IPv6 Addressing

Nodes

- Router
  - device that forwards packets not addressed to self
- Host
  - device that does not forward packets not addressed to self
- Hybrid
  - device that forwards to/from some interfaces but not others

Links

- loopback link
  - connects a node only to itself
- point-to-point link
  - can interconnect at most two nodes
- multi-access link
  - can interconnect more than two nodes

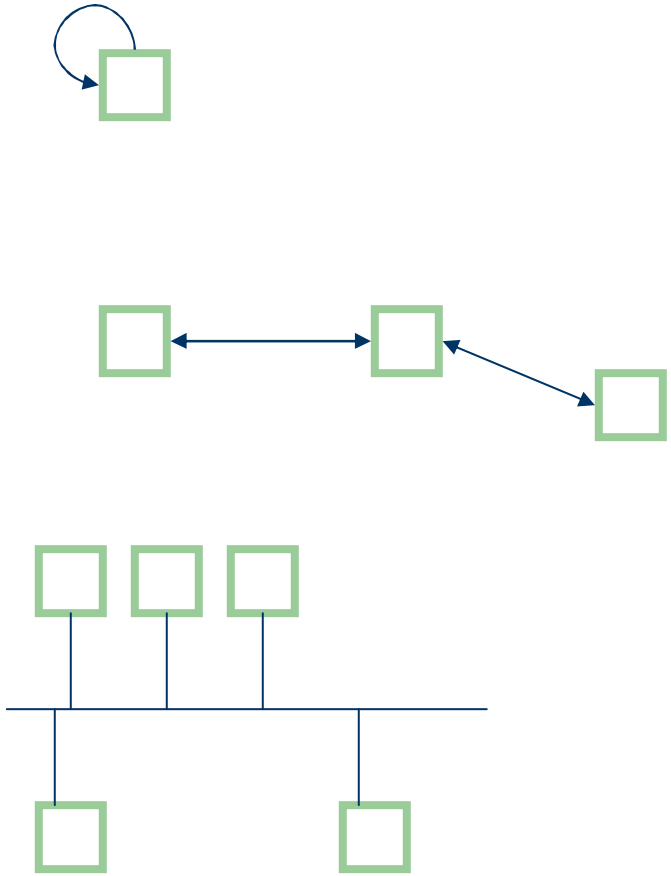

Interfaces

- an interface is a node's attachment to a link

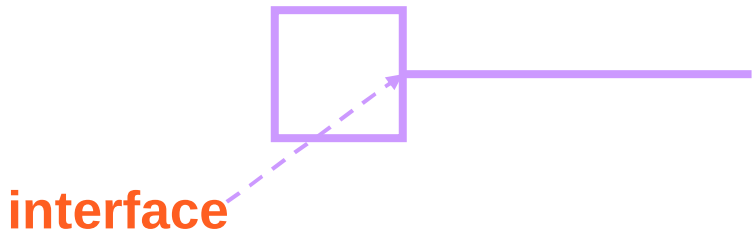

- an interface can be
  - bi-directional
  - input-only (e.g., to a receive-only satellite dish)
  - output-only (e.g., to a transmit-only satellite up-link)
  - loopback-only

Basic Address Types

- Unicast
  - Address of a single interface
  - Delivery to single interface
  - for one-to-one communication
- Multicast
  - Address of a set of interfaces
  - Delivery to all interfaces in the set
  - for one-to-many communication
- Anycast
  - Address of a set of interfaces
  - Delivery to a single interface in the set
  - for one-to-nearest communication
  - Nearest is defined as being closest in term of routing distance

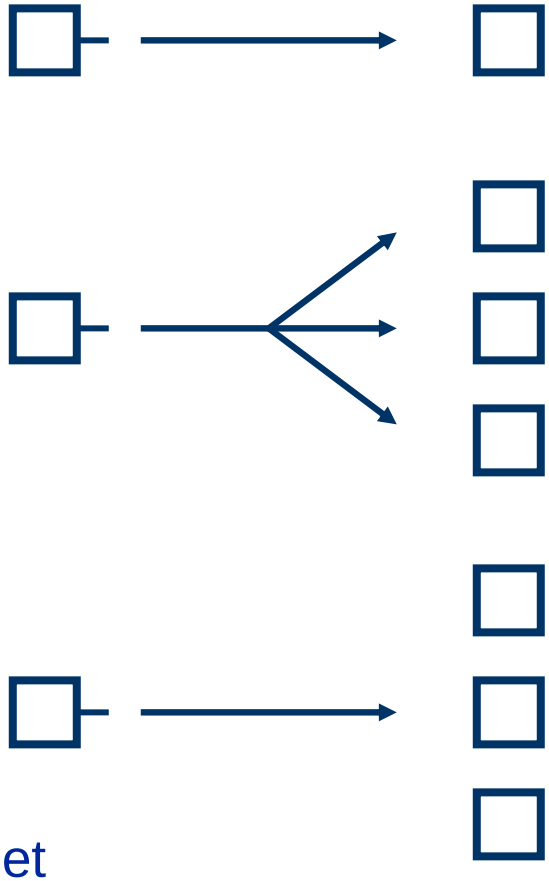

IPv6 - Addressing Model

- Addresses are assigned to interfaces, not hosts
  - No change from IPv4 Model
- Interface ‘expected’ to have multiple addresses
- Addresses have scope
  - Link Local
  - Site Local
  - Global
- Addresses have lifetime
  - Valid and Preferred lifetime

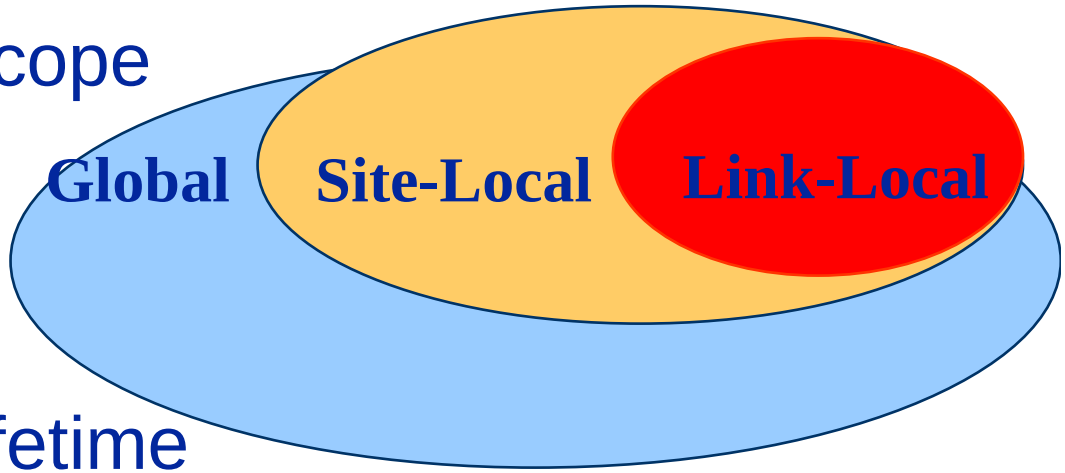

Text Representation of Addresses

- Colon-Hex
  - 3ffe:3600:2000:0800:0248:54ff:fe5c:8868
- Compressed Format:
  - 3ffe:0b00:0c18:0001:0000:0000:0000:0010
    becomes
    3ffe:b00:c18:1::10
- IPv4-compatible:
  - 0:0:0:0:0:0:140.110.60.46
  - or ::140.110.60.46
- 6to4 Address
  - 2002:8C6E:3C2E::8C6E:3C2E
    - 140.110.60.46 = 8C6E:3C2E

Address Type Prefixes

Address type	Binary prefix
IPv4-compatible	0000...0 (96 zero bits)
global unicast	001
link-local unicast	1111 1110 10
site-local unicast	1111 1110 11
multicast	1111 1111

- all other prefixes reserved (approx. 7/8ths of total)
- anycast addresses allocated from unicast prefixes

RFC 3587 - IPv6 Global Unicast Address Format

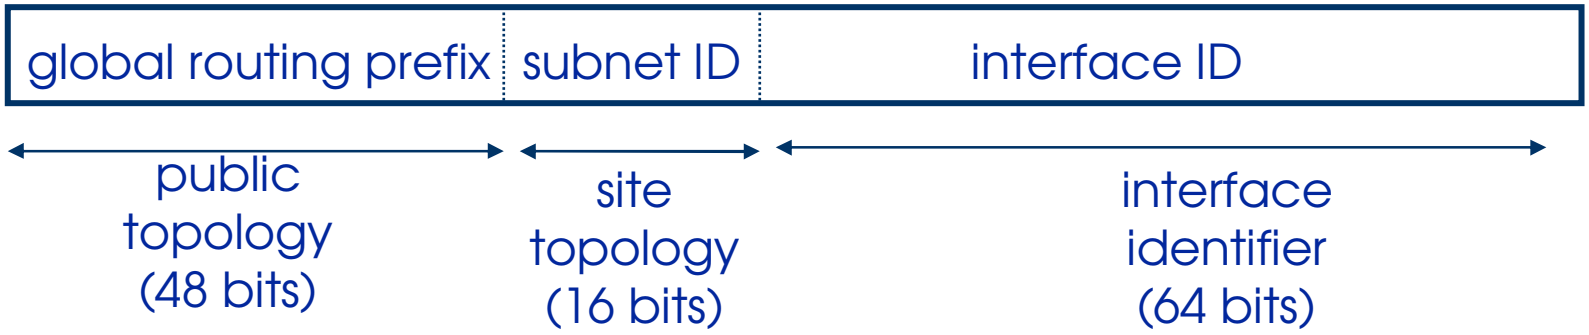

- global routing prefix
  - a (typically hierarchically-structured) value assigned to a site (a cluster of subnets/links)
- subnet ID
  - an identifier of a subnet within the site
- interface ID
  - constructed in Modified EUI-64 format

Link-Local Unicast Addresses

- meaningful only in a single link zone, and may be re-used on other links
- Link-local addresses for use during auto-configuration and when no routers are present
- Required for Neighbor Discovery process, always automatically configuration
- An IPv6 router never forwards link-local traffic beyond the link
- Prefix= FE80::/64

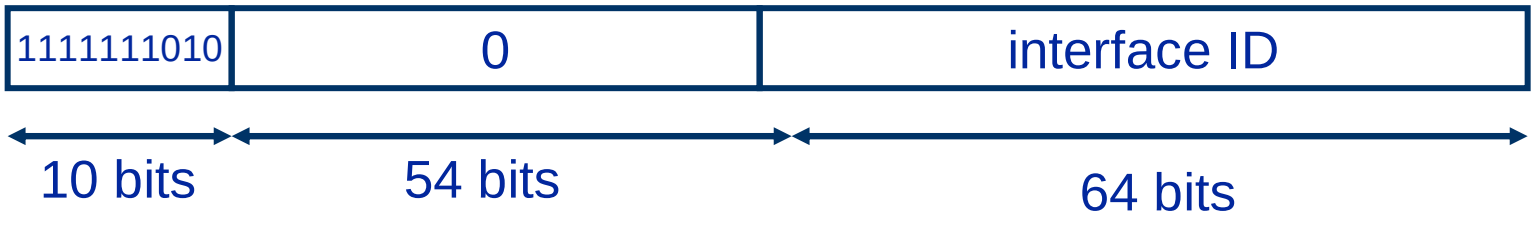

Site-Local Unicast Addresses

- meaningful only in a single site zone, and may be re-used in other sites
- Equivalent to the IPv4 private address space
- Address are not automatically configured and must be assigned
- Prefix= FEC0::/48

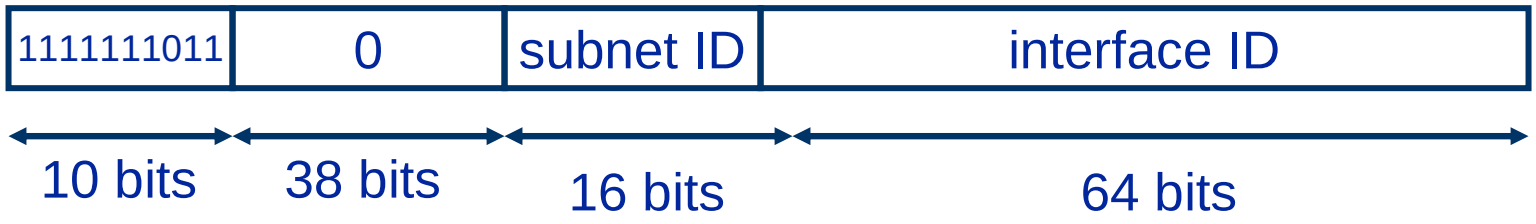

Special IPv6 address

- Loopback address (0:0:0:0:0:0:0:1 or ::1)
  - Identify a loopback interface
- IPv4-compatible address (0:0:0:0:0:0:w.c.x.z or ::w.c.x.z)
  - Used by dual-stack nodes
  - IPv6 traffic is automatically encapsulated with an IPv4 header and send to the destination using the IPv4 infrastructure
- IPv4 mapped address (0:0:0:0:0:FFFF:w.c.x.z or ::FFFF:w.c.x.z)
  - Represent an IPv4-only node to an IPv6 node
  - Only use a single listening socket (AF_INET6) to handle connections from client via both IPv6 and IPv4 protocols.
  - Never used as a source or destination address of IPv6 packet
  - Rarely implemented

Multicast IPv6 addresses

- Multicast address can not be used as source or as intermediate destination in a Routing header
- low-order Transient(T) flag indicates permanent (T=0) / transient(T=1) group; three other flags reserved
- Scope field
  - 1: node-local
  - 2: link-local
  - 5: site-local
  - 8: organization-local
  - E: global
  - Others: reserved

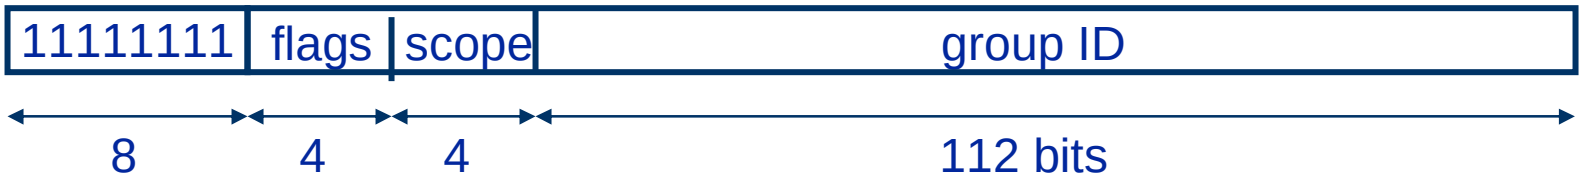

Multicast IPv6 addresses(cont.)

- Special multicast IPv6 address
  - FF01::1
    - Node-local scope all-nodes multicast address
  - FF02::1
    - Link-local scope all-nodes multicast address
  - FF01::2
    - Node-local scope all-routers multicast address
  - FF02::2
    - Link-local scope all-Routers multicast address
  - FF05::5
    - site-local scope all-routers multicast address
- Use low-order 32 bits, each group ID maps to a unique Ethernet MAC address(RFC 2373)

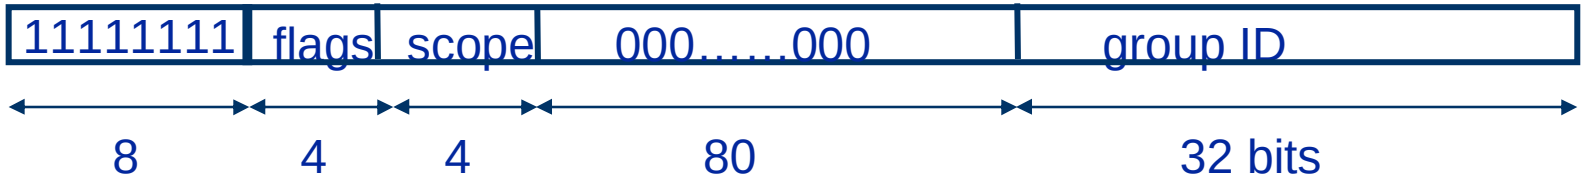

Example on FreeBSD

```
$ ping6 -c 5 FF02::2%em0
PING6(56=40+8+8 bytes) fe80::20f:eaff:fe4e:6a8c%em0 --> ff02::2%em0
16 bytes from fe80::20d:28ff:fe49:bea0%em0, icmp_seq=0 hlim=64 time=0.715 ms
16 bytes from fe80::20d:65ff:fee9:6c00%em0, icmp_seq=0 hlim=64 time=0.862 ms(DUP!)
16 bytes from fe80::20d:28ff:fe49:bea0%em0, icmp_seq=1 hlim=64 time=0.613 ms
16 bytes from fe80::20d:65ff:fee9:6c00%em0, icmp_seq=1 hlim=64 time=0.860 ms(DUP!)
16 bytes from fe80::20d:28ff:fe49:bea0%em0, icmp_seq=2 hlim=64 time=0.610 ms
16 bytes from fe80::20d:65ff:fee9:6c00%em0, icmp_seq=2 hlim=64 time=0.745 ms(DUP!)
16 bytes from fe80::20d:28ff:fe49:bea0%em0, icmp_seq=3 hlim=64 time=0.730 ms
16 bytes from fe80::20d:65ff:fee9:6c00%em0, icmp_seq=3 hlim=64 time=0.864 ms(DUP!)
16 bytes from fe80::20d:28ff:fe49:bea0%em0, icmp_seq=4 hlim=64 time=0.721 ms

--- FF02::2%em0 ping6 statistics ---
5 packets transmitted, 5 packets received, +4 duplicates, 0.0% packet loss
round-trip min/avg/max/std-dev = 0.610/0.747/0.864/0.093 ms
```

IPv6 address for a Host

- Unicast addresses
  - A link-local address for each interface
  - Unicast address for each interface
    - Site-local, or
    - One or multiple aggregatable global unicast
  - A loopback address(::1)
- Multicast addresses
  - Node-local all-nodes multicast address(FF01::1)
  - Link-local all-nodes multicast address(FF02::1)
  - Multicast address of joined group

IPv6 address for a Router

- Unicast addresses
  - A link-local address for each interface
  - Unicast address for each interface
    - Site-local, or
    - One or multiple aggregatable global unicast
  - Subnet-Router anycast address
  - Additional anycast address(optional)
  - A loopback address(::1)
- Multicast addresses
  - Node-local all-nodes multicast address(FF01::1)
  - Node-local all-routers multicast address(FF01::2)
  - Link-local all-nodes multicast address(FF02::1)
  - Link-local all-routers multicast address(FF02::2)
  - Site-local all-routers multicast address(FF05:2)
  - Multicast address of joined group

IPv6 interface identifier

- Lowest-order 64-bit field of unicast address
- Globally unique or locally unique within a subnet
- Future higher-layer protocols may take advantage of globally-unique interface IDs to identify nodes independently of their current location
- Configure interface identifier
  - manual configuration
  - DHCPv6 (configures whole address)
  - automatic derivation from MAC address or other hardware serial number
  - pseudo-random generation (for client privacy)
  - the latter two choices enable “serverless” or “stateless” autoconfiguration, when combined with high-order part of the address learned via Router Advertisements

The conversion of a universally administered, unicast IEEE 802 address to an IPv6 interface identifier

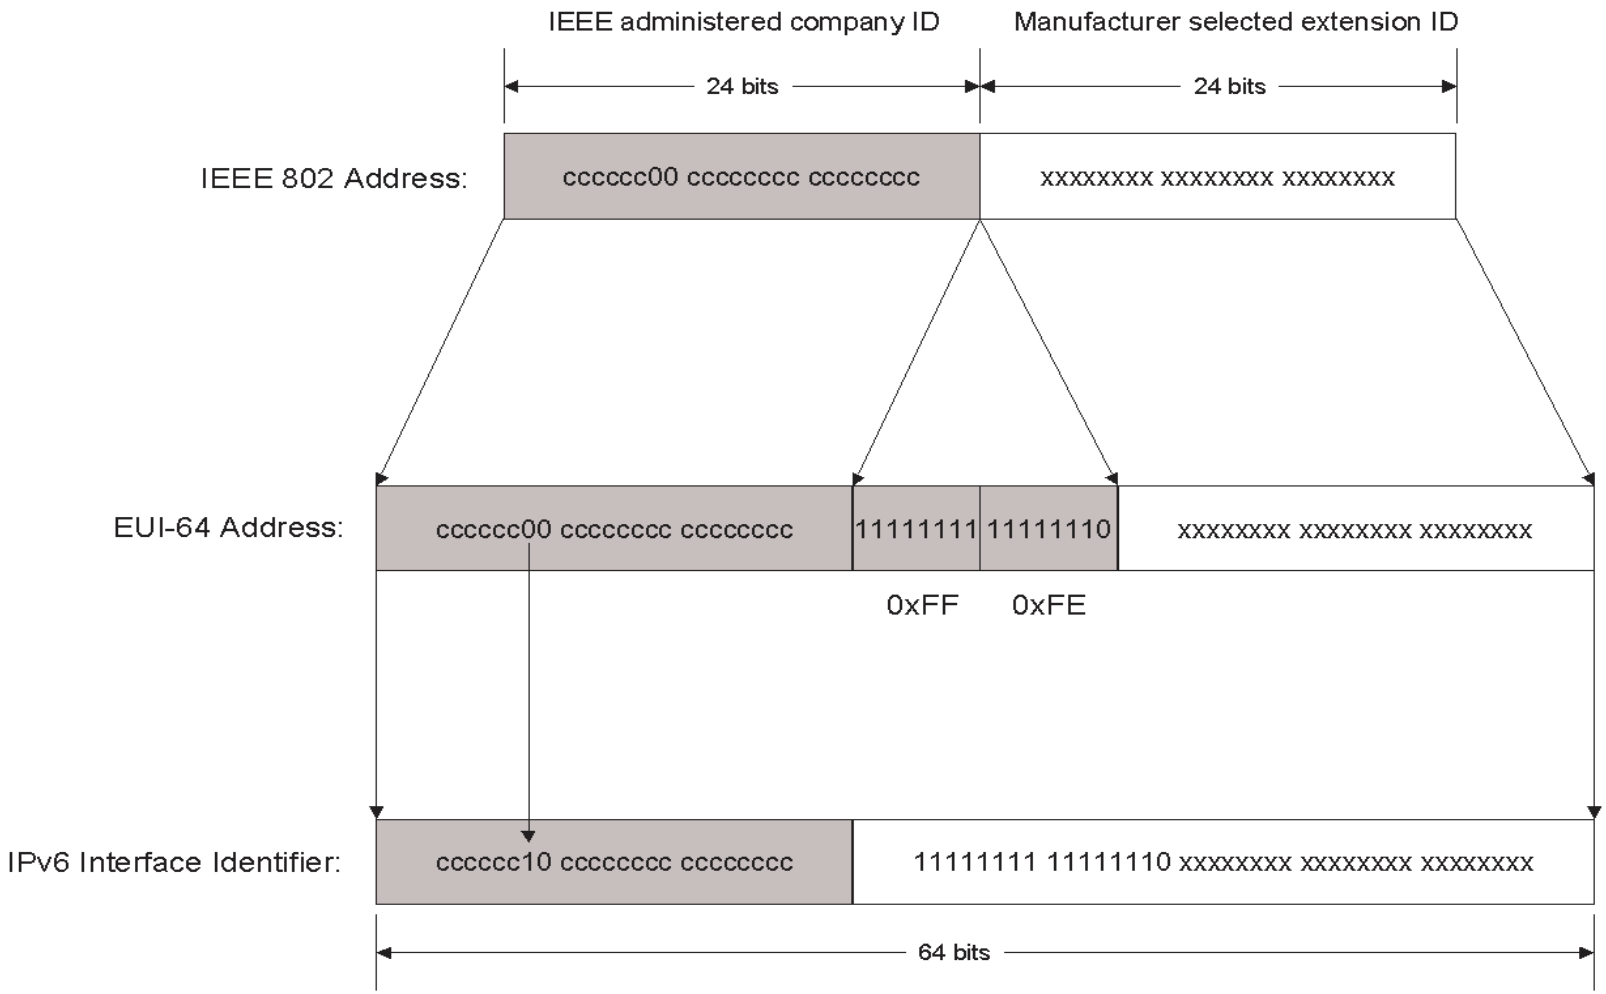

Difference between IPv4 and IPv6 addresses

Feature	IPv4	IPv6
Multicast address	224.0.0.0/4	FF00::/8
Unspecified address	0.0.0.0	::
Loopback address	127.0.0.1	::1
address	Public IP	Aggregatable global unicast
Broadcast address	Yes	No

Difference between IPv4 and IPv6 addresses (cont.)

Feature	IPv4	IPv6
Private IP address	10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16	Site-local(FEC0::/48)
DNS reverse resolution	IN-ADDR.ARPA domain	IP6.Arpa domain
DNS name resolution	IPv4 host address(A) resource record	IPv6 host address(AAAA) resource record
Text representation	Dotted decimal notation	Colon hexadecimal format with suppression of leading zero and zero compression. IPv4-compatible are expressed in Dotted decimal notation
Network bits representation	Subnet mask in dotted decimal notation or prefix length	Prefix length notation only
Autoconfigured addresses	169.254.0.0/16	Link-local(FE80::/64)

Introduction to IP version 6

Agenda

- Introduction
- IPv6 Addressing
- IPv6 Header
- Address Autoconfiguration
- Addressing Allocation Policy
- ICMPv6
- Neighbor Discovery

IPv4 Header

20 Octets+Options : 13 fields, include 3 flag bits

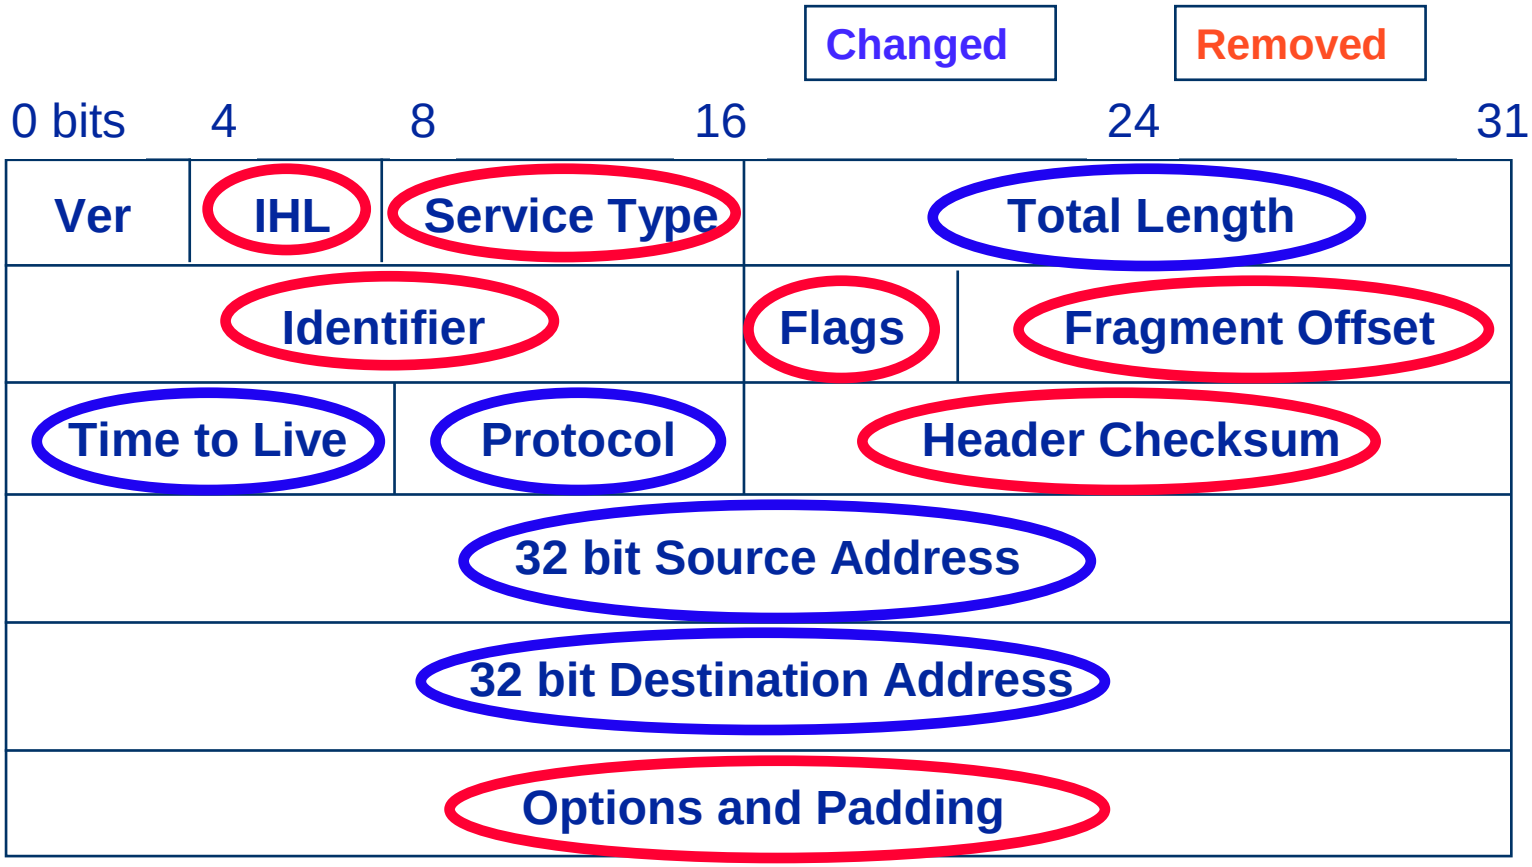

IPv6 Header

40 Octets, 8 fields

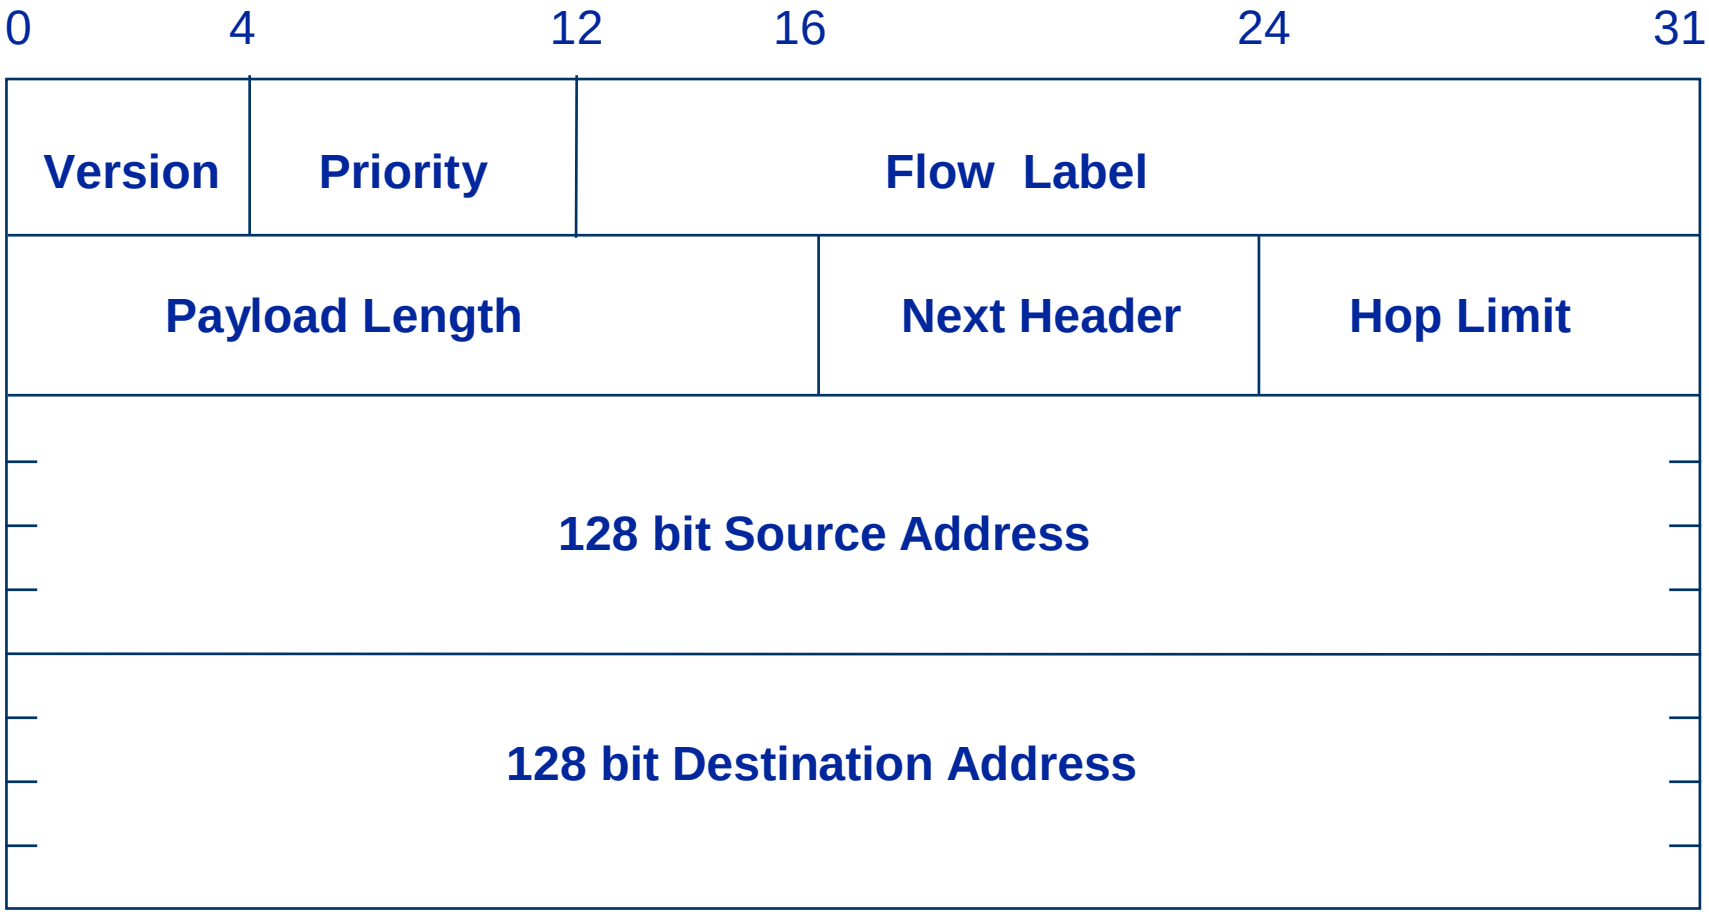

IPv6 vs. IPv4 Packet Data Unit

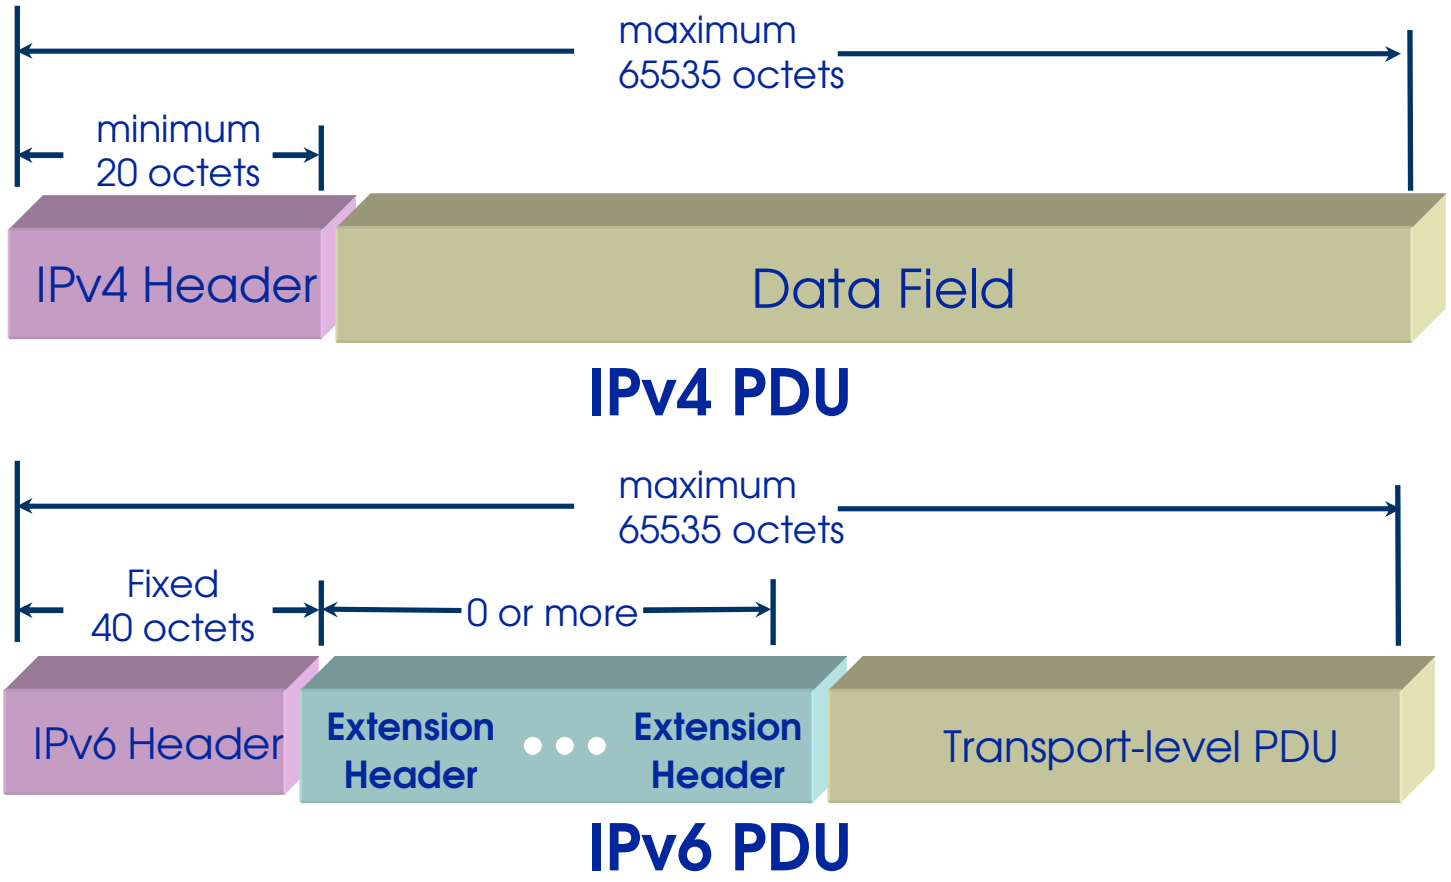

Comparison of IPv4 and IPv6 Header

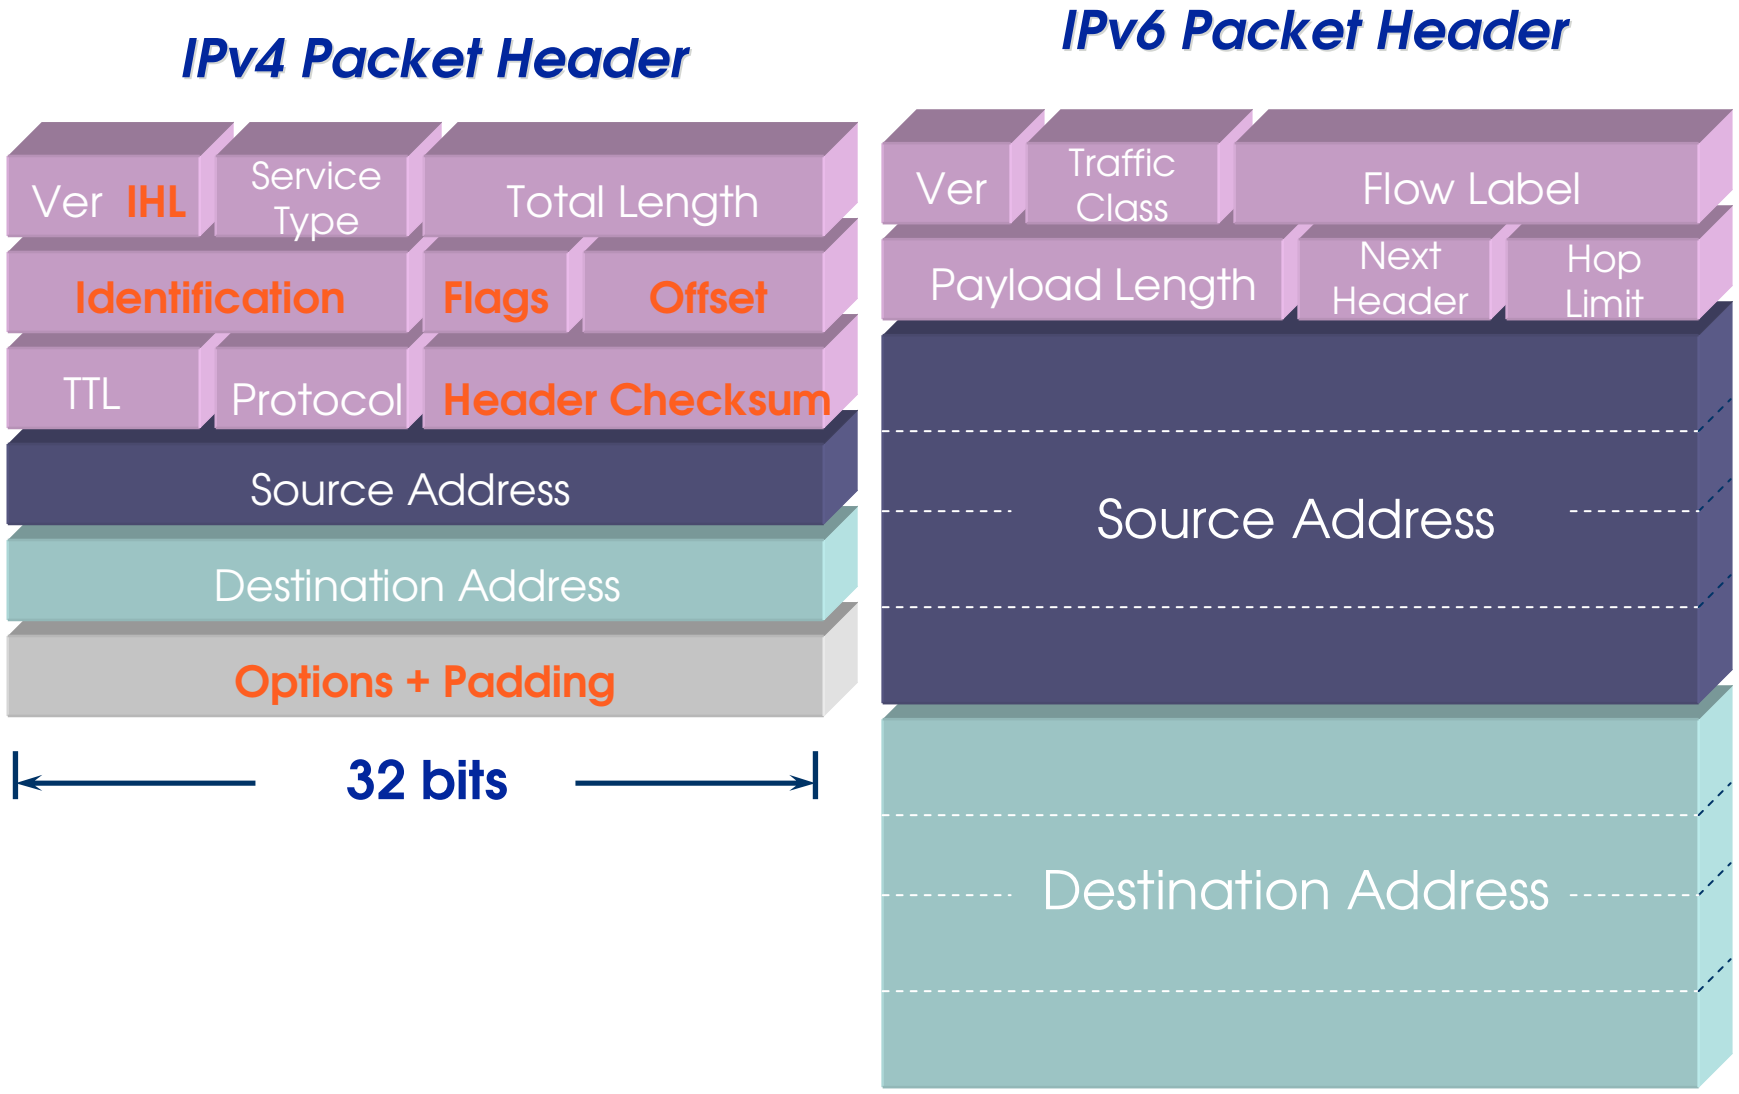

Summary of Header Changes between IPv4 & IPv6

- Streamlined
  - Fragmentation fields moved out of base header
  - IP options moved out of base header
  - Header Checksum eliminated
  - Header Length field eliminated
  - Length field excludes IPv6 header
  - Alignment changed from 32 to 64 bits
- Revised
  - Time to Live → Hop Limit
  - Protocol → Next Header
  - Precedence & TOS → Traffic Class
  - Addresses increased 32 bits → 128 bits
- Extended
  - Flow Label field added

Introduction to IP version 6

Agenda

- Introduction
- IPv6 Addressing
- IPv6 Header
- **Address Autoconfiguration**
- Addressing Allocation Policy
- ICMPv6
- Neighbor Discovery

IPv6 Auto-Configuration

- Stateless (RFC2462)
  - Host autonomously configures its own Link-Local address
  - Router solicitation are sent by booting nodes to request RAs for configuring the interfaces.
- Stateful
  - DHCPv6 (under definition at IETF)
- Renumbering

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

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

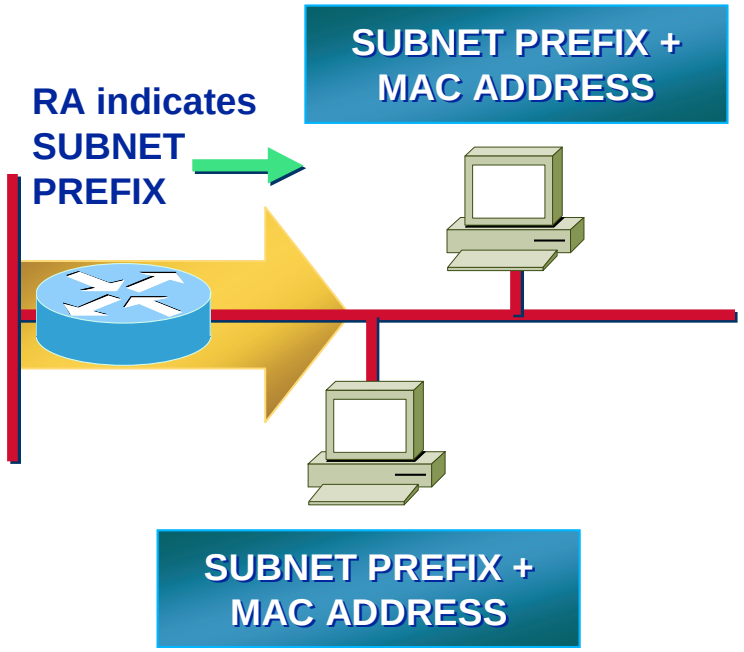

At boot time, an IPv6 host build a Link-Local address, then its global IPv6 address(es) from RA

Serverless Autoconfiguration ("Plug-n-Play")

- Hosts can construct their own addresses:
  - subnet prefix(es) learned from periodic multicast advertisements from neighboring router(s)
  - interface IDs generated locally
    1. MAC addresses
    2. pseudo-random temporary
- Other IP-layer parameters also learned from router adverts (e.g., router addresses, recommended hop limit, etc.)
- Higher-layer info (e.g., DNS server and NTP server addresses) discovered by multicast / anycast-based service-location protocol [details being worked out]
- DHCP also available for those who want more control

Auto-Reconfiguration ("Renumbering")

- New address prefixes can be introduced, and old ones withdrawn
  - we assume some overlap period between old and new, i.e., no "Flag Day"
  - hosts learn prefix lifetimes and preference order from router advertisements
  - old TCP connections can survive until end of overlap; new TCP connections can survive beyond overlap
- Router renumbering protocol, to allow domain-interior routers to learn of prefix introduction / withdrawal
- New DNS structure to facilitate prefix changes

Autoconfiguration address state

- Tentative(暫時的)
  - The address is in the process of being verified as unique
  - Verification is done through DAD (duplicate address detection)
- Preferred(偏好的)
- Deprecated(取代的)
- Valid
- Invalid

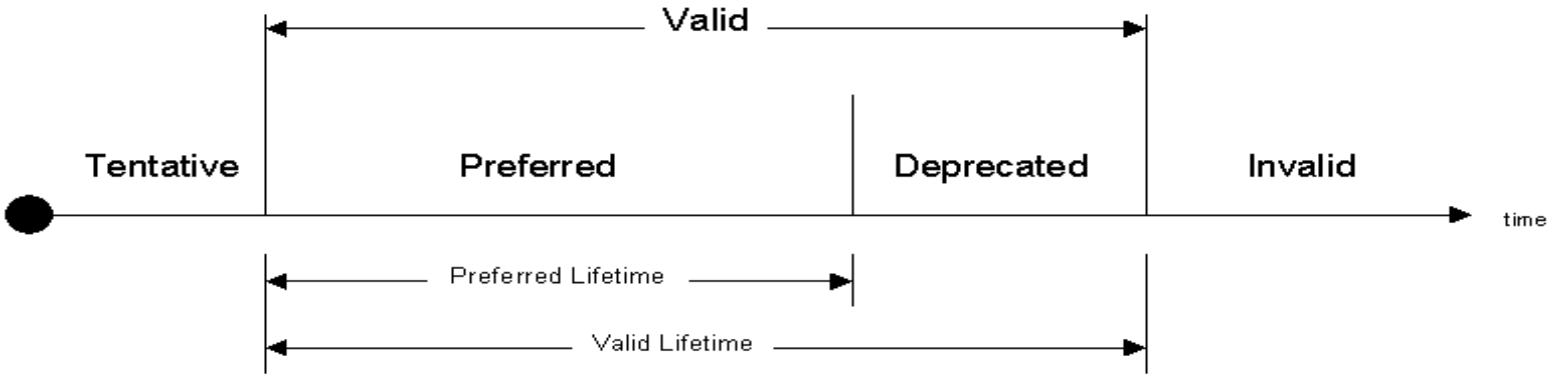

Introduction to IP version 6

Agenda

- Introduction
- IPv6 Addressing
- IPv6 Header
- Address Autoconfiguration
- Addressing Allocation Policy
- ICMPv6
- Neighbor Discovery

IPv6 Address Allocations

- RFC 3177
  - /32 for an ISP
  - /48 for an organization in the general case
  - /64 when it is known that one and only one subnet is needed
  - /128 when it is absolutely known that one and only one device is connecting

- This document also describe the advantage of the fixed boundary specifically at /48

International IPv6 Address Management (cont'd)

Initial IPv6 Prefix Allocation for RIRs

IPv6 Prefix Range	Assignment
2001:0000::/29-2001:01F8::/29	IANA
2001:0200::/29-2001:03F8::/29	APNIC
2001:0400::/29-2001:05F8::/29	ARIN

IPv6 Prefix Range	Assignment
2001:0600::/29-2001:07F8::/29	RIPE NCC
2001:1200::/29-2001:13F8::/29	LACNIC

IPv6 Prefix Allocation in Taiwan

- 2001:238::/32 - HiNet
- 2001:288::/32 - TANet
- 2001:C08::/32 - ASNet
- 2001:C50::/32 - TTN
- 2001:C58::/32 - 6REN
- 2001:CA0::/32 - CHT TL
- 2001:CD8::/32 - SeedNet
- 2001:D20::/32 - TFN
- 2001:D40::/32 - TW NTT
- 2001:E10::/32 - TWAREN
- 2001:ED8::/32 - ITRI
- 2001:F18::/32 - NCTU

Total number of allocated IPv6 prefixes on 02/10/2005

RIR	Size in /48s	Count
APNIC	1021935630	404
ARIN	10944560	215
LACNIC	2818049	37
RIPE NCC	1093492736	659

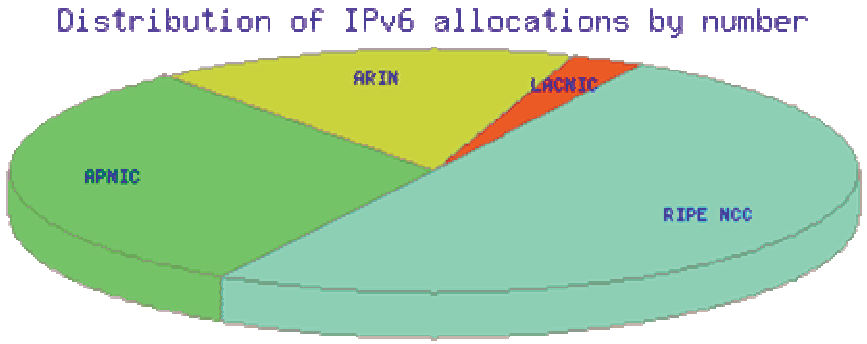

http://www.apnic.net/info/reports/index.html

Total IPv6 Allocations from RIRs to LIRs/ISPs Top 10 Countries

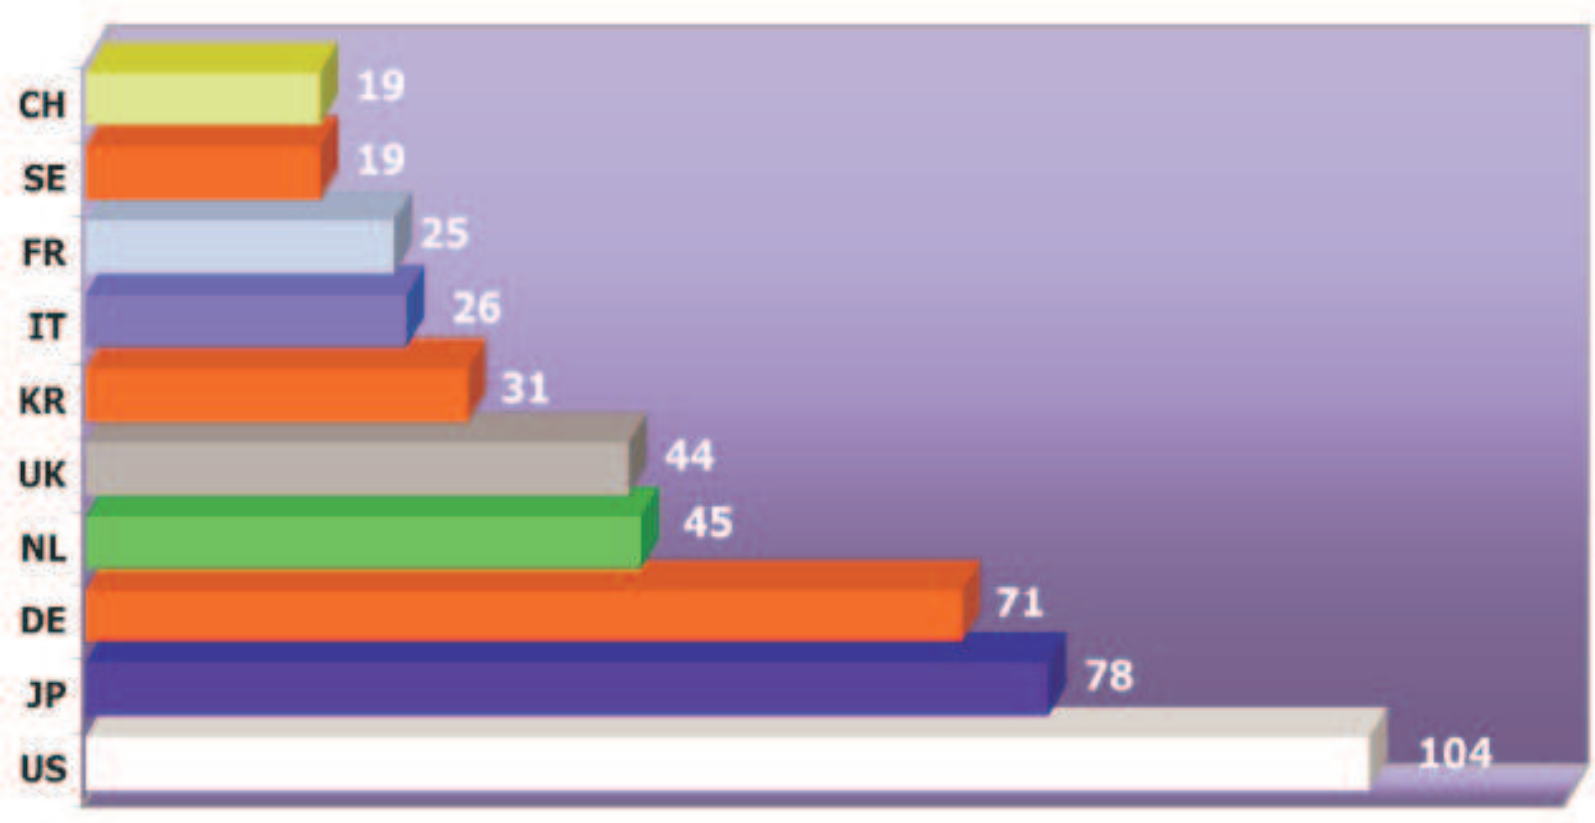

IPv4 /8 Address Space Status

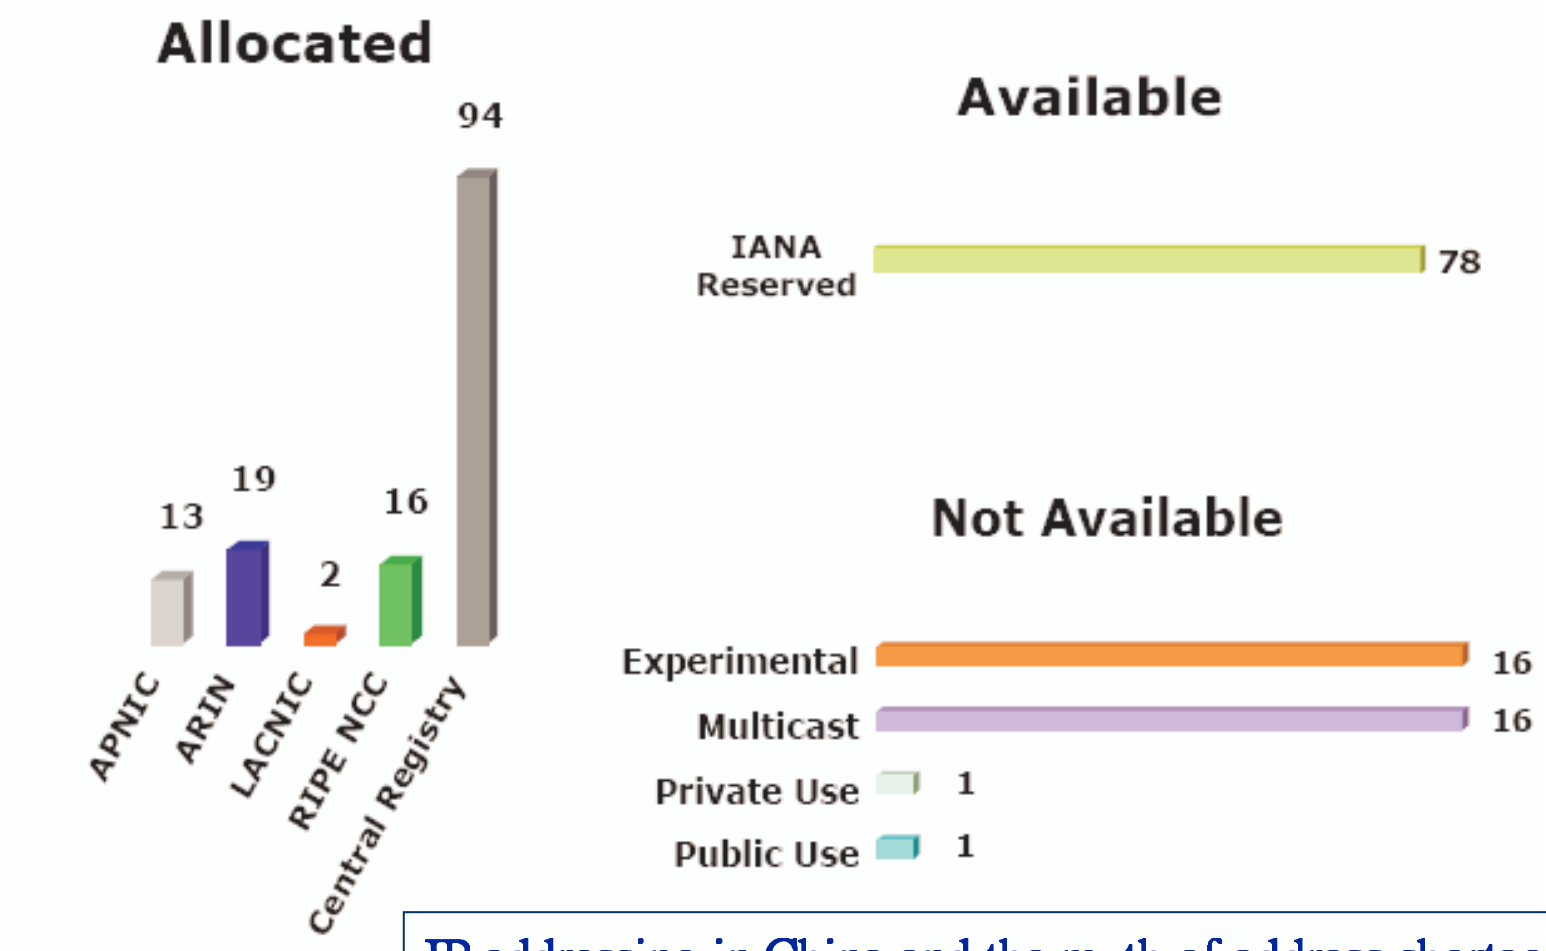

IP addressing in China and the myth of address shortage
http://www.apnic.net/news/hot-topics/index.html#ip-addressing

Introduction to IP version 6

Agenda

- Introduction
- IPv6 Addressing
- IPv6 Header
- Address Autoconfiguration
- Addressing Allocation Policy
- ICMPv6
- Neighbor Discovery

Features of ICMPv6

- An integral part of IPv6 and MUST be fully implement by every IPv6 node (RFC 2463)
- Next Header value= 58
- Report delivery or forwarding errors
- Provide simple echo service for troubleshooting
- Multicast Listener discovery(MLD) – 3 ICMP messages
- Neighbor Discovery(ND) – 5 ICMP messages

ICMPv6 message format

Type [8]	Code [8]	Checksum [16]
Message Body [N*32]		

Field Name	Bit Length	Description
Checksum	16	The checksum is a 16 bit CRC check calculated over the entire ICMPv6 message, including an IPv6 "Pseudo Header" as defined in Section 8.1 of [IPv6].
Code	8	The code field further defines the message content
Type	8	The Type field identifies the message contents. ICMPv6 messages are divided into two classes; Error messages and Informational messages. The Type values of Error messages go from 0 to 127. Informational messages have type values from 128 to 255.
Message Body	N*32	The message body depends on the message type. Every ICMPv6 error message includes as much of the offending packet as can be accommodated without exceeding the minimum IPv6 MTU (1280 bytes).

Two types of ICMP messages

- Error messages
  - Report error in the forwarding or delivery

Type	Meaning
1	Destination Unreachable
2	Packet too Big
3	Time exceeded
4	Parameter Problem

- Informational messages
  - Provide diagnostic function, MLD, and ND

Type	Meaning
128	Echo Request
129	Echo Reply

Error message (Destination Unreachable)

- Send by router or destination host

Type [8] = 1	Code [8]	Checksum [16]
Unused [32]		
Offending Packet [N*32]		

Code	Meaning
0	No route to destination
1	Communication with destination prohibited
3	Address unreachable
4	Port unreachable

Error message (Packet Too Big)

- Send when link MTU is smaller than the size of packet
- Used for IPv6 Path MTU Discovery process

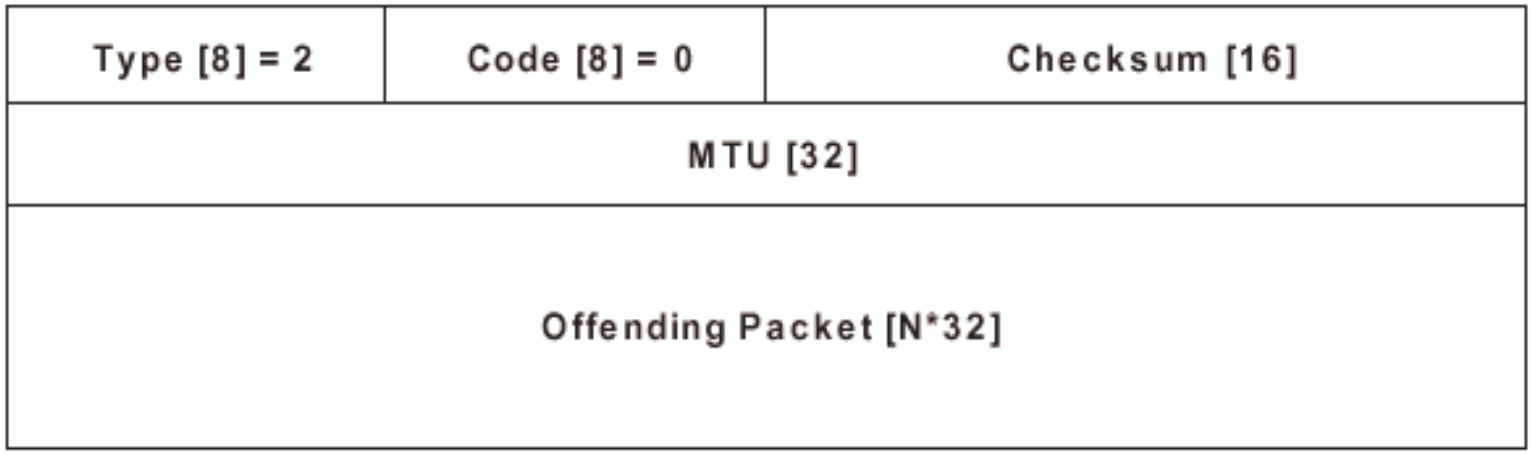

Error message (Time Exceeded)

- Send by router when Hop limit field is zero

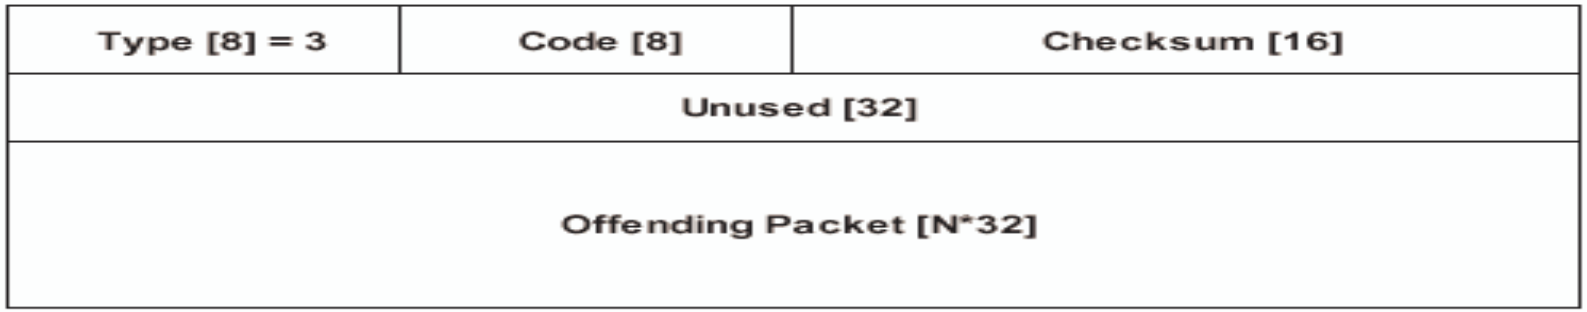

  - Code field:
    - 0: Hop limit= 0
      - Hop limit of outgoing packets is not large enough to reach destination, or
      - Routing loop exist
    - 1: fragmentation reassembly time of destination host is exceeded

Error message (Parameter Problem)

- Send by router or destination host when errors of IPv6 header or extension header

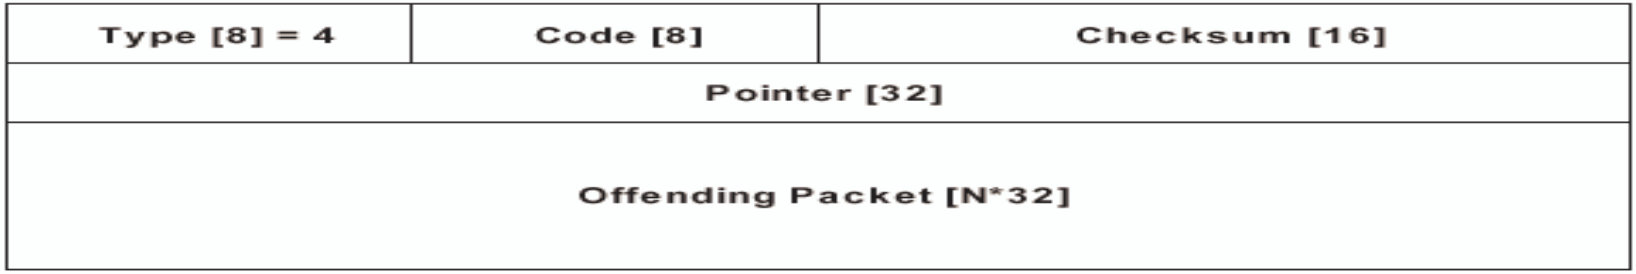

Code	Meaning
0	Erroneous Header Field
1	Unrecognized Next Header type
2	Unrecognized IPv6 Option

Informational message

- Echo Request message

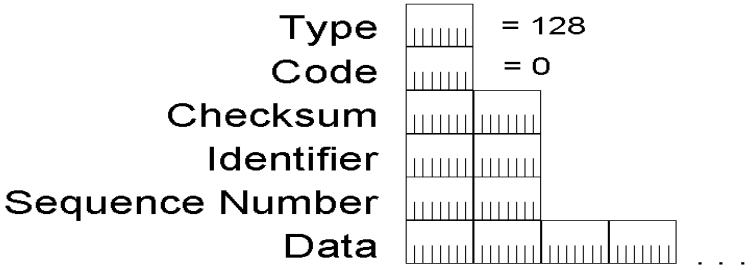

- Echo Reply message

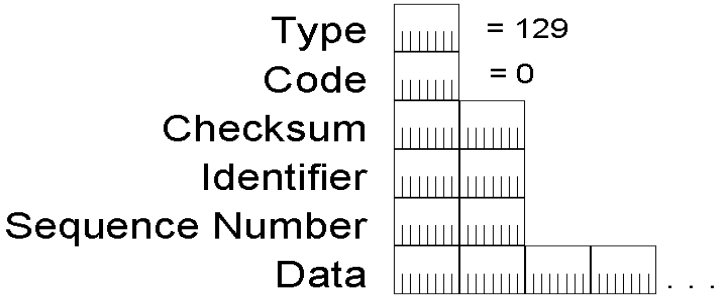

- Identifier and Sequence Number are send by host and used to match incoming Echo Reply with corresponding Echo Request(same as IPv4)
- Multicast Listener Query messages:
  - Query, Report, done(like IGMP for IPv4)

Comparing ICMPv4 and ICMPv6

ICMPv4	ICMPv6
Destination Unreachable-Network unreachable (Type 3, Code 1)	Destination Unreachable-No route to destination (Type 1, Code 0)
Destination Unreachable-Host unreachable (Type 3, Code 1)	Destination Unreachable-Address unreachable (Type 1, Code 3)
Destination Unreachable-Protocol unreachable (Type 3, Code 2)	Parameter Problem-Unrecognized Next Header field (Type 4, Code 1)
Destination Unreachable-Port unreachable (Type 3, Code 3)	Destination Unreachable-Port unreachable (Type 1, Code 4)
Destination Unreachable-Fragmentation needed and DF set (Type 3, Code 4)	Packet Too Big (Type 2, Code 0)
Destination Unreachable-Communication with destination host administratively prohibited (Type 3, Code 10)	Destination Unreachable-Communication with destination administratively prohibited (Type 1, Code 1)

Comparing ICMPv4 and ICMPv6(cont.)

ICMPv4	ICMPv6
Time Exceeded-TTL expired (Type 11, Code 0)	Time Exceeded-Hop Limit exceeded (Type 3, Code 0)
Time Exceeded-Fragmentation timer expired (Type 11, Code 1)	Time Exceeded-Fragmentation timer exceeded (Type 3, Code 1)
Parameter Problem (Type 12, Code 0)	Parameter Problem (Type 4, Code 0 or Code 2)
Source Quench (Type 4, Code 0)	N.A.
Redirect (Type 5, Code 0)	Neighbor Discovery Redirect message (Type 137, Code 0)

Minimum MTU

- Link MTU
  - A link's maximum transmission unit (ex: the max IP packet size that can be transmitted over the link)
- Path MTU
  - The minimum link MTU of all the links in a path between a source and a destination
- Minimum link MTU for IPv6 is 1280 octets vs 68 octets for IPv4
- On links that have a configurable MTU, it's recommended a MTU of 1500 bytes

RFC 1981 - Path MTU Discovery

- Implementations are expected to perform path MTU discovery to send packets bigger than 1280 octets
  - For each destination, start by assuming MTU of first-hop link
  - If a packet reach a link in which it can't fit, will invoke ICMP "packet too big" message to source, reporting the link's MTU; MTU is cached by source for specific destination
- Minimal implementation can omit path MTU discovery as long as all packets kept <= 1280 octets
  - Ex: in a boot ROM implementation
- The PMTU of a path may change over time, due to changes in the routing topology.
  - Reductions of the PMTU are detected by Packet Too Big messages.
  - Occasionally discard cached MTU to detect possible increase

Introduction to IP version 6

Agenda

- Introduction
- IPv6 Addressing
- IPv6 Header
- Address Autoconfiguration
- Addressing Allocation Policy
- ICMPv6
- Neighbor Discovery

RFC 2461 - Neighbor Discovery(ND)

- Node(Hosts and Routers) use ND to determinate the link-layer addresses for neighbors known to reside on attached links and quick purge cached valued that become invalid
- Hosts also use ND to find neighboring router that willing to forward packets on their behalf
- Nodes use the protocol to actively keep track of which neighbors are reachable and which are not, and to detect changed link-layer addresses
- Replace ARP, ICMP Router Discovery, and ICMP Redirect used in IPv4

IPv6 ND processes

- Router discovery
  - Discover the local hosts on an attached link
  - Equivalent to ICMPv4 Router Discovery
- Prefix discovery
  - Discovery the network prefix
  - Equivalent to ICMPv4 Address Mask Request/Reply
- Parameter discovery
  - Discovery additional parameter(ex: link MTU, default hop limit for outgoing packet)
- Address autoconfiguration
  - Configure IP address for interfaces
- Address resolution
  - Equivalent to ARP in IPv4

IPv6 ND processes(cont.)

- Next-hop determination
  - Destination address, or
  - Address of an on-link default router
- Neighbor unreachable detection(NUD)
- Duplicate address detection(DAD)
  - Determine that an address considered for use is not already in use by a neighboring node
- First-hop Redirect function
  - Inform a host of a better first-hop IPv6 address to reach a destination
  - Equivalent to ICMPv4 Redirect

ND message format

- 5 ND messages:
  - Router solicitation
  - Router Advertisement
  - Neighbor Solicitation
  - Neighbor Advertisement
  - Redirect
- All ND message are send with hop limit= 255.
  - If it is not set to 255, the message is silently discarded
  - Provide Protection from ND-based network attacks launched from off-link nodes
  - Router can not have forwarded the ND message from an off-link node

Neighbor Discovery options

- Source/Target link-layer address option
  - Source link-layer address
    - Indicate the link-layer address of the ND sender
    - Included in Neighbor Solicitation, Router Solicitation, and Router Advertisement
    - Type = 1
  - Target link-layer address
    - Indicate the link-layer address of the neighbor node
    - Included in Neighbor Advertisement and Redirect
    - Type = 2

Neighbor Discovery options(cont.)

- Prefix information option
  - Indicate both address prefixes and information about address autoconfiguration
  - Included in Router Advertisement
  - Can be multiple prefix information options in Router Advertisement message

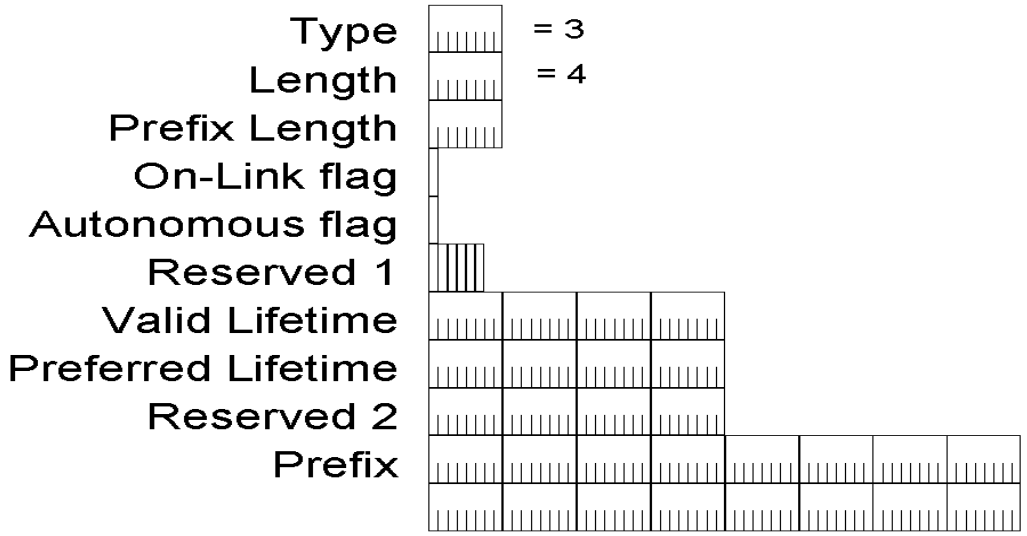

  - Autonomous flag: stateless address configuration

Neighbor Discovery options(cont.)

- Redirect header option

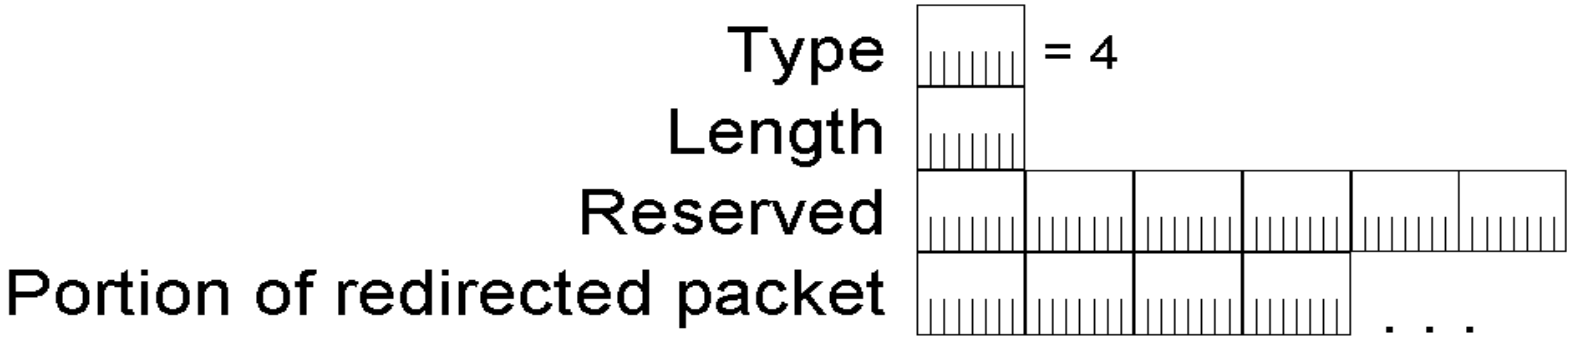

- MTU option

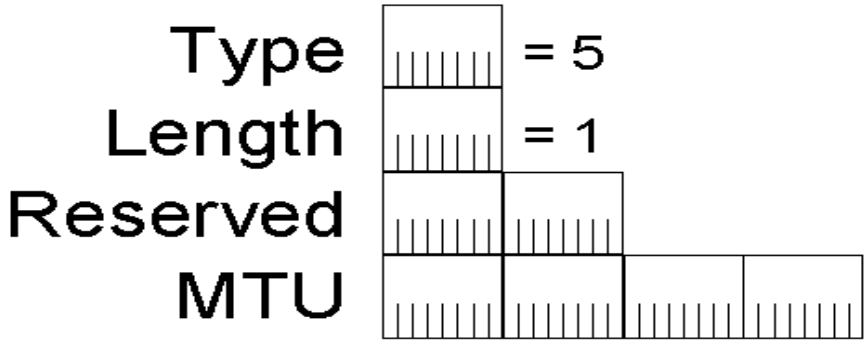

ND Autoconfiguration, Prefix & Parameter Discovery

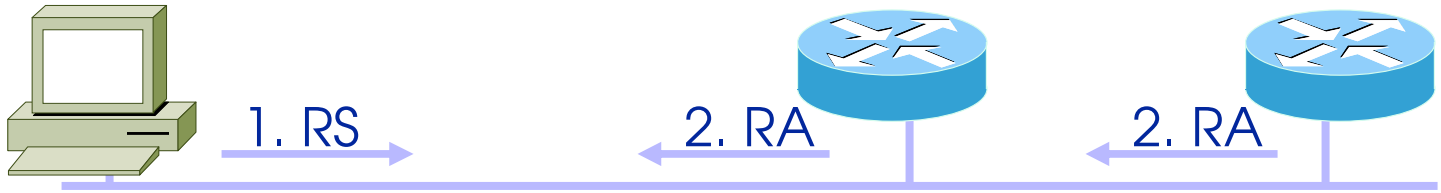

1. RS:

ICMP Type = 133

Src = ::

Dst = All-Routers multicast Address (FF02::2)

query= please send RA

2. RA:

ICMP Type = 134

Src = Router Link-local Address

Dst = All-nodes multicast address

Data= options, prefix, lifetime, autoconfig flag

- Router solicitation are sent by booting nodes to request RAs for configuring the interfaces.

ND Address Resolution & Neighbor Unreachability Detection

ICMP type = 135 (NS)
Src = A

Dst = Solicited-node multicast of B

Data = link-layer address of A
Query = what is your link address?

ICMP type = 136 (NA)
Src = B
Dst = A
Data = link-layer address of B

A and B can now exchange packets on this link

ND Redirect

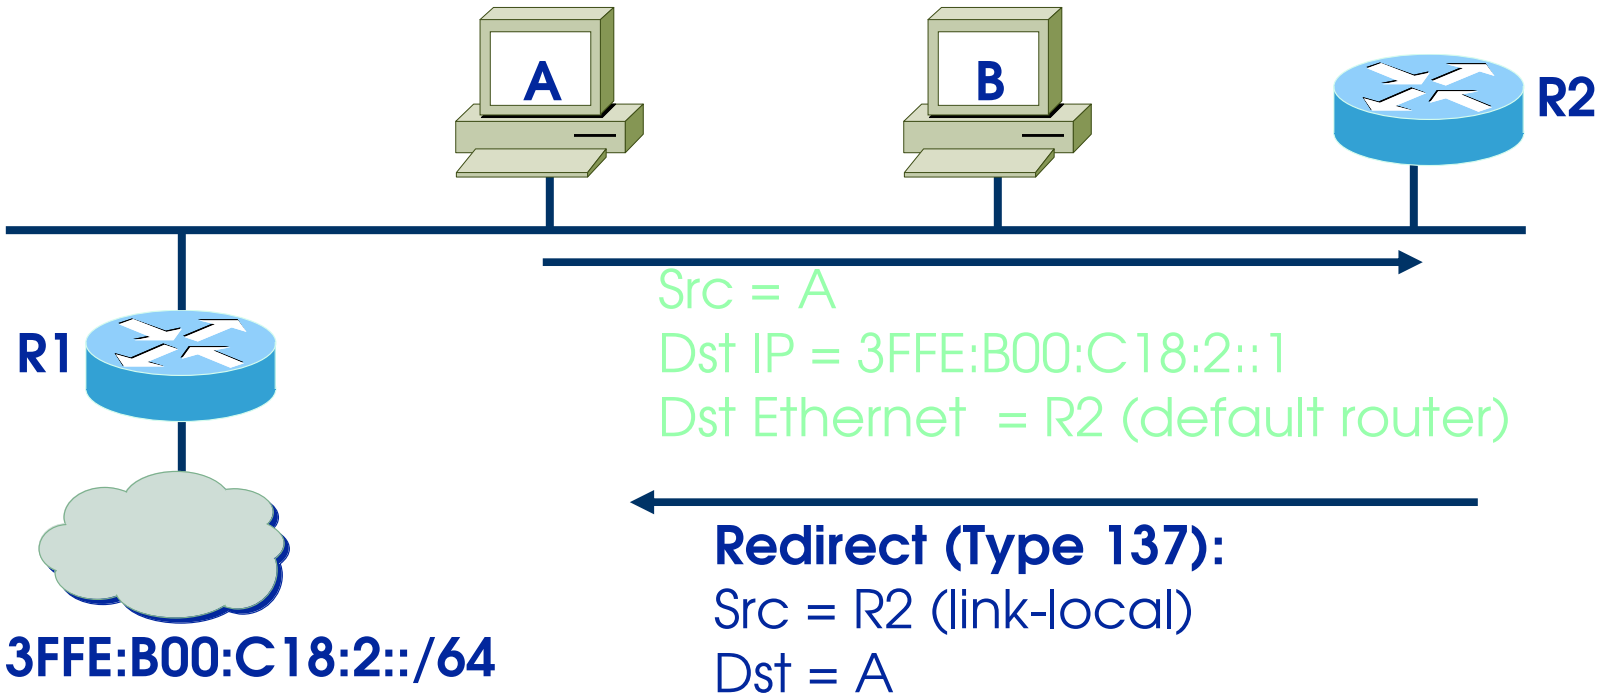

- Redirect is used by a router to signal the reroute of a packet to an onlink host to a better router or to another host on the link

ND Duplicate Address Detection

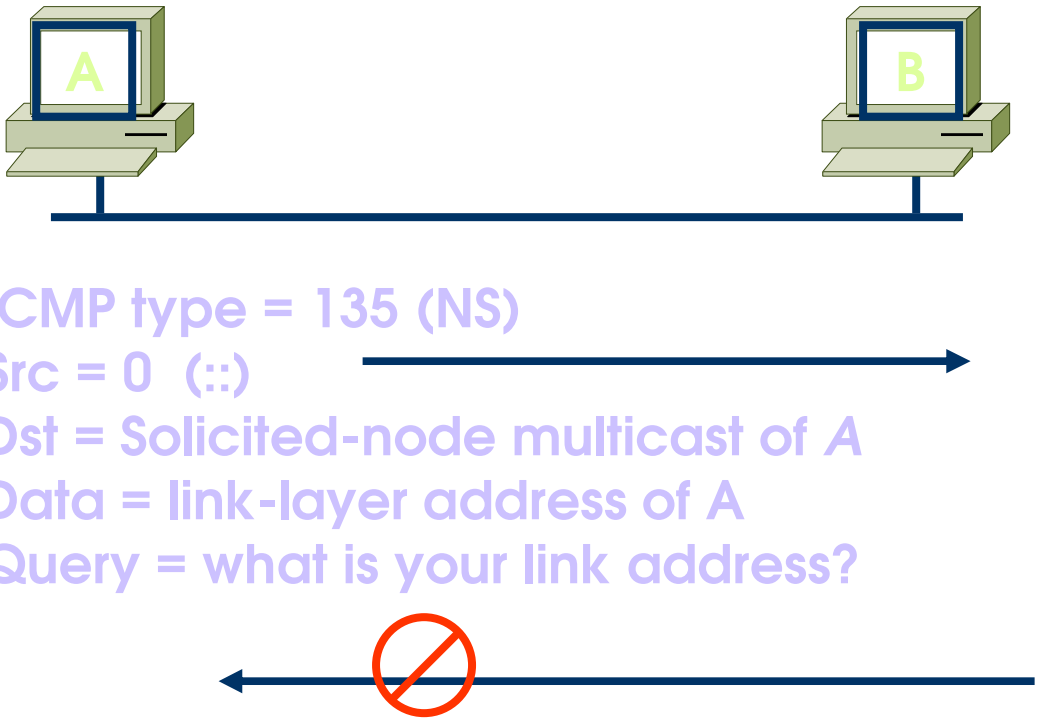

- Duplicate Address Detection (DAD) uses neighbor solicitation to verify the existence of an address to be configured.

List of Acronyms

- ARP：Address Resolution Protocol
- DAD：Duplicate address detection
- DHCP：Dynamic Host Configuration Protocol
- ICMP：Internet Control Message Protocol
- IEEE：Institute of Electrical and Electronic Engineers
- IGMP：Internet Group Management Protocol
- IPv6：Internet Protocol version 6
- MLD：Multicast Listener Discovery
- MTU：Maximum Transmission Unit
- ND：Neighbor Discovery
- NUD：Neighbor unreachable detection
- PDU：Packet Data Unit