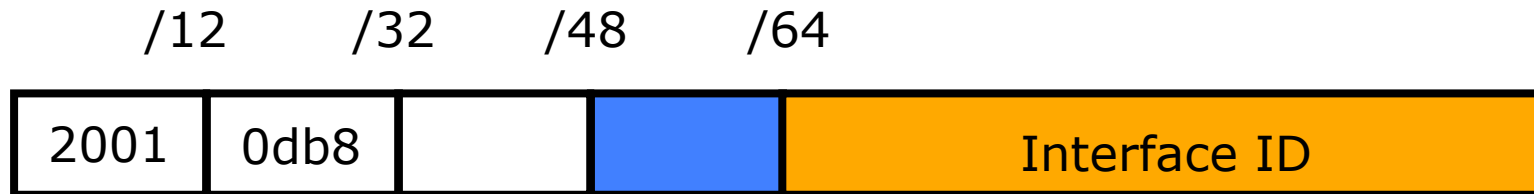
- EUI-64 address is formed by inserting FFFE between the **company-id** and the **manufacturer extension**, and setting the "u" bit to indicate scope

- Global scope: for IEEE 48-bit MAC
- Local scope: when no IEEE 48-bit MAC is available (eg serials, tunnels)

IPv6 Addressing Examples



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

IPv6 Renumbering

□ Renumbering Hosts

■ Stateless:

- Hosts renumbering is done by modifying the RA to announce the old prefix with a short lifetime and the new prefix

■ Stateful:

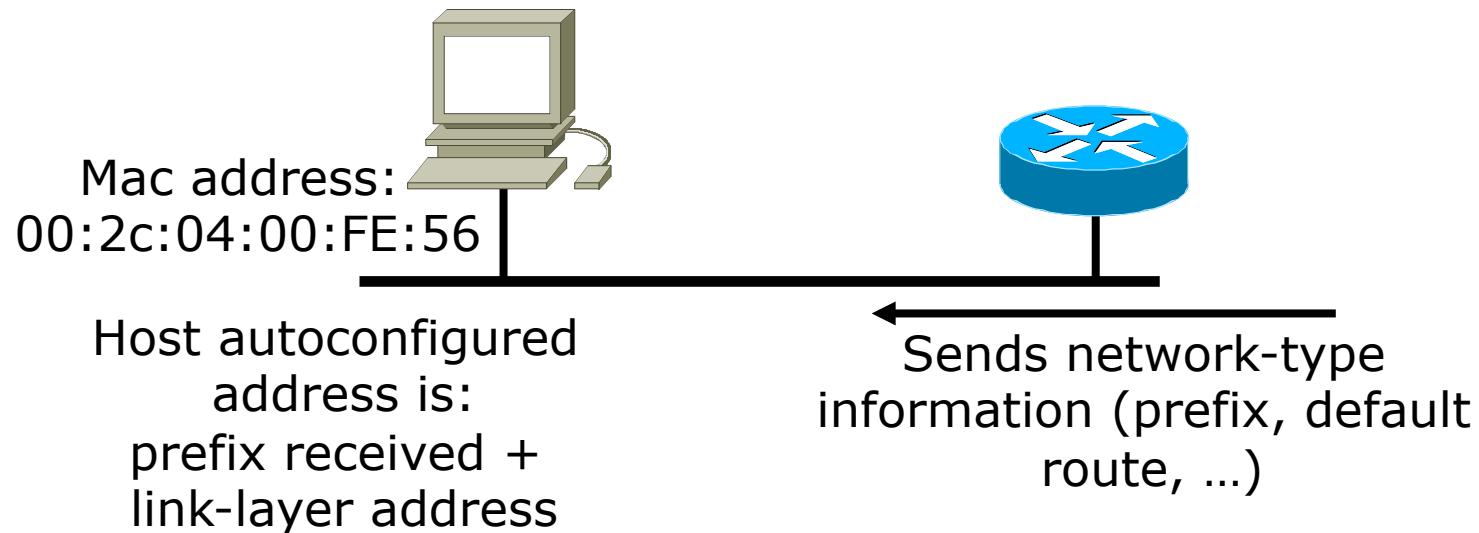
- DHCPv6 uses same process as DHCPv4

□ Renumbering Routers

- Router renumbering protocol was developed (RFC 2894) to allow domain-interior routers to learn of prefix introduction / withdrawal

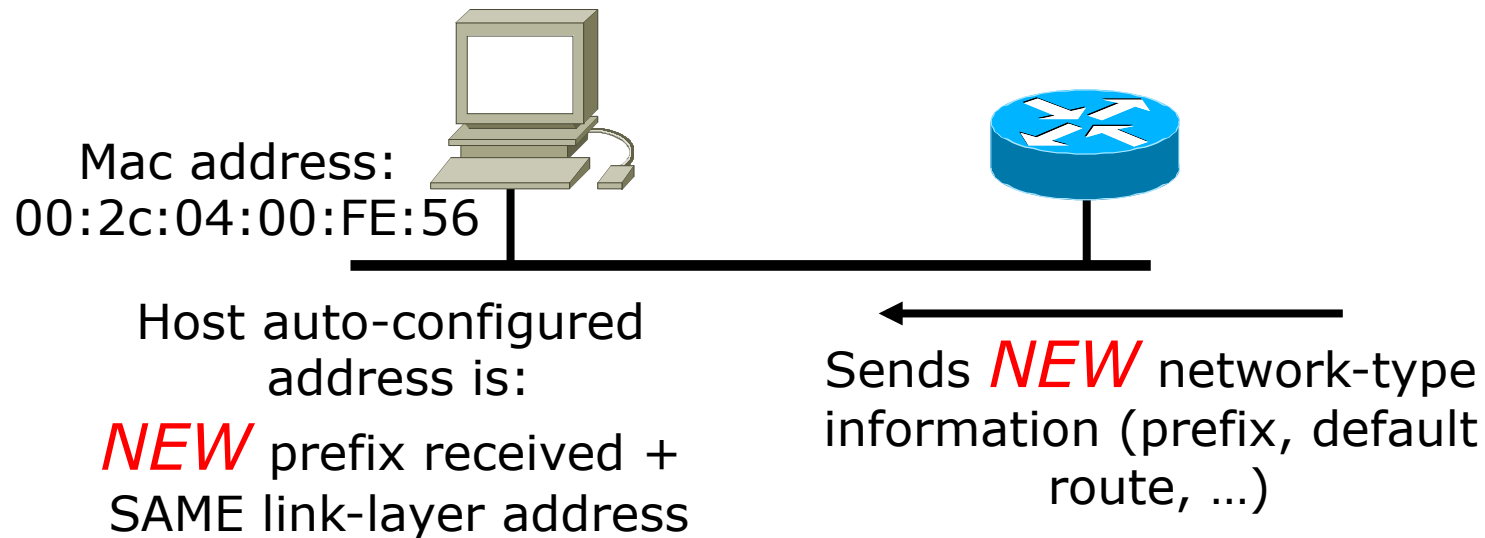
- **No known implementation!**

Auto-configuration



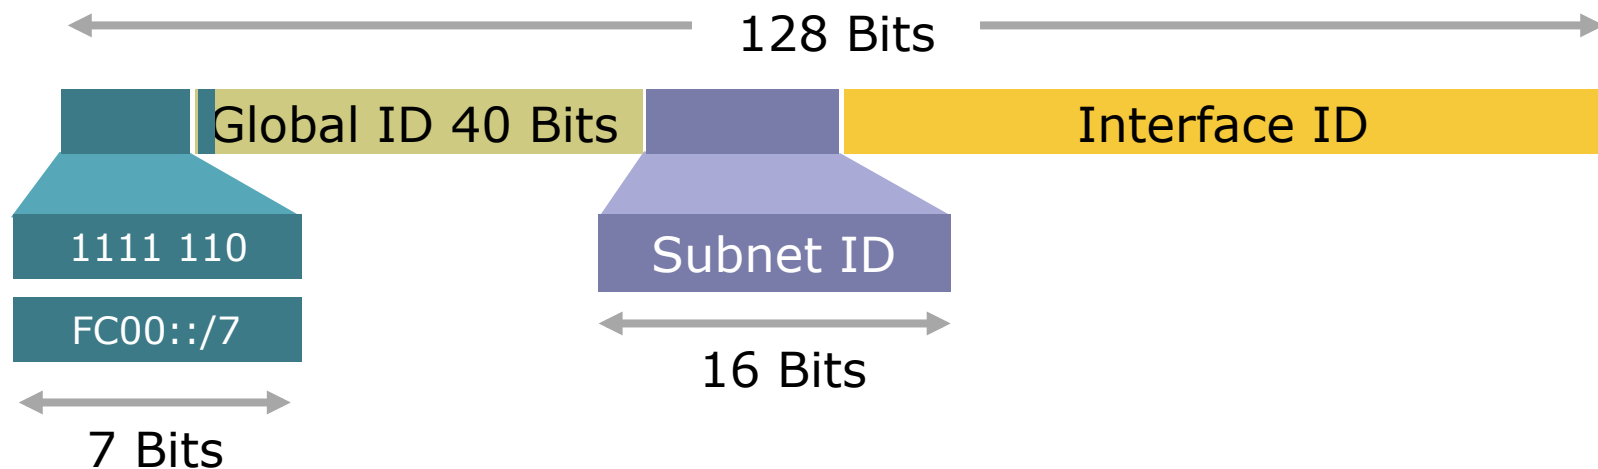
- ❑ PC sends router solicitation (RS) message
- ❑ Router responds with router advertisement (RA)
 - This includes prefix and default route
 - RFC6106 adds DNS server option
- ❑ PC configures its IPv6 address by concatenating prefix received with its EUI-64 address

Renumbering



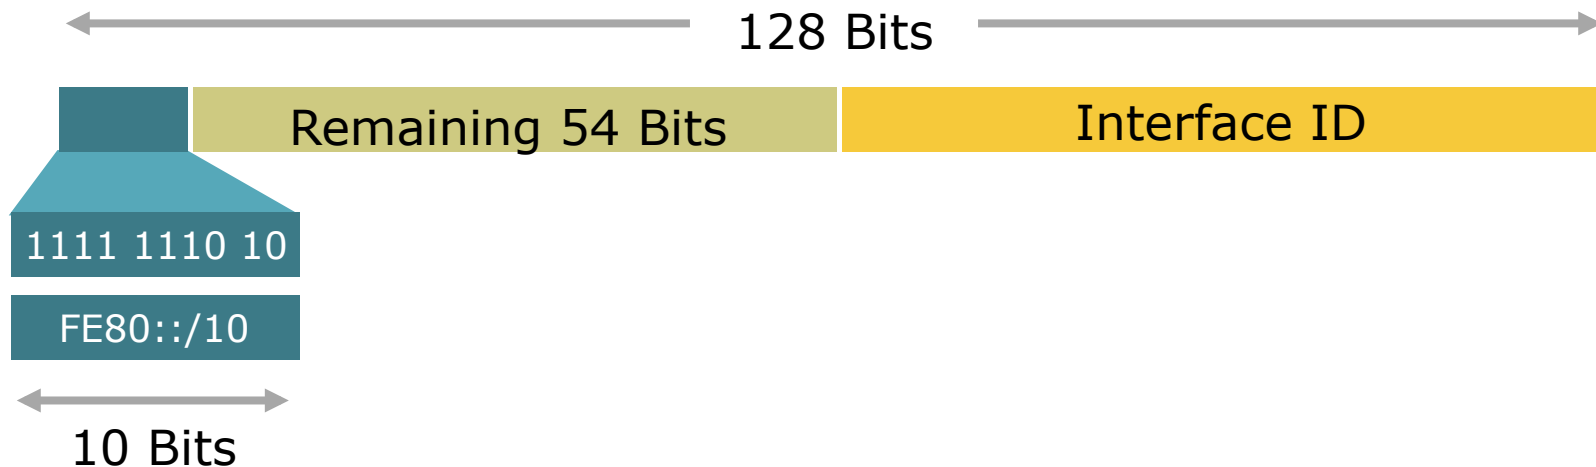
- ❑ Router sends router advertisement (RA)
 - This includes the new prefix and default route (and remaining lifetime of the old address)
- ❑ PC configures a new IPv6 address by concatenating prefix received with its EUI-64 address
 - Attaches lifetime to old address

Unique-Local



- Unique-Local Addresses Used For:
 - Local communications & inter-site VPNs
 - Local devices such as printers, telephones, etc
 - Site Network Management systems connectivity
- Not routable on the Internet
- Reinvention of the deprecated site-local?

Link-Local



- ❑ 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 configured₃₁ value

Multicast use

- Broadcasts in IPv4
 - Interrupts all devices on the LAN even if the intent of the request was for a subset
 - Can completely swamp the network (“broadcast storm”)
- Broadcasts in IPv6
 - Are not used and replaced by multicast
- Multicast
 - Enables the efficient use of the network
 - Multicast address range is much larger

IPv6 Multicast Address

- IP multicast address has a prefix FF00::/8
- The second octet defines the lifetime and scope of the multicast address.

8-bit	4-bit	4-bit	112-bit
1111 1111	Lifetime	Scope	Group-ID

Lifetime	
0	If Permanent
1	If Temporary

Scope	
1	Node
2	Link
5	Site
8	Organisation
E	Global

IPv6 Multicast Address Examples

□ RIPng

- The multicast address AllRIPRouters is **FF02::9**
 - Note that 02 means that this is a permanent address and has link scope

□ OSPFv3

- The multicast address AllSPFRouters is **FF02::5**
- The multicast address AllDRouters is **FF02::6**

□ EIGRP

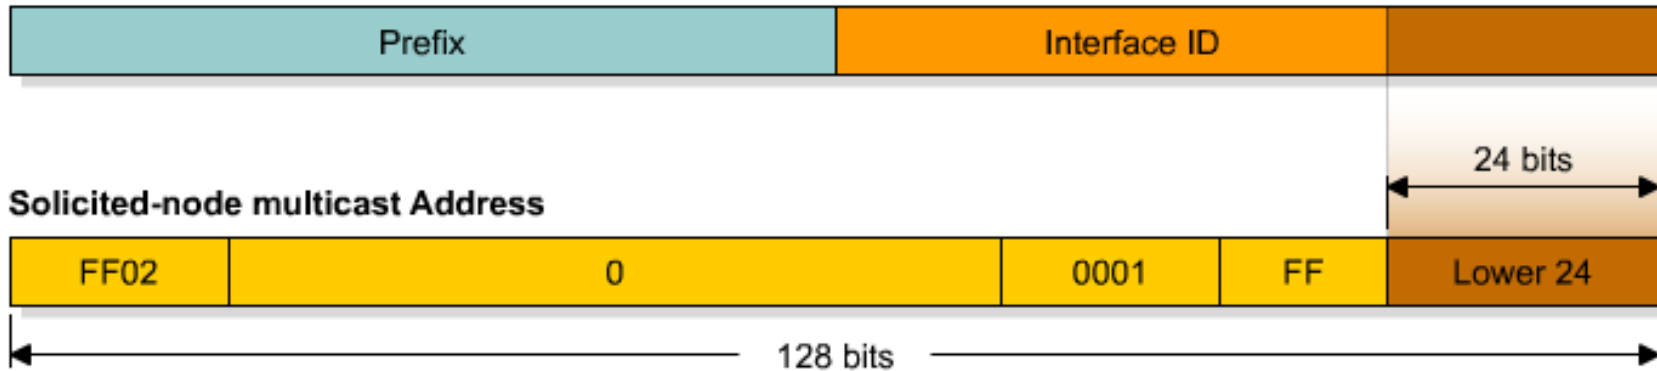
- The multicast address AllEIGRPRouters is **FF02::A**

Solicited-Node Multicast

- Solicited-Node Multicast is used for Duplicate Address Detection
 - Part of the Neighbour Discovery process
 - Replaces ARP
 - Duplicate IPv6 Addresses are rare, but still have to be tested for
- For each unicast and anycast address configured there is a corresponding solicited-node multicast address
 - This address is only significant for the local link

Solicited-Node Multicast Address

IPv6 Address



- Solicited-node multicast address consists of FF02:0:0:0:0:1:FF::/104 prefix joined with the lower 24 bits from the unicast or anycast IPv6 address

Solicited-Node Multicast

```
R1#sh ipv6 int 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.
R1#
```

Solicited-Node Multicast Address

IPv6 Anycast

- An IPv6 anycast address is an identifier for a set of interfaces (typically belonging to different nodes)
 - A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the “nearest” one, according to the routing protocol’s measure of distance).
 - [RFC4291 describes IPv6 Anycast in more detail](#)
- In reality there is no known implementation of IPv6 Anycast as per the RFC
 - Most operators have chosen to use IPv4 style anycast instead

Anycast on the Internet

- A global unicast address is assigned to all nodes which need to respond to a service being offered
 - This address is routed as part of its parent address block
- The responding node is the one which is closest to the requesting node according to the routing protocol
 - Each anycast node looks identical to the other
- Applicable within an ASN, or globally across the Internet
- Typical (IPv4) examples today include:
 - Root DNS and ccTLD/gTLD nameservers
 - SMTP relays and DNS resolvers within ISP autonomous systems

MTU Issues

- ❑ Minimum link MTU for IPv6 is 1280 octets (versus 68 octets for IPv4)
 - ⇒ on links with MTU < 1280, link-specific fragmentation and reassembly must be used
- ❑ Implementations are expected to perform path MTU discovery to send packets bigger than 1280
- ❑ Minimal implementation can omit PMTU discovery as long as all packets kept ≤ 1280 octets
- ❑ A Hop-by-Hop Option supports transmission of “jumbograms” with up to 2^{32} octets of payload

IPv6 Neighbour Discovery

- Neighbour Discovery Protocol defines mechanisms for the following problems:
 - Router discovery
 - Prefix discovery
 - Parameter discovery
 - Address autoconfiguration
 - Address resolution
 - Next-hop determination
 - Neighbour unreachability detection
 - Duplicate address detection
 - Redirects

IPv6 Neighbour Discovery

- ❑ Defined in RFC 4861
- ❑ Protocol built on top of ICMPv6 (RFC 4443)
 - Combination of IPv4 protocols (ARP, ICMP, IGMP,...)
- ❑ Fully dynamic, interactive between Hosts & Routers
- ❑ Defines 5 ICMPv6 packet types:
 - Router Solicitation
 - Router Advertisement
 - Neighbour Solicitation
 - Neighbour Advertisement
 - Redirect

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 2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0.8.1.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

Serverless,
Reconfiguration, DHCP

Security

IPSec

IPSec Mandated,
works End-to-End

Mobility

Mobile IP

Mobile IP with Direct
Routing

Quality-of-
Service

Differentiated Service,
Integrated Service

Differentiated Service,
Integrated Service

IP Multicast

IGMP/PIM/Multicast
BGP

MLD/PIM/Multicast
BGP, Scope Identifier

What does IPv6 do for:

□ Security

- Nothing IPv4 doesn't do – IPSec runs in both
- But IPv6 mandates IPSec

□ QoS

- Nothing IPv4 doesn't do –
 - Differentiated and Integrated Services run in both
 - So far, Flow label has no real use

IPv6 Security

- ❑ IPsec standards apply to both IPv4 and IPv6
- ❑ All implementations required to support authentication and encryption headers (“IPsec”)
- ❑ Authentication separate from encryption for use in situations where encryption is prohibited or prohibitively expensive
- ❑ Key distribution protocols are not yet defined (independent of IP v4/v6)
- ❑ Support for manual key configuration required

IP Quality of Service Reminder

- Two basic approaches developed by IETF:
 - “Integrated Service” (int-serv)
 - Fine-grain (per-flow), quantitative promises (e.g., x bits per second), uses RSVP signalling
 - “Differentiated Service” (diff-serv)
 - Coarse-grain (per-class), qualitative promises (e.g., higher priority), no explicit signalling
 - Signalled diff-serv (RFC 2998)
 - Uses RSVP for signalling with course-grained qualitative aggregate markings
 - Allows for policy control without requiring per-router state overhead

IPv6 Support for Int-Serv

- 20-bit Flow Label field to identify specific flows needing special QoS
 - Each source chooses its own Flow Label values; routers use Source Addr + Flow Label to identify distinct flows
 - Flow Label value of 0 used when no special QoS requested (the common case today)
- Originally standardised as RFC 3697

IPv6 Flow Label

- Flow label has not been used since IPv6 standardised
 - Suggestions for use in recent years were incompatible with original specification (discussed in RFC6436)
- Specification updated in RFC6437
 - RFC6438 describes the use of the Flow Label for equal cost multi-path and link aggregation in Tunnels

IPv6 Support for Diff-Serv

- 8-bit Traffic Class field to identify specific classes of packets needing special QoS
 - Same as new definition of IPv4 Type-of-Service byte
 - May be initialized by source or by router enroute; may be rewritten by routers enroute
 - Traffic Class value of 0 used when no special QoS requested (the common case today)

IPv6 Status – Standardisation

- IPv6 Standardisation:
 - Specification (RFC2460)
 - ICMPv6 (RFC4443)
 - RIP (RFC2080)
 - IGMPv6 (RFC2710)
 - Router Alert (RFC2711)
 - Autoconfiguration (RFC4862)
 - DHCPv6 (RFC3315 & 4361)
 - IPv6 Mobility (RFC3775)
 - GRE Tunnelling (RFC2473)
 - DAD for IPv6 (RFC4429)
 - ISIS for IPv6 (RFC5308)
 - Neighbour Discovery (RFC4861)
 - IPv6 Addresses (RFC4291 & 3587)
 - BGP (RFC2545)
 - OSPF (RFC5340)
 - Jumbograms (RFC2675)
 - Radius (RFC3162)
 - Flow Label (RFC6436/7/8)
 - Mobile IPv6 MIB (RFC4295)
 - Unique Local IPv6 Addresses (RFC4193)
 - Teredo (RFC4380)
 - VRRP (RFC5798)
- IPv6 available over:
 - PPP (RFC5072)
 - FDDI (RFC2467)
 - NBMA (RFC2491)
 - Frame Relay (RFC2590)
 - IEEE1394 (RFC3146)
 - Facebook (RFC5514)
 - Ethernet (RFC2464)
 - Token Ring (RFC2470)
 - ATM (RFC2492)
 - ARCnet (RFC2497)
 - FibreChannel (RFC4338)

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

- IPv6 is “ready to go”
- The core components have already seen many years field experience
 - Many network operators have silently deployed IPv6 over the last decade
 - Even if not made available to end-users