INTRODUCTION TO IPV6

ITU IPv6 and IoT Workshop

GLOBAL INTERNET MANAGEMENT

- Several organizations form a framework for global Internet governance.
  - Internet Assigned Numbers Authority (IANA)
  - 5 Regional Internet Registries (RIR)
- The 5 RIRs are geographical distributed

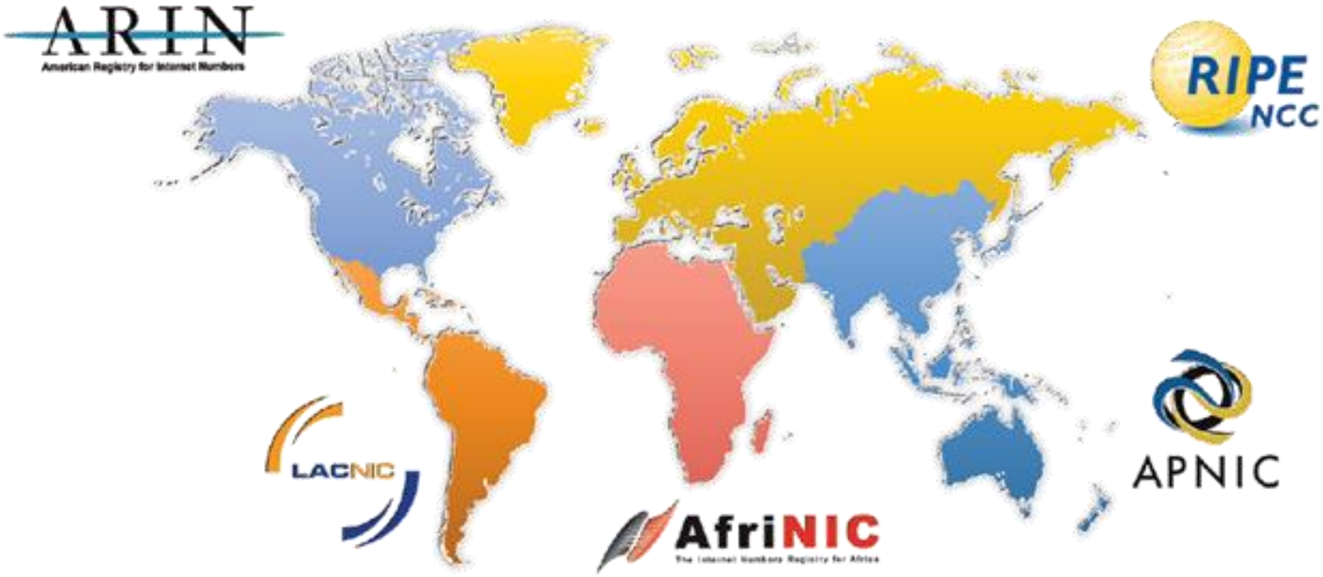
ARIN
American Registry for Internet Numbers
RIPE NCC
LACNIC
AfriNIC
The Internet Numbers Registry for Africa
APNIC

ROLES OF AN RIR

- ❑IPv6 address allocation, management, and deployment measurement
- ❑Research, education, and information distribution about IPv6
- ❑Community outreach and liaison.
- ❑Representation in forums, such as the ITU, OECD, the Internet Governance Forum (IGF), and ICANN

WHAT IS THE HIERARCHY FOR THE GLOBAL ADDRESS ALLOCATION?

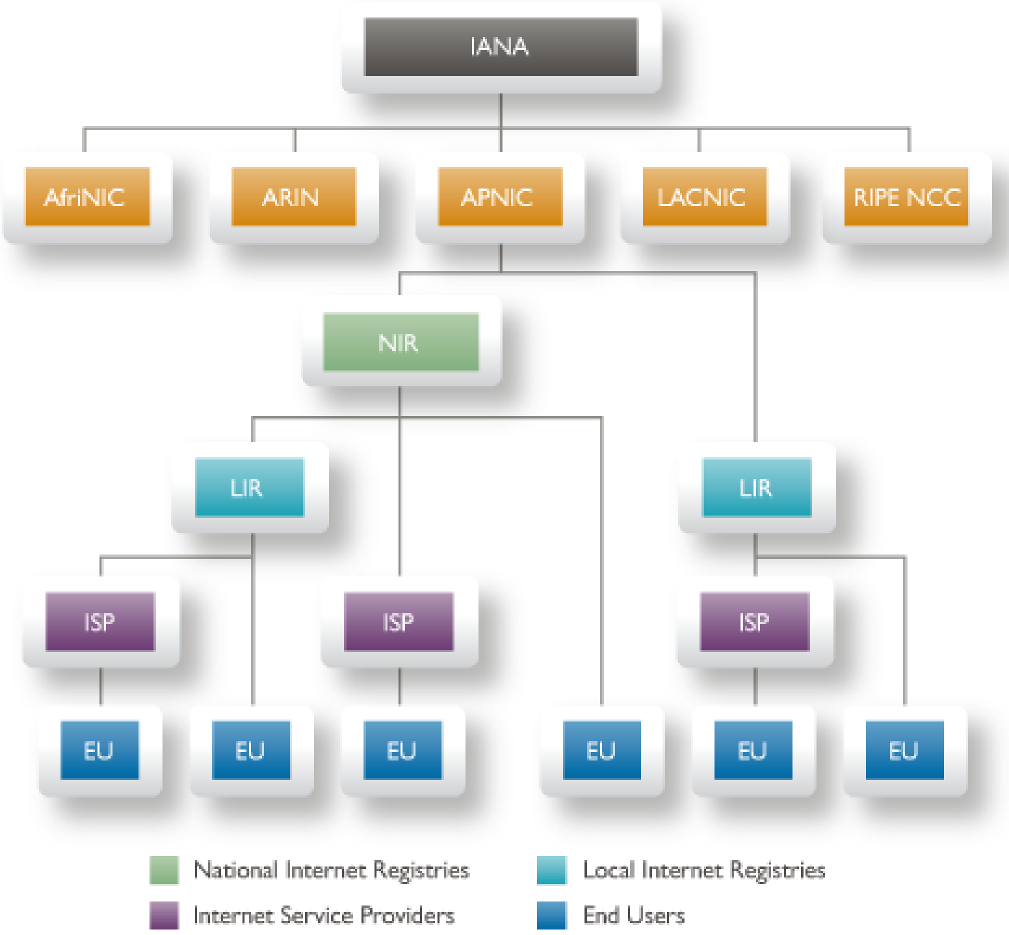
IANA
AfriNIC
ARIN
APNIC
LACNIC
RIPE NCC
NIR
LIR
LIR
ISP
ISP
ISP
EU
EU
EU
EU
EU
EU
National Internet Registries
Local Internet Registries
Internet Service Providers
End Users

http://www.apnic.net/policy/ipv6-address-policy

ISSUES WITH IPV4

Can IPv4 address be depleted? *Answer:* **Yes & No**

http://www.iana.org/assignments/ipv4-address-space/

Last updated: **April 2011**

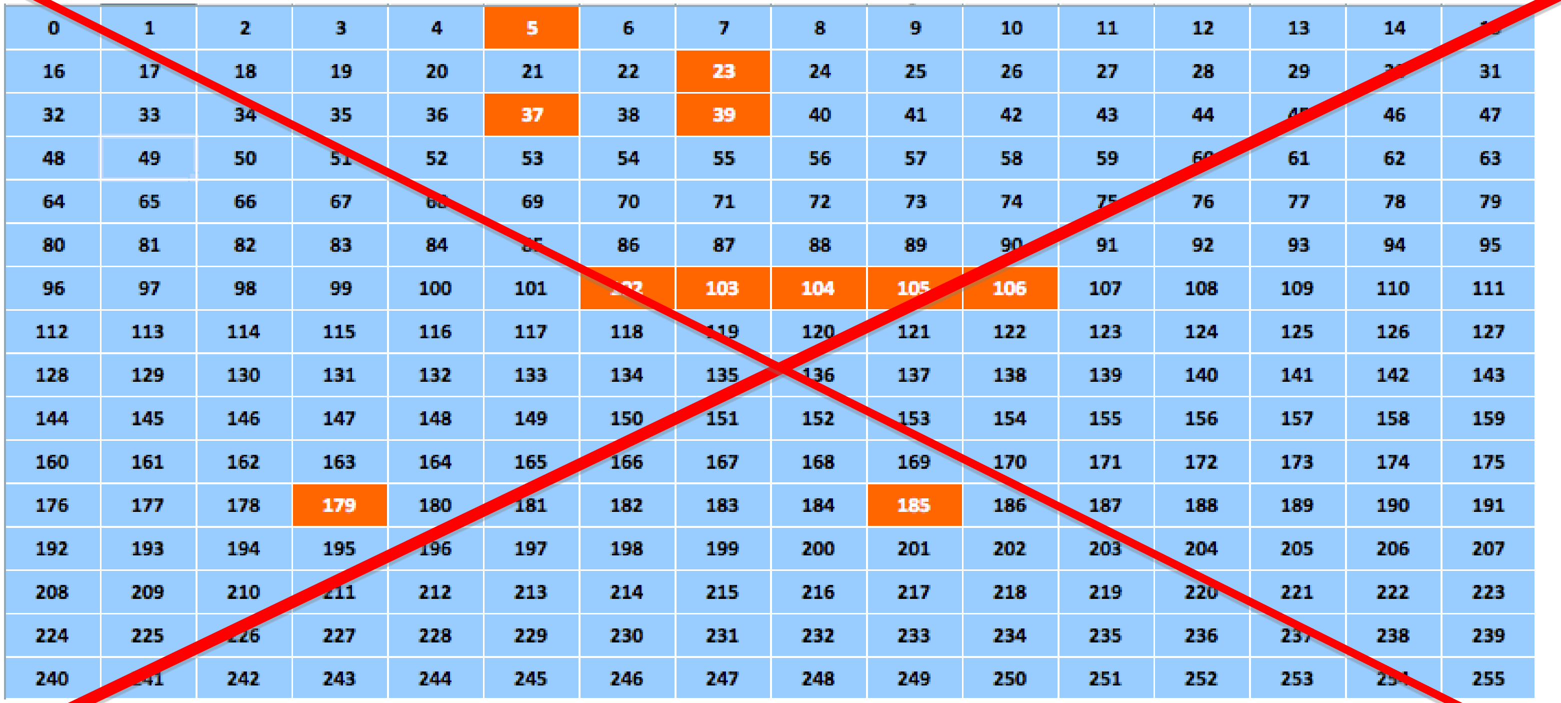

0	1	2	3	4	🟧 5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	🟧 23	24	25	26	27	28	29	30	31
32	33	34	35	36	🟧 37	38	🟧 39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	🟧 102	🟧 103	🟧 104	🟧 105	🟧 106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	🟧 179	180	181	182	183	184	🟧 185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

🟧 **Unallocated**

🟦 **Allocated**

IPV4 ADDRESSES ARE RUNNING OUT.

- The big blocks of IPv4 addresses that are assigned by IANA was exhausted around April 2011.
- RIRs running out IPv4 that cause ISPs, wireless carriers, governments, and major corporations suffers from lack of IPv4 address
- Old address blocks will have to be better managed, and split the old address blocks even further, i.e. further subnets
  - This causes more routing fragmentation and performance issues.
- Organizations try to extend IPv4 lifetime using CIDR and NAT

IPCALYSE

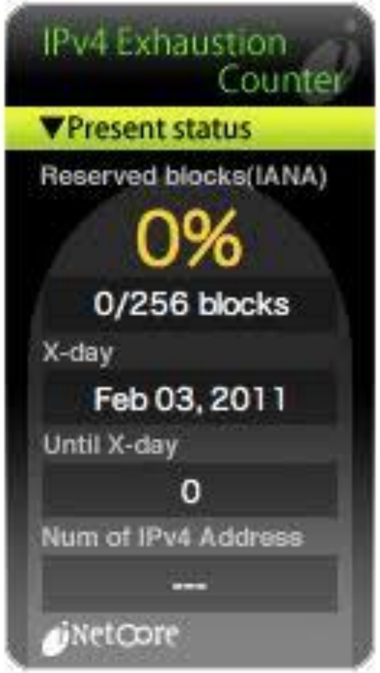

IPv4 Exhaustion Counter

▼Present status

Reserved blocks(IANA)

0%

0/256 blocks

X-day

Feb 03, 2011

Until X-day

0

Num of IPv4 Address

---

iNetCore

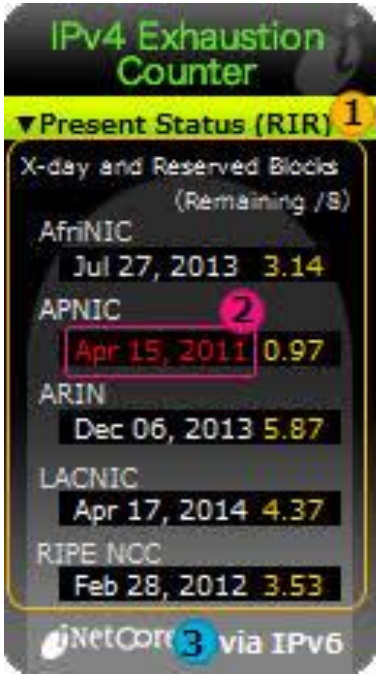

IPv4 Exhaustion Counter

▼Present Status (RIR)

X-day and Reserved Blocks (Remaining /8)

AfriNIC	Jul 27, 2013	3.14
APNIC	Apr 15, 2011	0.97
ARIN	Dec 06, 2013	5.87
LACNIC	Apr 17, 2014	4.37
RIPE NCC	Feb 28, 2012	3.53

iNetCore via IPv6

https://inetcore.com/project/ipv4ec/images/ipv4ec_iana_en.png

IPV4 ADDRESS REPORT

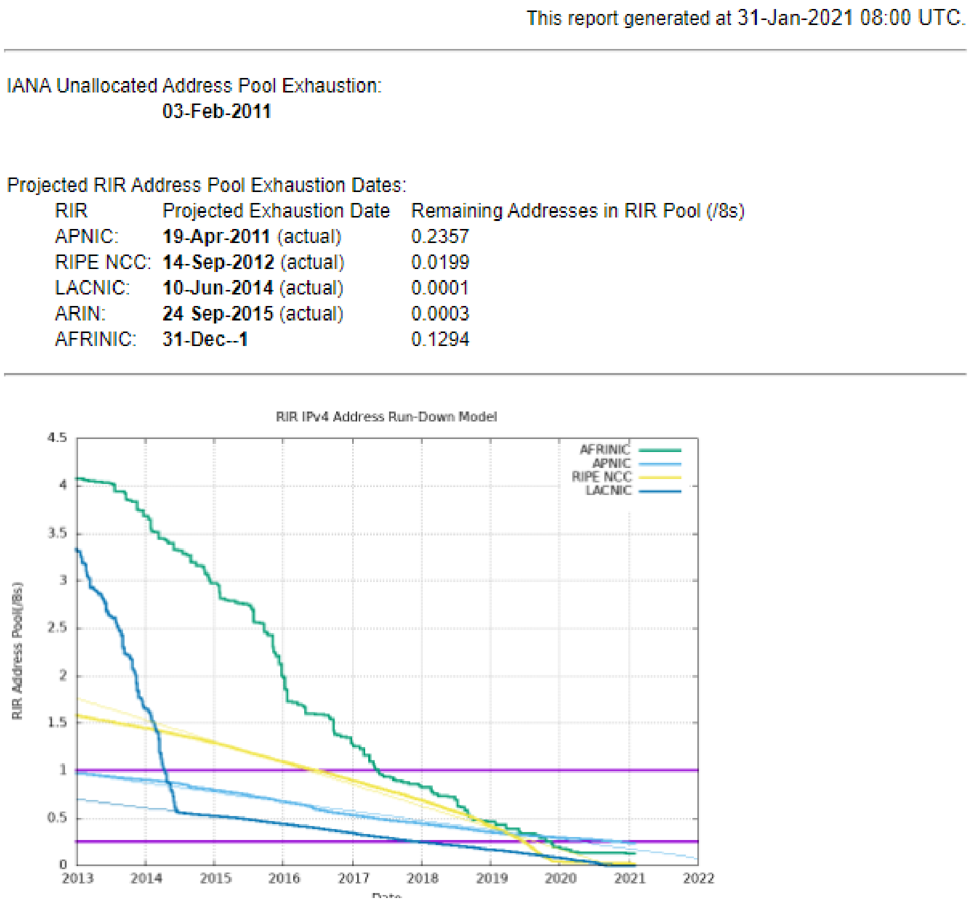

This report generated at 31-Jan-2021 08:00 UTC.

IANA Unallocated Address Pool Exhaustion:
03-Feb-2011

Projected RIR Address Pool Exhaustion Dates:

RIR	Projected Exhaustion Date	Remaining Addresses in RIR Pool (/8s)
APNIC:	**19-Apr-2011** (actual)	0.2357
RIPE NCC:	**14-Sep-2012** (actual)	0.0199
LACNIC:	**10-Jun-2014** (actual)	0.0001
ARIN:	**24 Sep-2015** (actual)	0.0003
AFRINIC:	**31-Dec--1**	0.1294

RIR IPv4 Address Run-Down Model

https://www.potaroo.net/tools/ipv4/index.html

WHAT IS IPV6?

Developed in the 1990s

IPv6 is…?

IPv1, v2, v3, v5.. IPv9?

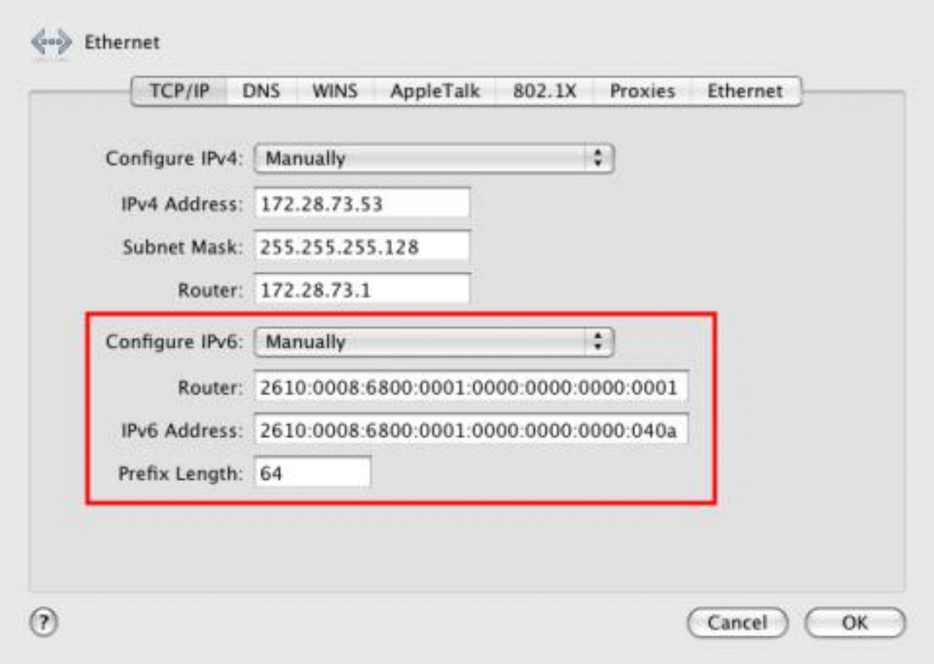

Configure IPv4: Manually

IPv4 Address: 172.28.73.53

Subnet Mask: 255.255.255.128

Router: 172.28.73.1

Configure IPv6: Manually

Router: 2610:0008:6800:0001:0000:0000:0000:0001

IPv6 Address: 2610:0008:6800:0001:0000:0000:0000:040a

Prefix Length: 64

Features	IPv4	IPv6
Size	32 bits	128 bits
Space	4,294,967,296	340,282,366,920,938,463,463 ,374,607,431,768,211,456
Notation	dotted decimal notation	hexadecimal with colons

GOALS IN DESIGNING IPV6

- Larger Address Space
- Better Management of Address Space
- Elimination of “Addressing Kludges”
- Easier TCP/IP Administration
- Modern Design for Routing
- Better Support for Multicasting
- Better Support for Security
- Better Support for Mobility

GOALS IN DESIGNING IPV6

	IPV4	IPv6
Address	32 bits (4 bytes) 12.34.56.78	128 bits (16 bytes) 1234:5678:9abc:def0
Packet size	576 bytes required fragmentation optional	80 bytes required without fragmentation
Packet fragmentation	Routers and sending host	Sending hosts only
Packet header	Does not Identify packet flow for QoS handling includes a checksum. Includes options up to 40 bytes	Contains Flow Label field that specifies packet flow for QoS handling. Does not Include a checksum. Extension headers used for optional data.

GOALS IN DESIGNING IPV6

	IPV4	IPv6
DNS records	Address (A) records. Maps host names. Pointer (PTR) records IN-ADDR.ARPA DNS domain.	Address (AAAA) records. Maps host names Pointer (PTR) records IP6.ARPA DNS domain.
Address Configuration	Manual or via DHCP	Stateless address auto configuration (SLAAC) using Internet Control Message Protocol version 6 (ICMPv6) or DHCPv6.
IP to MAC resolution	Broadcast ARP	Multicast Neighbor Solicitation
Local subnet group management	Internet Group Management Protocol (IGMP)	Multicast listener Discovery (MLD)

GOALS IN DESIGNING IPV6………

	IPV4	IPv6
Broadcast	Yes	No
Multicast	Yes	Yes
IPSec	Optional external	Required

COMMON MISCONCEPTIONS

- The introduction of IPv6 puts our current IP infrastructure—our networks and services—at risk.
- The IPv6 protocol is immature and hasn’t proven that it stands the test of time or whether it is capable of handling the requirements.
- The costs of introducing IPv6 are too high
- With Stateless Address Autoconfiguration, we will not be able to control or monitor network access
- Our Internet Service Provider (ISP) does not offer IPv6 services, so we can’t use it.
- It would be too expensive and complex to upgrade our backbone.”

COMMON MISCONCEPTIONS…….

- It would be too complex and expensive to port all of our applications to IPv6.
- We have enough IPv4 addresses; we don’t need IPv6.

CAN WE USE IPV6 NOW?

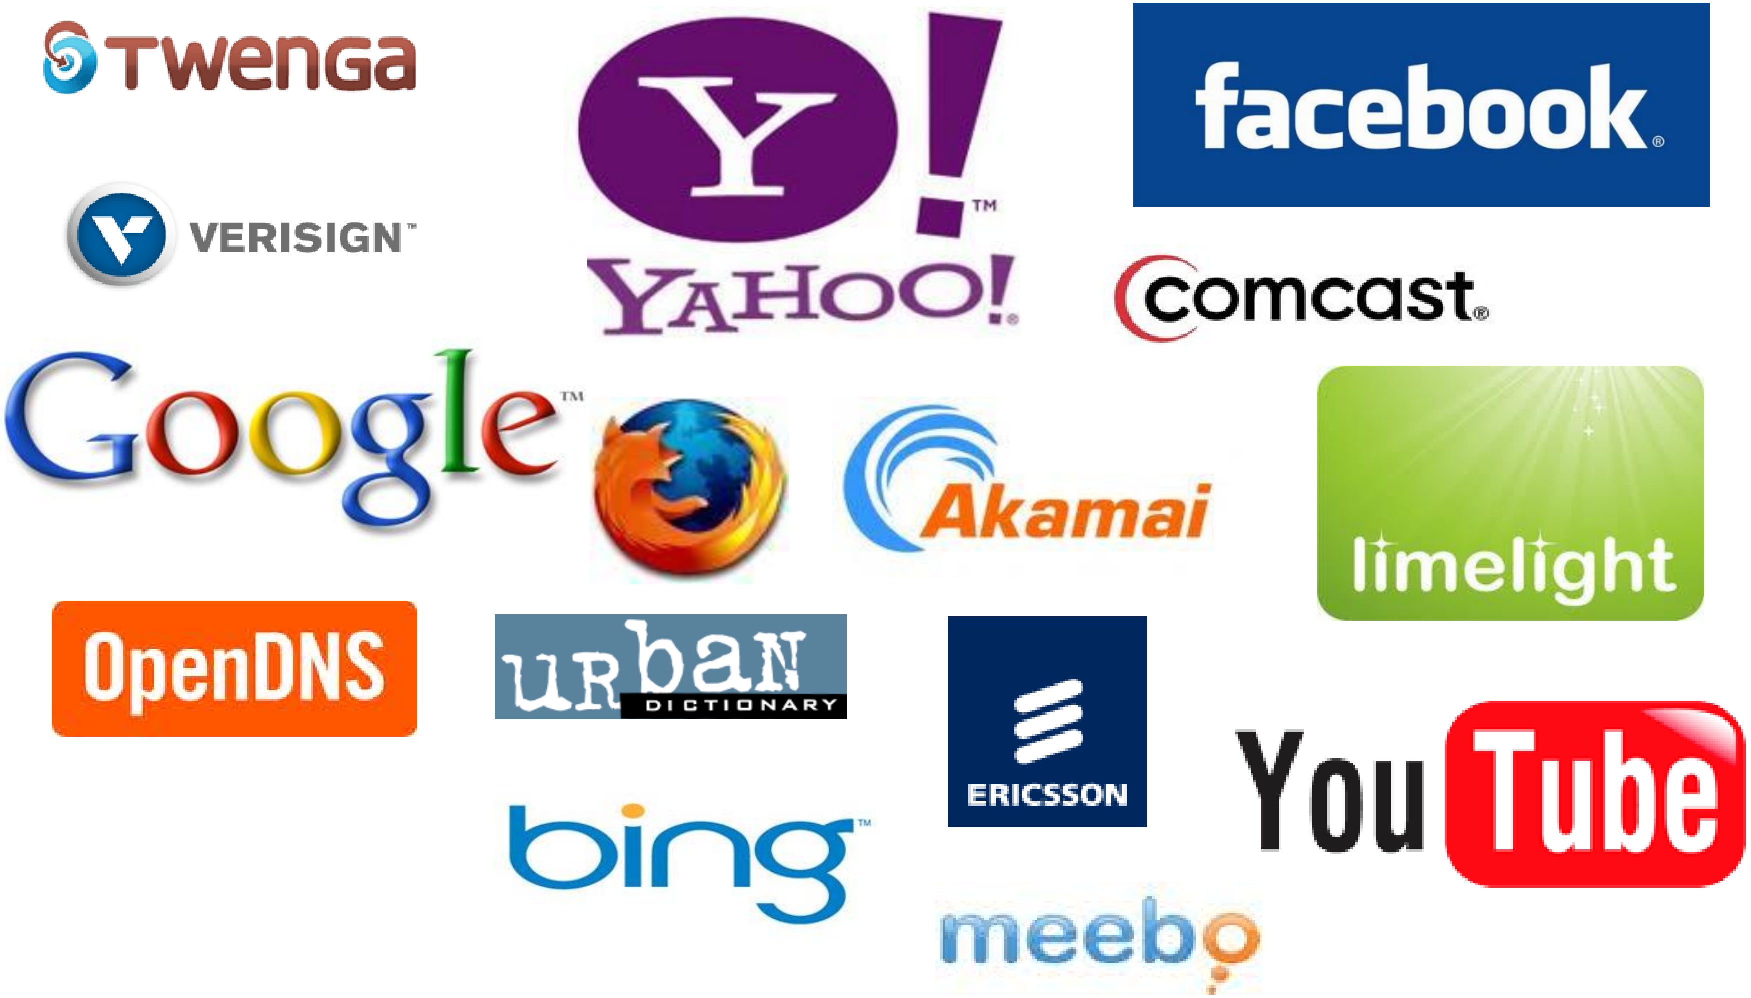

CAN WE USE IPV6 NOW?

CAN WE USE IPV6 NOW?

Source - http://www.sfc.wide.ad.jp/InternetCAR/about/more.html

InternetCAR

Kitchen Appliances

Botanics v6

FIVE STEPS ON THE PATH TO IPV6

- ❑ **Focus on IP address design and management.**
    - ❑ Start the IPv6 prefix assignment application process now. Stop worrying about conserving addresses and start thinking about adding meaning to individual hex digits.

- ❑ **Update network support systems**
    - ❑ Do you have an internal DNS infrastructure? Can nameservers support both IPv4 A and IPv6 AAAA records? If they're dual stacked, how do they respond to a name query when there are both IPv4 and IPv6 addresses assigned?

- ❑ **Budget for security updates and expertise**
    - ❑ End-to-end IPsec notwithstanding, security systems tend to be the problem children in IPv6 deployments. Not everything will survive the transition, so allocate some funds here.

FIVE STEPS ON THE PATH TO IPV6

- **Understand the lingo**
  - Tools for monitoring, logging, alarms, configuration management, and change management have to understand IPv6, not speak it.

- **Have end-to-end training**
  - Don't limit IPv6 education to IT. Going all-IPv6 positions your company as a technology leader. Make sure customer-facing personnel can tell the story.

PACKET STRUCTURE & HEADER FORMAT

UNDERSTANDING THE IPV6 HEADER

The IPv6 Header is designed to

Minimize the header overhead

Reduce the header process

NOTE* IPv4 & IPv6 headers are not interoperable

HEADER COMPARISON

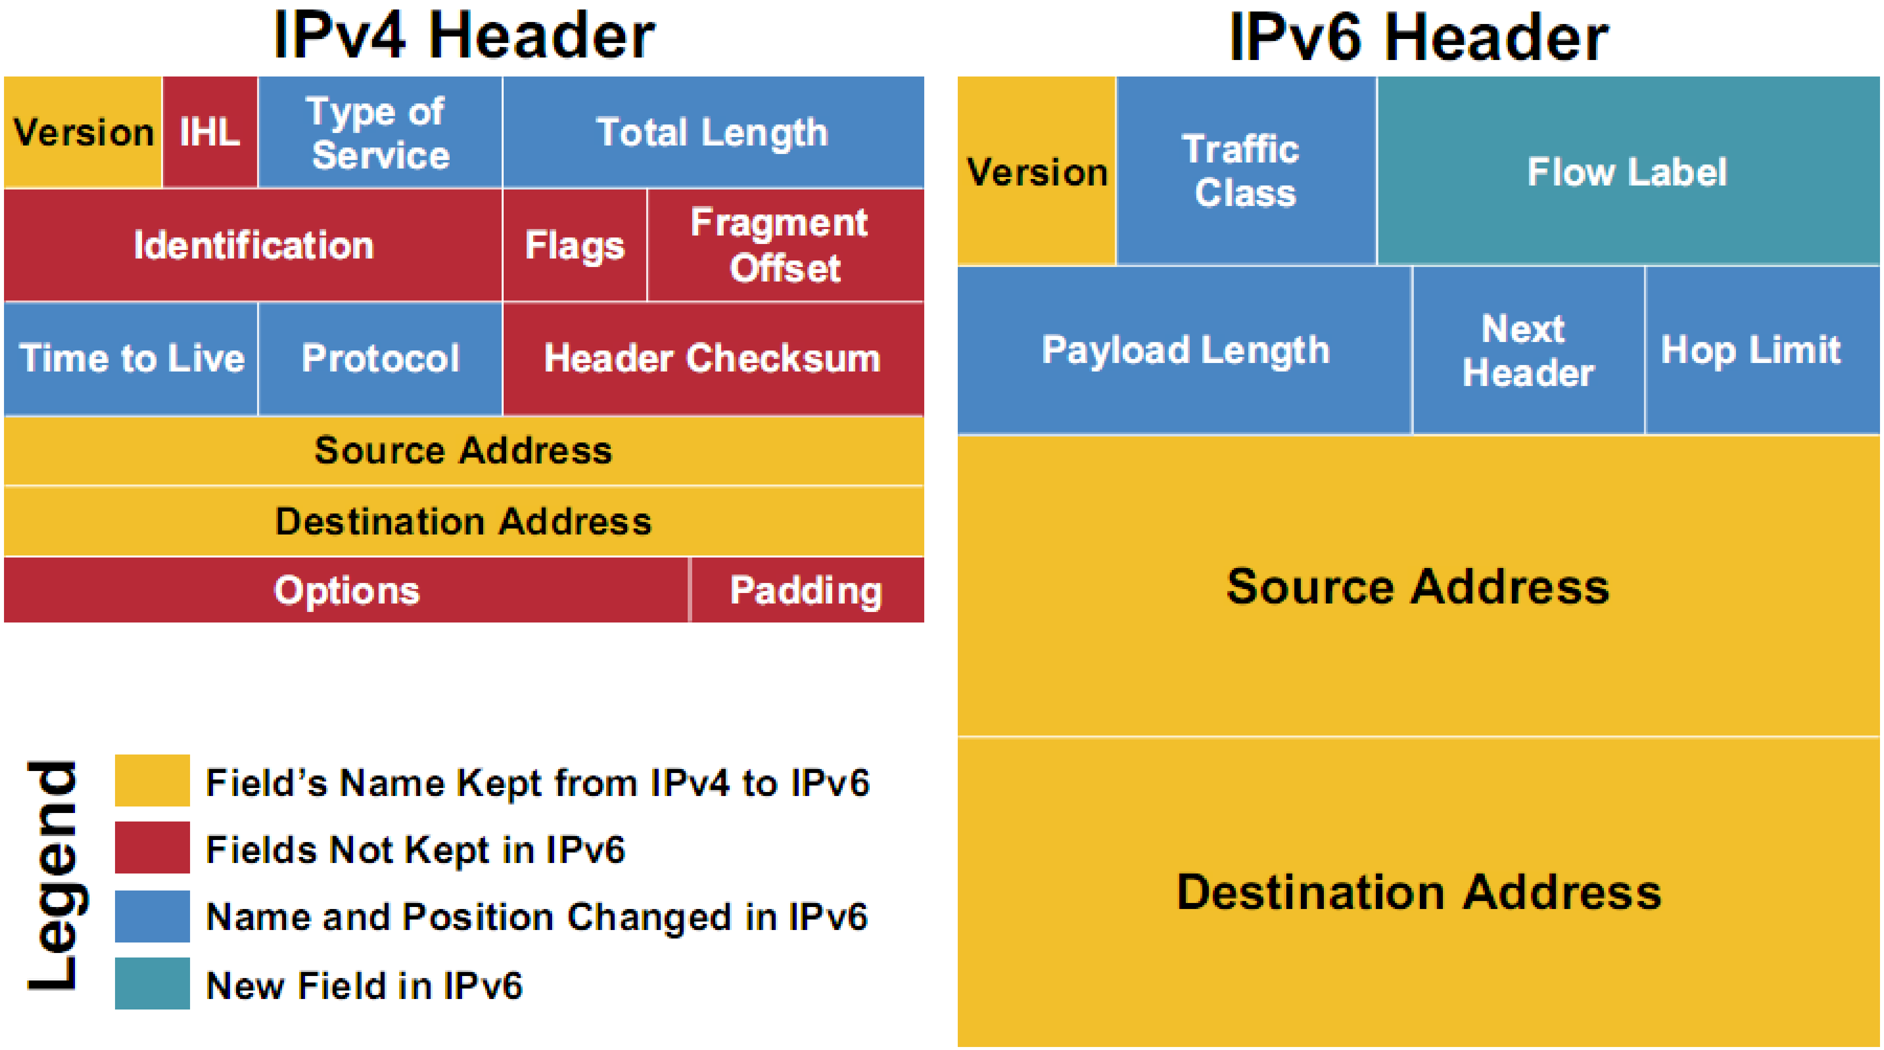

IPv4 Header

Version | IHL | Type of Service | Total Length

Identification | Flags | Fragment Offset

Time to Live | Protocol | Header Checksum

Source Address

Destination Address

Options | Padding

IPv6 Header

Version | Traffic Class | Flow Label

Payload Length | Next Header | Hop Limit

Source Address

Destination Address

Legend

- 🟨 Field's Name Kept from IPv4 to IPv6
- 🟥 Fields Not Kept in IPv6
- 🟦 Name and Position Changed in IPv6
- 🟩 New Field in IPv6

KEY CHANGES

- **Unchanged Fields**
  - Three fields are used the same way and ***retain*** the same name (though they have different content and/or size): **Version**, **Source Address** and **Destination Address**.
- **Renamed Fields**
  - Two fields are used the same way but ***renamed***: **Traffic Class** and **Hop Limit**.
- **Modified Fields**
  - Two fields are used in a way similar way in IPv4 but are slightly different in meaning and also renamed: **Payload Length** and **Next Header**.
- **Added Fields**
  - There is one new field: **Flow Label**.
- **Removed Fields**
  - To cut down on header length and unnecessary work, **five (5)** IPv4 header fields are removed from the IPv6 header

WHAT’S MISSING FROM V4

Options

- Moved to be separate headers (discussed later)

Fragmentation fields

- MTU discovery is a better approach
- available in ***Next Header***

Checksum

- Redundant with Layer-2 CRC

Length fields simplified

- No fragmentation
- no options

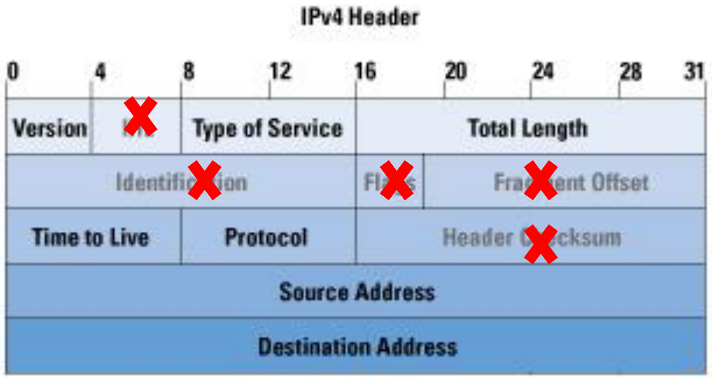

- Internet Header Length (IHL)
- Total Length

WHAT'S MISSING - EXPLAINED

- **Internet Header Length**
  No longer needed, as the main IPv6 header is **fixed in length at 40 bytes**.
- **Identification, Flags, Fragment Offset (*PMTU)**
  These are used for fragmentation, which is **done less** in IPv6 than IPv4, so these fields are now found **only when needed** in the '***Fragmentation***' extension header.

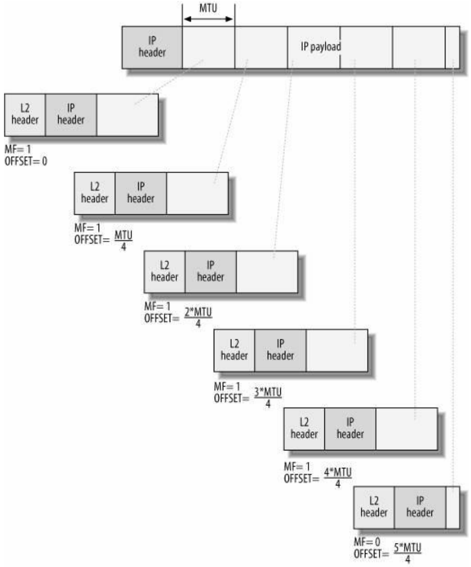

- **Header Checksum**
  It was viewed as redundant with higher-layer error-checking and data link layer CRC calculations.
  This **saves processing time for routers** and 2 bytes in the datagram header.

* IPv6 focuses on the PMTU process to eliminate the need for fragmentation

UNDERSTANDING EXTENSION HEADERS

In IPv6, optional Internet layer information is encoded in separate headers that may be placed **between** the **IPv6 header** and the **upper-layer header** in a packet.

Benefits of IPv6 Extension Headers

- **Extension headers are external to IPv6 header**
- **Routers do not look at these options except for Hop-by-hop options**
- **No negative impact on router‘s forwarding performance**
- **Easy to extend with new headers and option**

CATEGORIES OF EXTENSION HEADERS

Processed by EVERY hop

- Must be processed by **EVERY** node on the packet's path.
- Must always **appear immediately** after IPv6 header.
- Two Hop-by-hop options already defined:
  - **Router alert** option
  - **Jumbo payload** option

Processed ONLY by destination

- Meant to carry information intended to be **examined by the destination node.**

HOW EXTENSION HEADER WORKS

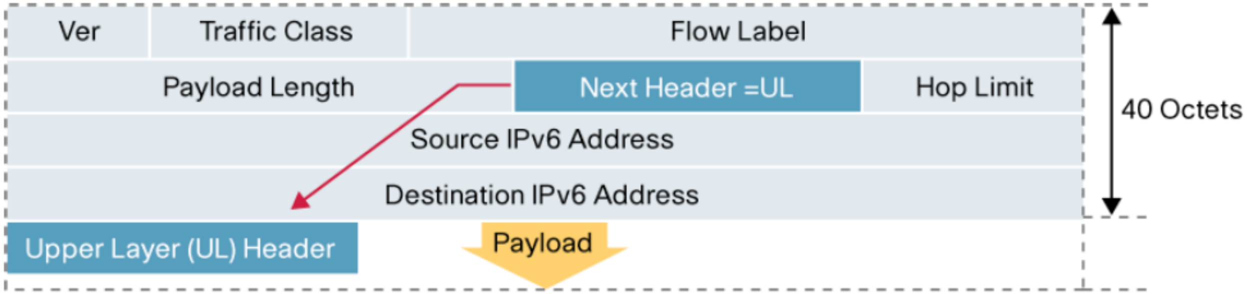

Before chaining Extension Headers

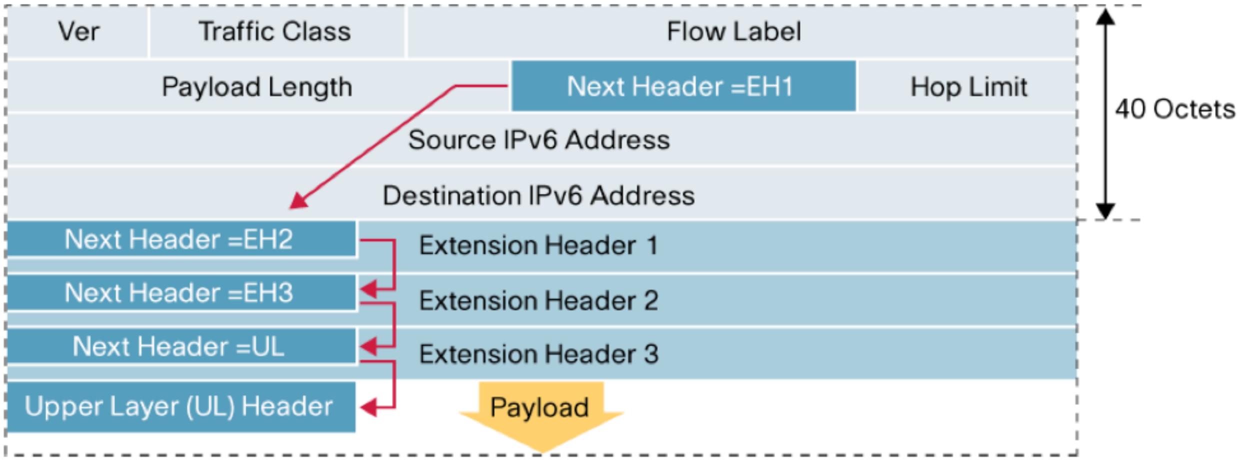

After Extension Headers in IPv6 Packets

EXTENSION HEADER PROCESSING

Extension headers are **NOT** *examined* or *processed* by any node along a packet' s delivery path.

ONLY hop-by-hop extension headers is processed by every node along a packet's delivery path. **(including source and destination)**

Hop-by-hop header (if present) **must** immediately follow IPv6 header

Extension headers **are processed** strictly **in order** they appear in the packet.

EXTENSION HEADER PROCESSING

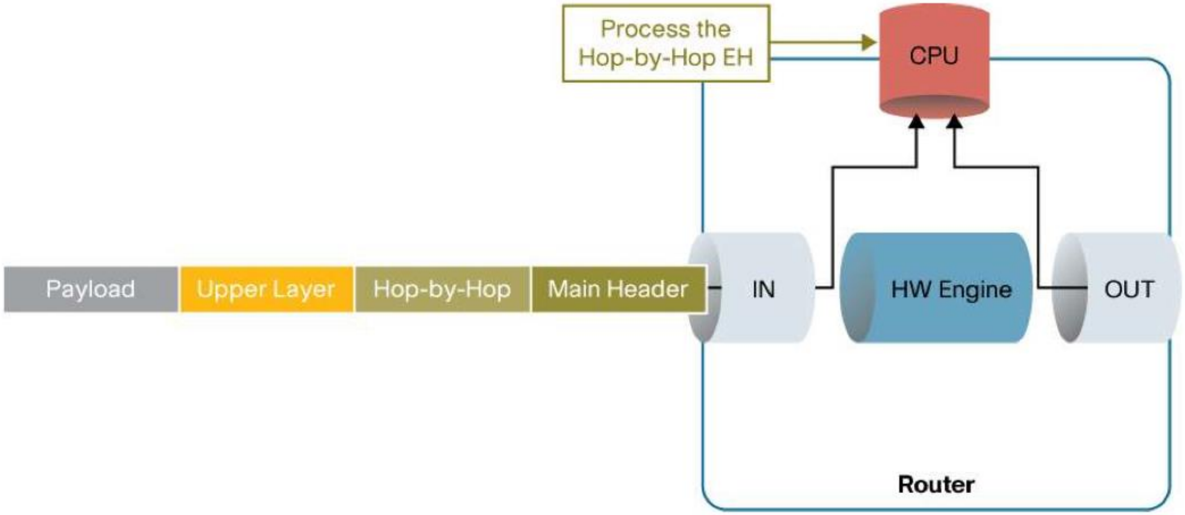

Forwarding IPv6 Packets with the Hop-by-Hop Extension Header

Forwarding IPv6 Packets with Extension Headers other than Hop-by-Hop in the Absence of ACLs

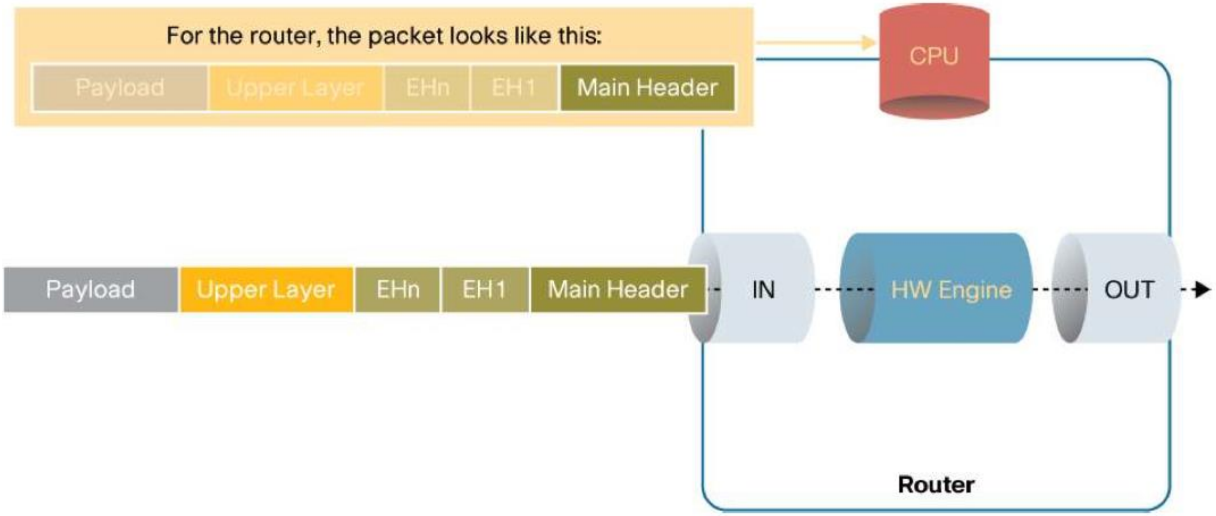

EXTENSION HEADER ORDER

RFC 2460 recommends following order:

1. IPv6 header (main header)
2. Hop-by-hop options header
3. Destination options header
4. Routing header
5. Fragment header
6. Authentication header
7. ESP header
8. Destination options header
9. Upper-layer header

EXTENSION HEADER CODES

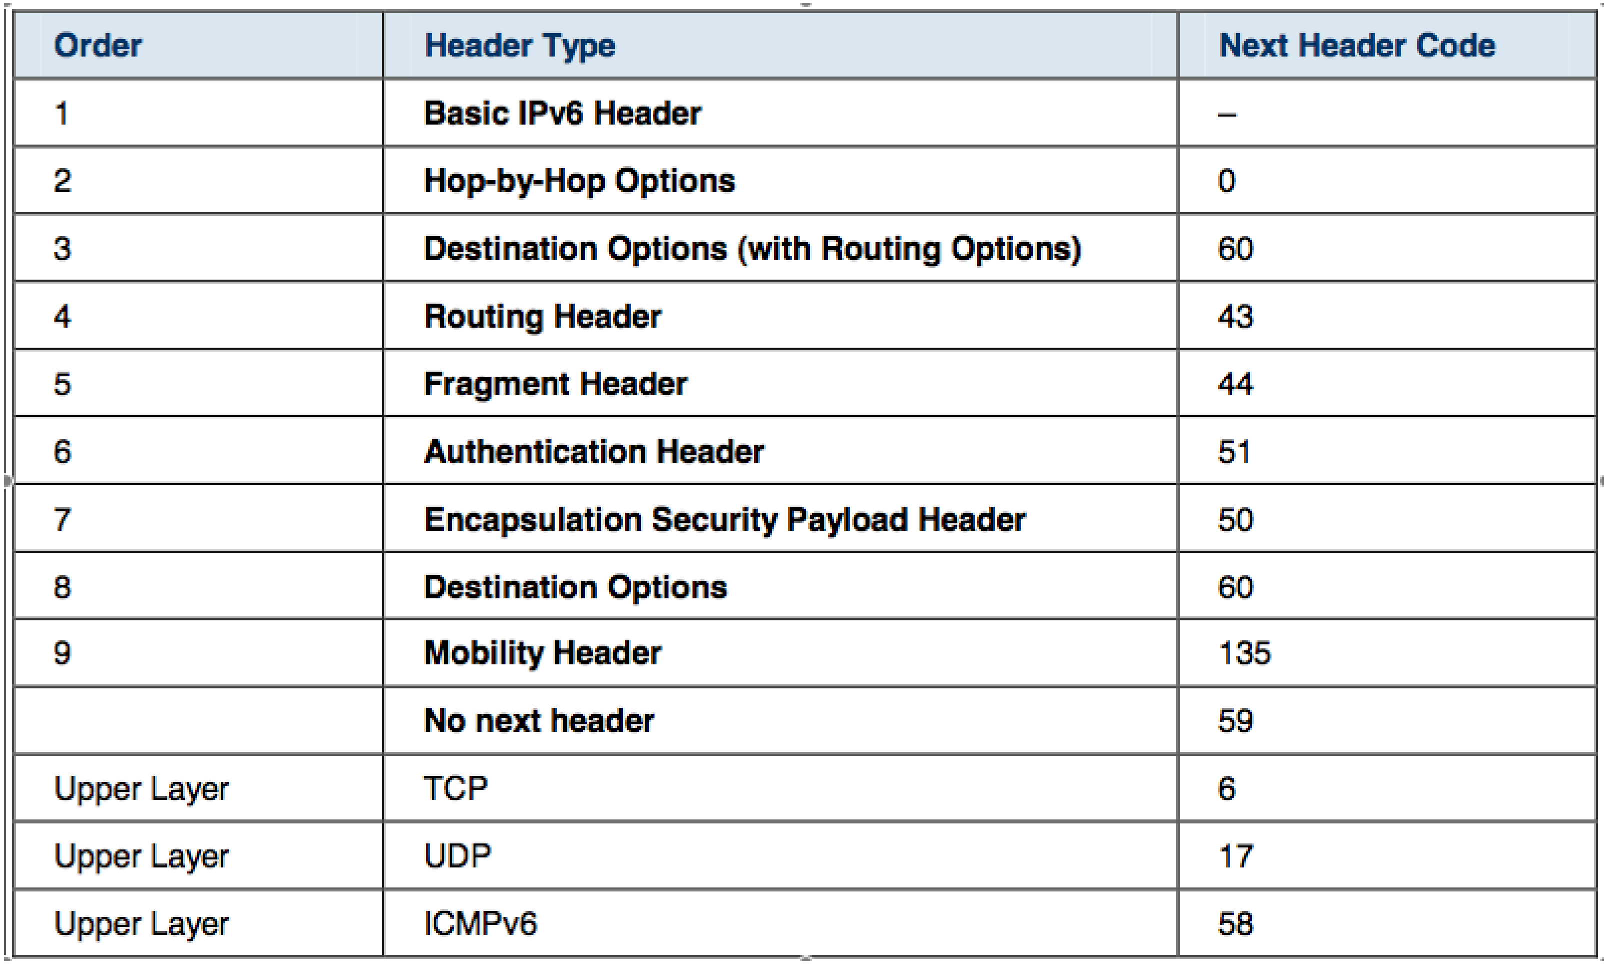

Order	Header Type	Next Header Code
1	**Basic IPv6 Header**	–
2	**Hop-by-Hop Options**	0
3	**Destination Options (with Routing Options)**	60
4	**Routing Header**	43
5	**Fragment Header**	44
6	**Authentication Header**	51
7	**Encapsulation Security Payload Header**	50
8	**Destination Options**	60
9	**Mobility Header**	135
	No next header	59
Upper Layer	TCP	6
Upper Layer	UDP	17
Upper Layer	ICMPv6	58

E.G. ROUTING HEADER

Next header value: 43

Provides "source-routing" functionality

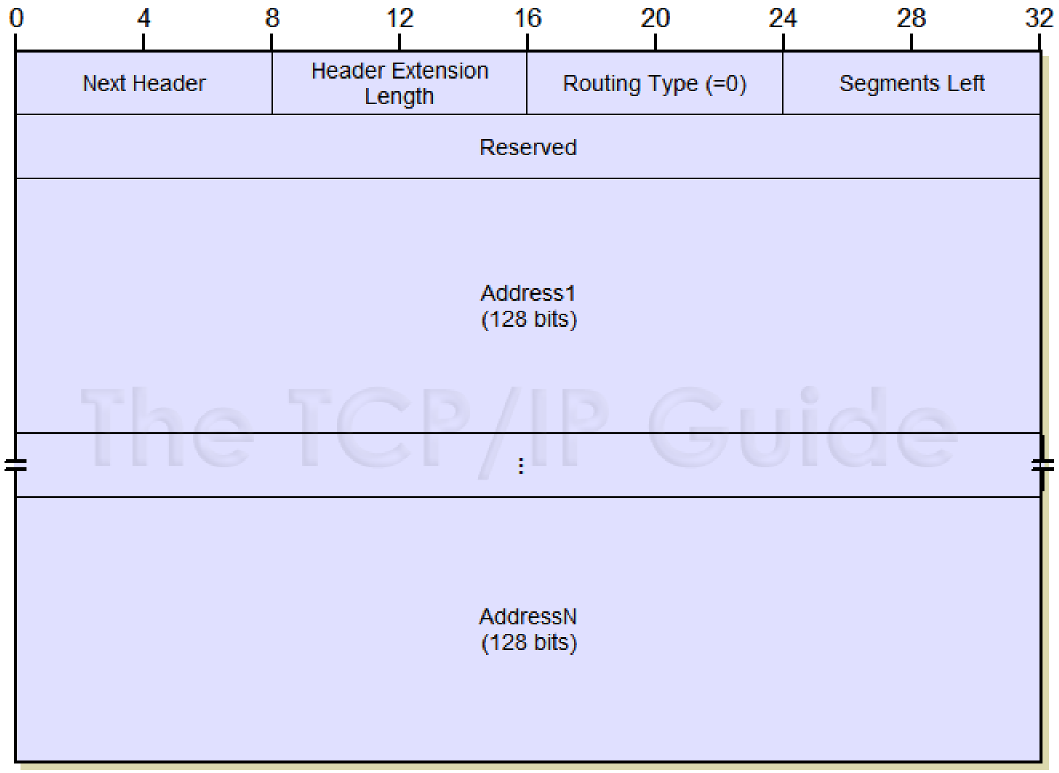

Next Header: Contains the protocol number of the next header after the Routing header. Used to link headers together as described above.

Header Extension Length: The length of the Routing header in 8-byte units, not including the first 8 bytes of the header. For a Routing Type of 0, this value is thus two times the number addresses embedded in the header.

Routing Type: This field allows multiple routing types to be defined; at present, the only value used is 0.

Segments Left: Specifies the number of explicitly-named nodes remaining in the route until the destination**.**

Reserved: Not used; set to zeroes.

Addresses: A set of IPv6 addresses that specify the route to be used.

E.G. FRAGMENT HEADER

Next header value: 44

Used to provide datagram fragmentation

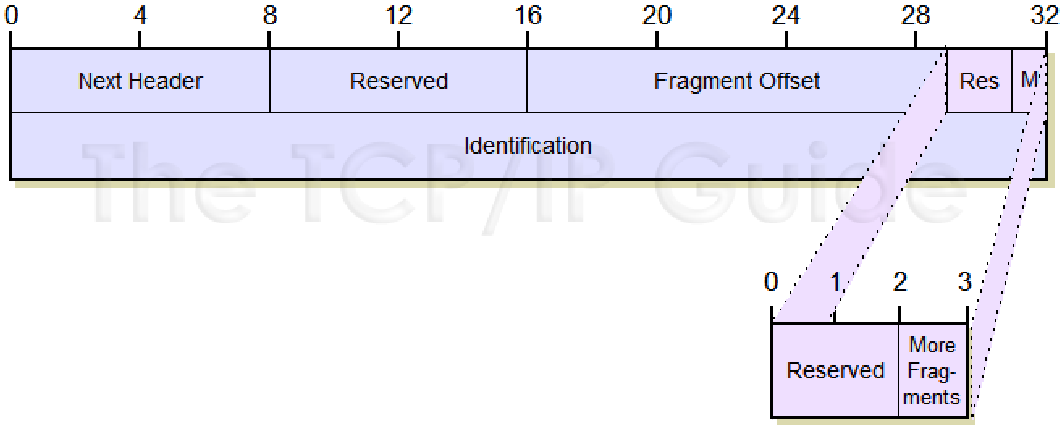

Next Header: Contains the protocol number of the next header after the Fragment header. Used to link headers together as described above**.**

Reserved: Not used; set to zeroes.

Fragment Offset: Specifies the offset, or position, in the overall message where the data in this fragment goes. It is specified in units of 8 bytes (64 bits) and used in a manner very similar to the field of the same name in the IPv4 header.

(Res) Reserved: Not used; set to zeroes.

More Fragments Flag: Same as the flag of the same name in the IPv4 header—when set to 0, indicates the last fragment in a message; when set to 1, indicates that more fragments are yet to come in the fragmented message.

Identification: Same as the field of the same name in the IPv4 header, but expanded to 32 bits. It contains a specific value that is common to each of the fragments belonging to a particular message, to ensure that pieces from different fragmented messages are not mixed together.

EXTENSION HEADER IN ACTION

- In this example, the packet is sent from **Mobile Node A** to **Mobile Node B** over the route optimized path [RFC3775], hence the use of the **Routing EH (43)** and the **Destination Options EH (60)**. It is sent over a path that has an Maximum Transmission Unit (MTU) smaller than that of Mobile Nodes (MNs) access link, hence the use of the **Fragmentation EH (44)**.

EXTENSION HEADER IN ACTION

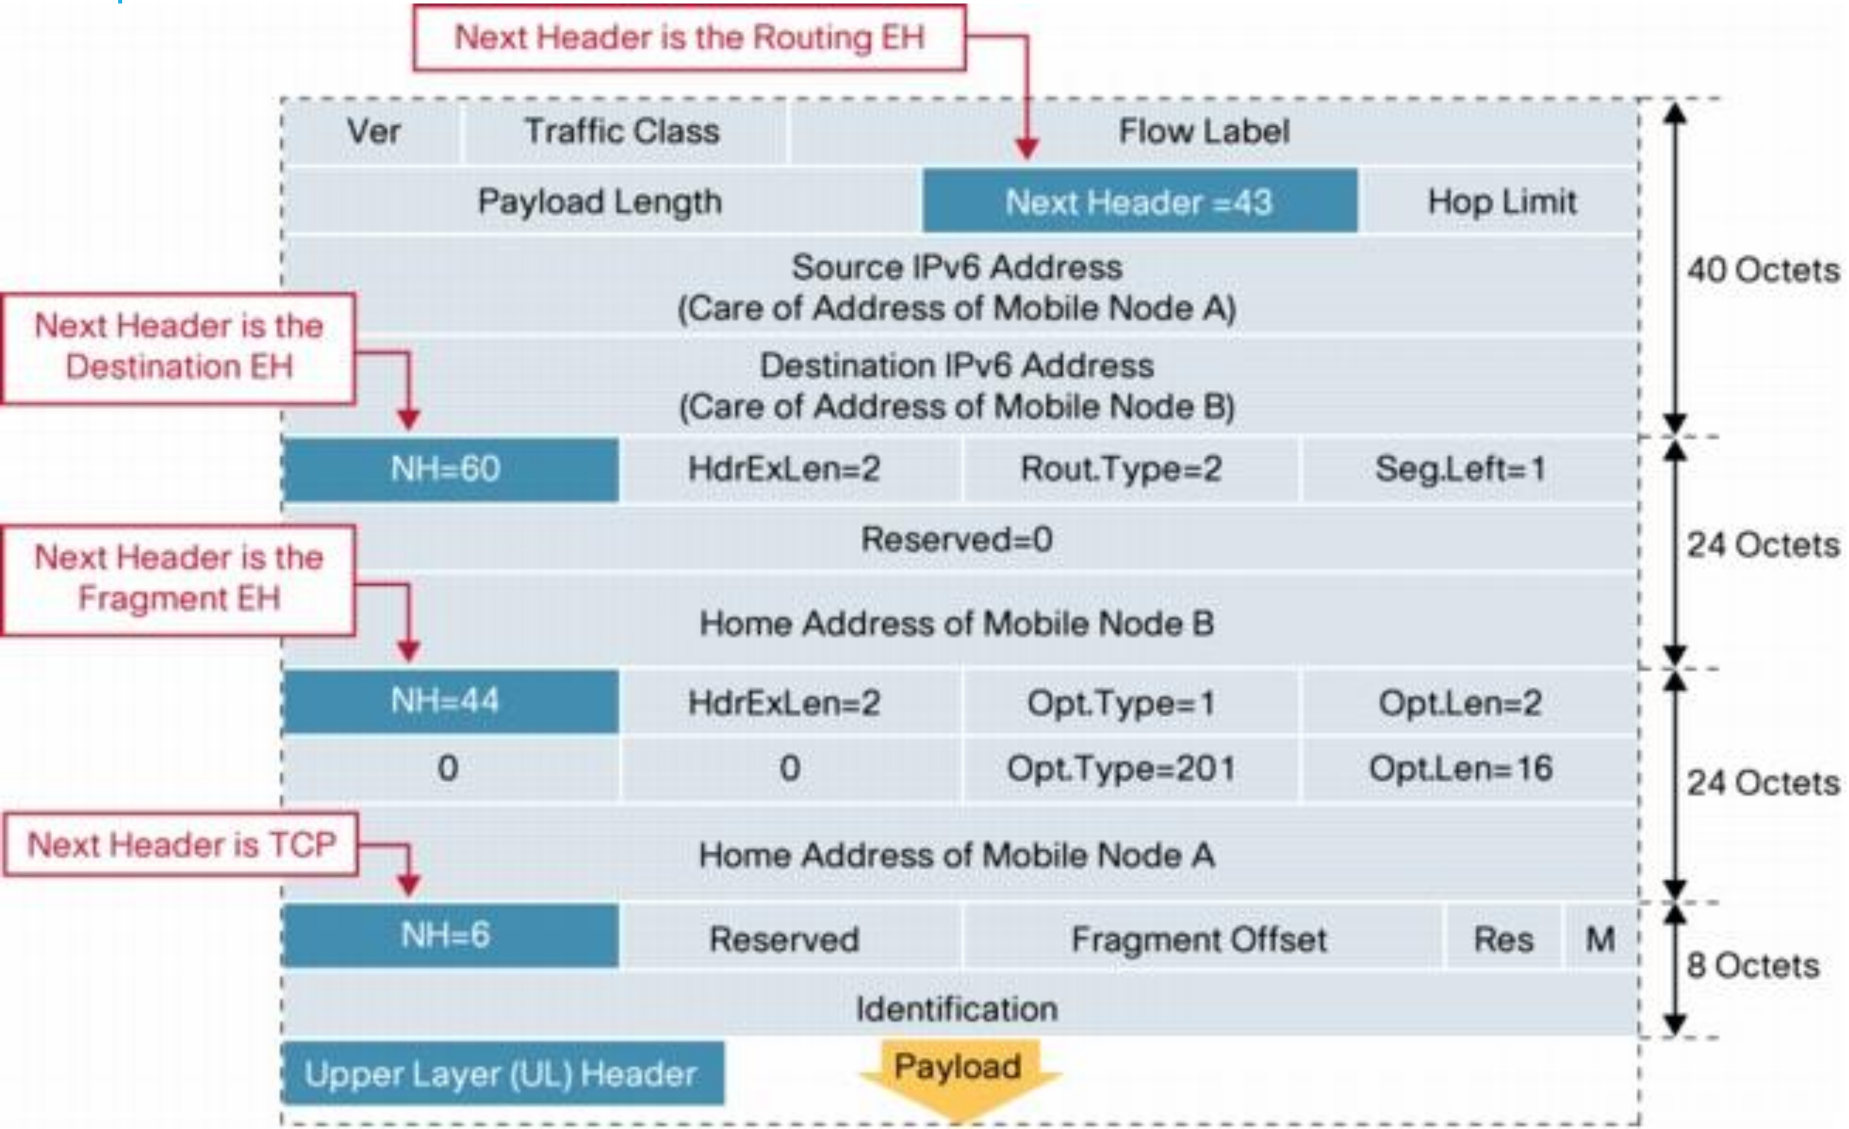

Next Header is the Routing EH

Ver	Traffic Class	Flow Label	
Payload Length		Next Header =43	Hop Limit
Source IPv6 Address (Care of Address of Mobile Node A)			
Destination IPv6 Address (Care of Address of Mobile Node B)			

40 Octets

Next Header is the Destination EH

NH=60	HdrExLen=2	Rout.Type=2	Seg.Left=1
Reserved=0			
Home Address of Mobile Node B			

24 Octets

Next Header is the Fragment EH

NH=44	HdrExLen=2	Opt.Type=1	Opt.Len=2
0	0	Opt.Type=201	Opt.Len=16
Home Address of Mobile Node A			

24 Octets

Next Header is TCP

NH=6	Reserved	Fragment Offset	Res	M
Identification				

8 Octets

Upper Layer (UL) Header

Payload

IPV6 ADDRESSING ARCHITECTURE

ADDRESSING MODEL

IPv6 addresses of all types are assigned to interfaces, **not nodes**.

An IPv6 unicast address **refers to a single interface**.
Since each interface **belon**
that node's interfaces' uni
an identifier for the node.

Source=http://tools.ietf.org/html/rfc4291

ADDRESSING MODEL

A unicast address or **a set of unicast addresses** may be assigned to multiple physical interfaces if the implementation treats the multiple physical interfaces as one interface when presenting it to the internet layer.

This is useful for **load-sharing** over multiple physical interfaces.

IPv6 Address = **2001:db8:a:b::c/128**

Source=http://tools.ietf.org/html/rfc4291

IPV6 TEXT REPRESENTATION (1)

There are three conventional forms for representing IPv6 addresses as text strings:

- The preferred form is **x:x:x:x:x:x:x:x**, where the 'x's are **one to four hexadecimal** digits of the eight 16-bit pieces of the address.

Examples:
ABCD:EF01:2345:6789:ABCD:EF01:2345:6789
2001:DB8:0:0:8:800:200C:417A

Note* that it is not necessary to write the leading zeros in an individual field, but there must be at least one numeral in every field

Source=http://tools.ietf.org/html/rfc4291

IPV6 TEXT REPRESENTATION (2)

- In order to make writing **addresses containing zero bits** easier, a special syntax is available to compress the zeros.
- The use of **"::"** indicates **one or more** groups of 16 bits of zeros.
- The **"::"** can ***only appear once*** in an address. The **"::"** can also be used to compress **leading** or **trailing** zeros in an address.

For example, the following addresses

Before compression	After Compression
2001:DB8:0:0:8:800:200C:417A	2001:DB8::8:800:200C:417A
FF01:0:0:0:0:0:0:101	FF01::101
0:0:0:0:0:0:0:1	::1

Source=http://tools.ietf.org/html/rfc4291

IPV6 TEXT REPRESENTATION (3)

An alternative form that is sometimes more convenient when dealing with a ***mixed environment of IPv4 and IPv6*** nodes is:

x:x:x:x:x:x:d.d.d.d

where the 'x's are the **hexadecimal** values of the six high-order 16-bit pieces of the address

and the 'd's are the **decimal** values of the four low-order 8-bit pieces of the address (standard IPv4 representation)

Examples:

0:0:0:0:0:0:13.1.68.3
0:0:0:0:0:FFFF:129.144.52.38

or in compressed form:

::13.1.68.3
::FFFF:129.144.52.38

Source=http://tools.ietf.org/html/rfc4291

TEXT REPRESENTATION OF ADDRESS PREFIXES

- The text representation of IPv6 **address prefixes** is similar to the way IPv4 address prefixes are written in **Classless Inter-Domain Routing** (CIDR) notation.
- An IPv6 address prefix is represented by the notation:

ipv6-address/**prefix-length**

ipv6-address	is an IPv6 address in any of the notations listed earlier
prefix-length	is a decimal value specifying how many of the leftmost contiguous bits of the address comprise the prefix.

For example, the following are legal representations of the 60-bit prefix:

2001:0DB8:0000:CD30:0000:0000:0000:0000/60
2001:0DB8::CD30:0:0:0:0/60
2001:0DB8:0:CD30::/60

Source=http://tools.ietf.org/html/rfc4291

RECOMMENDATION FOR IPV6 ADDRESS TEXT REPRESENTATION

THE PROBLEM

- ❑ **Leading Zeros in a 16-Bit Field**
    - ❑ It is not necessary to **write OR omit** the leading zeros in an individual field.
- ❑ **Zero Compression**
    - ❑ The use of "::" indicates one or more groups of 16 bits of zeros. It is possible to select **whether OR not** to omit just one 16-bit 0 field.
- ❑ **Uppercase or Lowercase**
    - ❑ [RFC4291] does **not** mention any preference of uppercase or lowercase.

Source=http://tools.ietf.org/html/rfc5952

RECOMMENDATION FOR IPV6 ADDRESS TEXT REPRESENTATION

THE PROPOSED SOLUTION

- ❑ **Handling Leading Zeros in a 16-Bit Field**
- ❑ Leading zeros **MUST** be suppressed
- ❑ **"::" Usage**
  - ❑ "::" **MUST** be used to its maximum capability.
  - ❑ "::" **MUST NOT** be used to shorten just one 16-bit 0 field.
  - ❑ The longest run of consecutive 16-bit 0 fields MUST be shortened.
- ❑ **Lowercase**
  - ❑ The characters "a", "b", "c", "d", "e", and "f" in an IPv6 address **MUST** be represented in lowercase.

Source=http://tools.ietf.org/html/rfc5952

RECOMMENDATION FOR IPV6 ADDRESS TEXT REPRESENTATION

Combining IPv6 Addresses with Port Numbers

There are many different ways to combine **IPv6 addresses and port numbers** that are represented in text. Examples are shown below.

- [2001:db8::1]:80
- 2001:db8::1:80 **(NOT RECOMENDED)**
- 2001:db8::1.80
- 2001:db8::1 port 80
- 2001:db8::1p80
- 2001:db8::1#80

Source=http://tools.ietf.org/html/rfc5952

IPV6 ADDRESS TYPES (RFC 4291)

Unicast

- An identifier for a **single interface**. A packet sent to a unicast address is *delivered to the interface identified by that address*.

Anycast

- 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 protocols' measure of distance).

Multicast

- An identifier for a **set of interfaces** (typically belonging to different nodes). A packet sent to a multicast address is *delivered to all interfaces identified by that address*.

IPv6 has no broadcast address.

ADDRESS TYPES

The type of an IPv6 address is identified by the high-order bits of the address, as follows:

Address type	Binary prefix	IPv6 notation
Unspecified	00...0 (128 bits)	::/128
Loopback	00...1 (128 bits)	::1/128
Multicast	11111111	FF00::/8
Link-Local unicast	1111111010	FE80::/10
Global Unicast	(everything else)	

Note* Anycast addresses are taken from the **unicast** address spaces (of any scope) and are **not syntactically distinguishable from unicast addresses**.

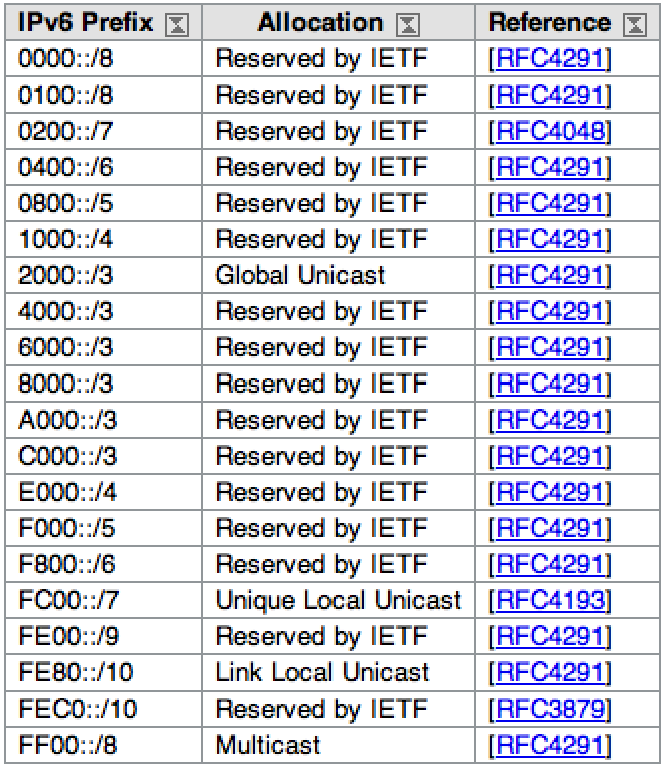

IPv6 Prefix	Allocation	Reference
0000::/8	Reserved by IETF	[RFC4291]
0100::/8	Reserved by IETF	[RFC4291]
0200::/7	Reserved by IETF	[RFC4048]
0400::/6	Reserved by IETF	[RFC4291]
0800::/5	Reserved by IETF	[RFC4291]
1000::/4	Reserved by IETF	[RFC4291]
2000::/3	Global Unicast	[RFC4291]
4000::/3	Reserved by IETF	[RFC4291]
6000::/3	Reserved by IETF	[RFC4291]
8000::/3	Reserved by IETF	[RFC4291]
A000::/3	Reserved by IETF	[RFC4291]
C000::/3	Reserved by IETF	[RFC4291]
E000::/4	Reserved by IETF	[RFC4291]
F000::/5	Reserved by IETF	[RFC4291]
F800::/6	Reserved by IETF	[RFC4291]
FC00::/7	Unique Local Unicast	[RFC4193]
FE00::/9	Reserved by IETF	[RFC4291]
FE80::/10	Link Local Unicast	[RFC4291]
FEC0::/10	Reserved by IETF	[RFC3879]
FF00::/8	Multicast	[RFC4291]

Source=http://tools.ietf.org/html/rfc4291

ADDRESS SCOPE

- The notion of scope allows IPv6 address to be **limited to specific spheres of influence**.
- As the **Scope ID** value increases, the scope expands to cover the local network, site, organization, and finally, the entire Internet.

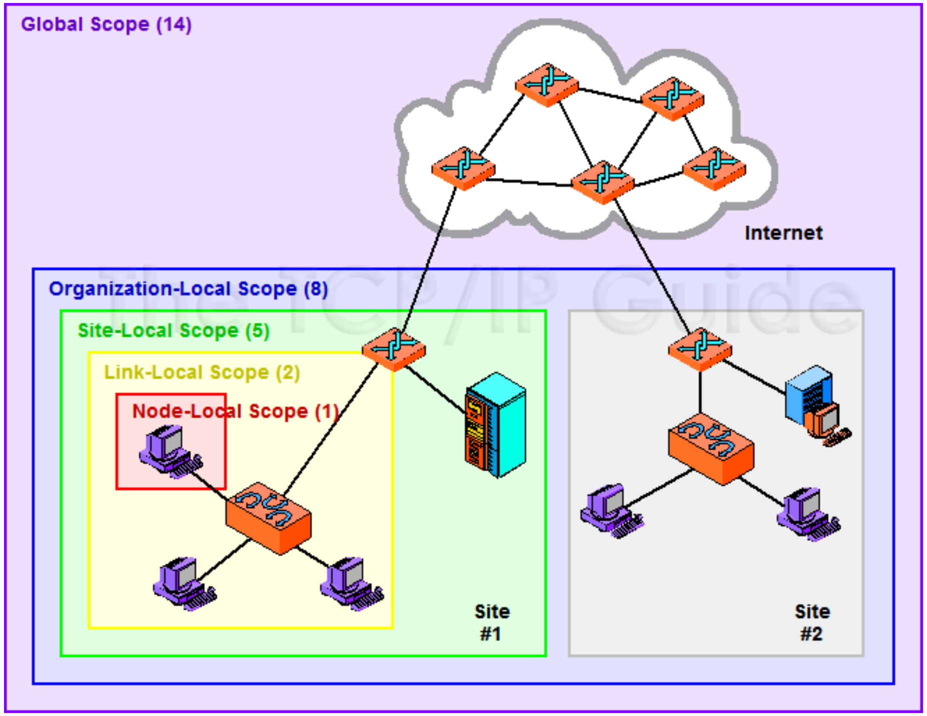

ADDRESS SCOPE

Address Scope/ Reachability	Description
Node-local addresses to reach same node	Used to send protocol data units (PDUs) to the **same node**: • Loopback address • Node-local multicast address
Link-local addresses to reach local link	Used to communicate between host devices (e.g., servers, VoIP devices, etc.) **on the link**; **these addresses are always configured automatically** • Unspecified address • Link-local unicast address • Link-local multicast address
Site-local addresses to reach the private intranet (internetwork)	Used between nodes that communicate with other nodes in the **same site**. • Site-local Unicast address • Site-local Multicast address
Global addresses to reach the Internet **a.k.a aggregatable global unicast addresses**	Globally **routable** and **reachable** addresses on the IPv6 portion of the Internet. • Global Unicast address • Other scope Multicast address

Source=Handbook of IPv4 to IPv6 Transition

UNDERSTANDING ADDRESS TYPES

IPv6 Address Type 1

IPv6 Unicast Address

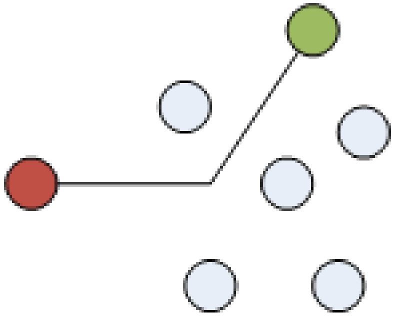

UNICAST ADDRESS

- There are **several types of unicast** addresses in IPv6:
    - Global Unicast
    - Site-local unicast (deprecated)
    - Link-Local unicast
- There are also some special-purpose subtypes of Global Unicast, such as **IPv6 addresses with embedded IPv4 addresses**.
- Additional address types or subtypes can be defined in the future.

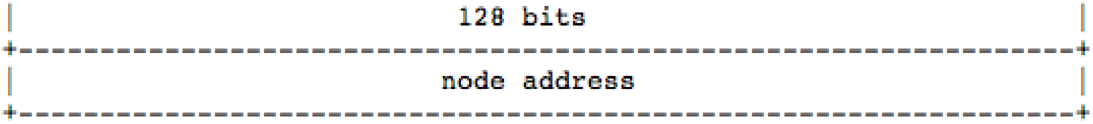

```
|                             128 bits                              |
+-------------------------------------------------------------------+
|                           node address                            |
+-------------------------------------------------------------------+
```

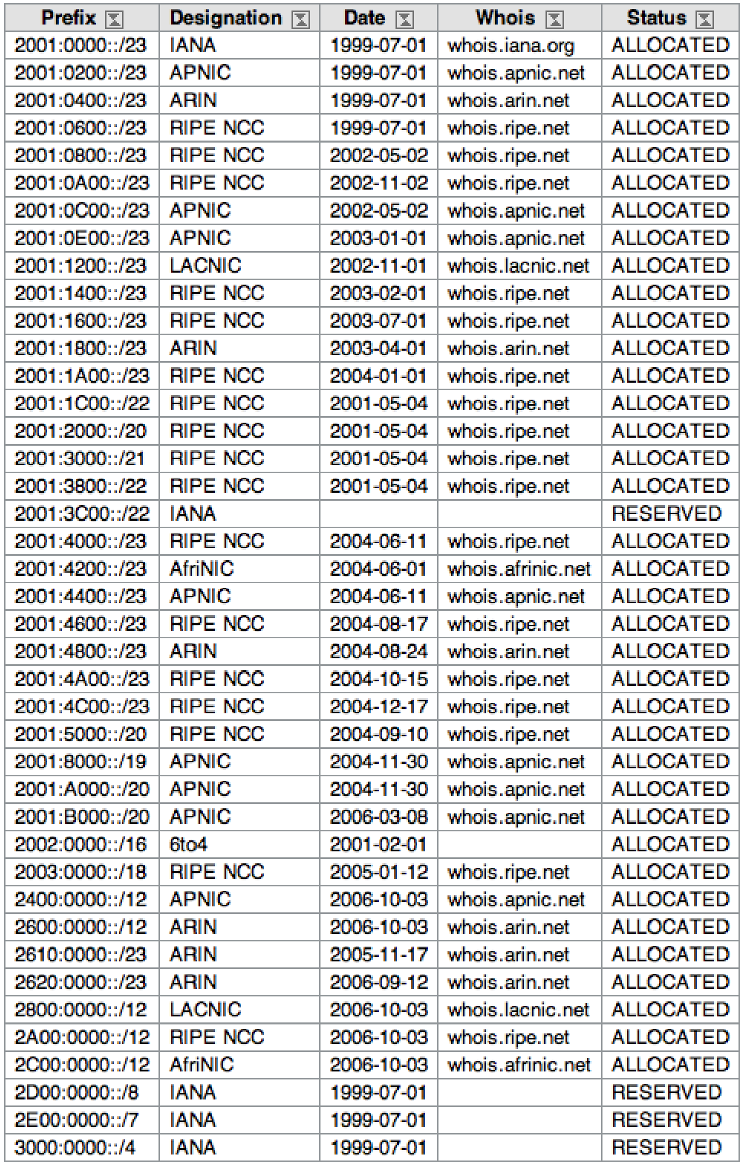

Prefix ☒	Designation ☒	Date ☒	Whois ☒	Status ☒
2001:0000::/23	IANA	1999-07-01	whois.iana.org	ALLOCATED
2001:0200::/23	APNIC	1999-07-01	whois.apnic.net	ALLOCATED
2001:0400::/23	ARIN	1999-07-01	whois.arin.net	ALLOCATED
2001:0600::/23	RIPE NCC	1999-07-01	whois.ripe.net	ALLOCATED
2001:0800::/23	RIPE NCC	2002-05-02	whois.ripe.net	ALLOCATED
2001:0A00::/23	RIPE NCC	2002-11-02	whois.ripe.net	ALLOCATED
2001:0C00::/23	APNIC	2002-05-02	whois.apnic.net	ALLOCATED
2001:0E00::/23	APNIC	2003-01-01	whois.apnic.net	ALLOCATED
2001:1200::/23	LACNIC	2002-11-01	whois.lacnic.net	ALLOCATED
2001:1400::/23	RIPE NCC	2003-02-01	whois.ripe.net	ALLOCATED
2001:1600::/23	RIPE NCC	2003-07-01	whois.ripe.net	ALLOCATED
2001:1800::/23	ARIN	2003-04-01	whois.arin.net	ALLOCATED
2001:1A00::/23	RIPE NCC	2004-01-01	whois.ripe.net	ALLOCATED
2001:1C00::/22	RIPE NCC	2001-05-04	whois.ripe.net	ALLOCATED
2001:2000::/20	RIPE NCC	2001-05-04	whois.ripe.net	ALLOCATED
2001:3000::/21	RIPE NCC	2001-05-04	whois.ripe.net	ALLOCATED
2001:3800::/22	RIPE NCC	2001-05-04	whois.ripe.net	ALLOCATED
2001:3C00::/22	IANA			RESERVED
2001:4000::/23	RIPE NCC	2004-06-11	whois.ripe.net	ALLOCATED
2001:4200::/23	AfriNIC	2004-06-01	whois.afrinic.net	ALLOCATED
2001:4400::/23	APNIC	2004-06-11	whois.apnic.net	ALLOCATED
2001:4600::/23	RIPE NCC	2004-08-17	whois.ripe.net	ALLOCATED
2001:4800::/23	ARIN	2004-08-24	whois.arin.net	ALLOCATED
2001:4A00::/23	RIPE NCC	2004-10-15	whois.ripe.net	ALLOCATED
2001:4C00::/23	RIPE NCC	2004-12-17	whois.ripe.net	ALLOCATED
2001:5000::/20	RIPE NCC	2004-09-10	whois.ripe.net	ALLOCATED
2001:8000::/19	APNIC	2004-11-30	whois.apnic.net	ALLOCATED
2001:A000::/20	APNIC	2004-11-30	whois.apnic.net	ALLOCATED
2001:B000::/20	APNIC	2006-03-08	whois.apnic.net	ALLOCATED
2002:0000::/16	6to4	2001-02-01		ALLOCATED
2003:0000::/18	RIPE NCC	2005-01-12	whois.ripe.net	ALLOCATED
2400:0000::/12	APNIC	2006-10-03	whois.apnic.net	ALLOCATED
2600:0000::/12	ARIN	2006-10-03	whois.arin.net	ALLOCATED
2610:0000::/23	ARIN	2005-11-17	whois.arin.net	ALLOCATED
2620:0000::/23	ARIN	2006-09-12	whois.arin.net	ALLOCATED
2800:0000::/12	LACNIC	2006-10-03	whois.lacnic.net	ALLOCATED
2A00:0000::/12	RIPE NCC	2006-10-03	whois.ripe.net	ALLOCATED
2C00:0000::/12	AfriNIC	2006-10-03	whois.afrinic.net	ALLOCATED
2D00:0000::/8	IANA	1999-07-01		RESERVED
2E00:0000::/7	IANA	1999-07-01		RESERVED
3000:0000::/4	IANA	1999-07-01		RESERVED

GLOBAL UNICAST ADDRESS EXPLAINED

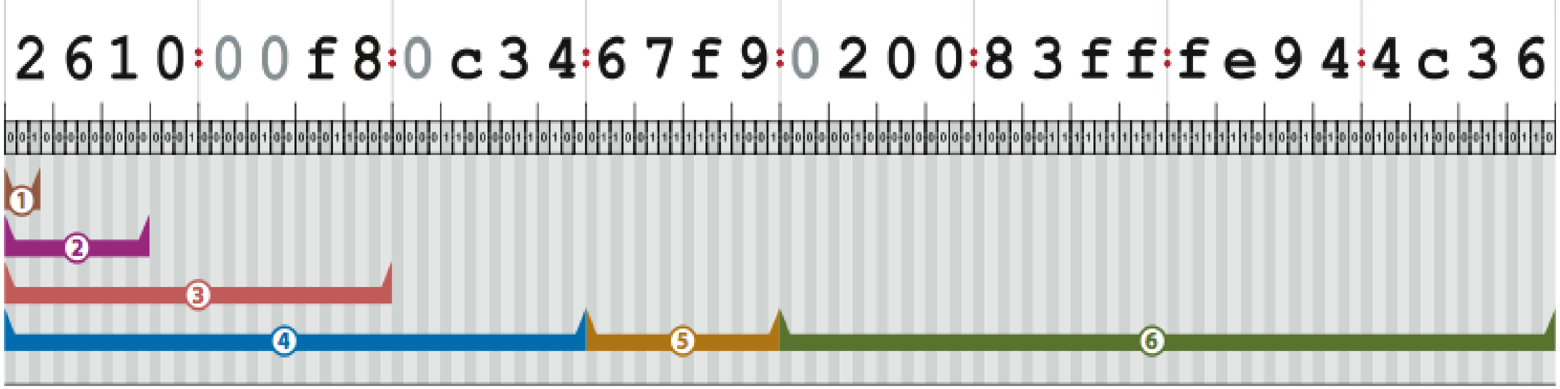

1 **2000::/3**
The current IPv6 address space for unicast allocations is 1/8 of the total address space.

2 **IANA Allocation to Registries (Varies)**
IANA makes assignments to regional registries. New allocations are /12 bits, previous assignment have varied.
For example: "2a01:0000::/16" was assigned by IANA to RIPE NCC (the European and Middle East registry) in December 2005.

3 **"ISP Allocations"**
Regional registries make assignments to local ISPs. A typical assignment is /32 bits, but more space may be assigned.
For example: RIPE NCC assigned "2a01:c000::/19 to France Telecom in December 2005.

4 **"End-Site Allocations"**
ISPs make assignments to their customers. The amount of address space varies, but a /48 bit allocation is common. Organizations can get larger assignments, based on need (Command Information has a /32 allocation), smaller organizations may get less space (for small companies a /56 is common).

5 **"Subnet Assignments"**
Organizations make assignment to individual subnets, where the most common size is /64.
With 16 bits subnetting bits available, an organization can deploy as many as 65,536 subnets.

6 **"Interface ID"**
Interfaces must have a unique identifier on the subnet – often created by embedding the underlying 48-bit (L2) MAC address.
Theoretically then, a single subnet could support 2^64 active hosts – clearly far beyond the practical limit.

SLAAC – EUI-64

The first step is to convert the 48-bit MAC address to a 64-bit value. The 16-bit hex value 0xFFFE is then inserted between these two halves to form a 64-bit address.

Why 0xFFFE? As explained in the IEEE's Guidelines for EUI-64 Registration Authority, this is a reserved value which equipment manufacturers cannot include in "real" EUI-64 address assignments. In other words, any EUI-64 address having 0xFFFE immediately following its OUI portion can be recognized as having been generated from an EUI-48 (or MAC) address.

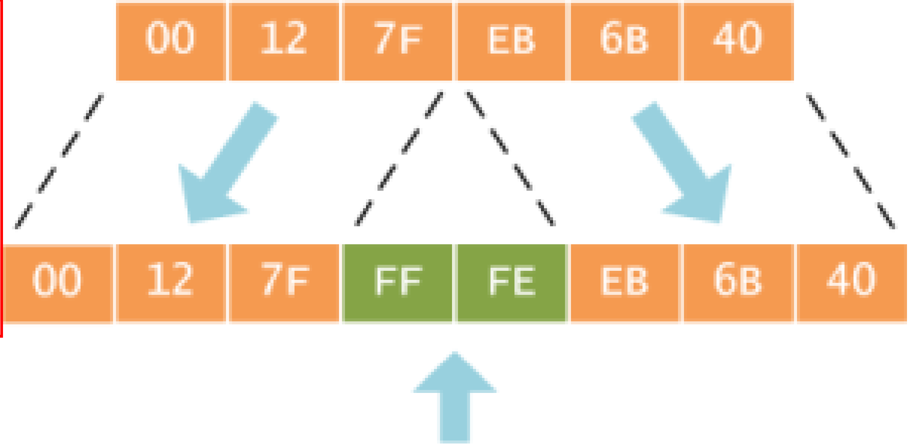

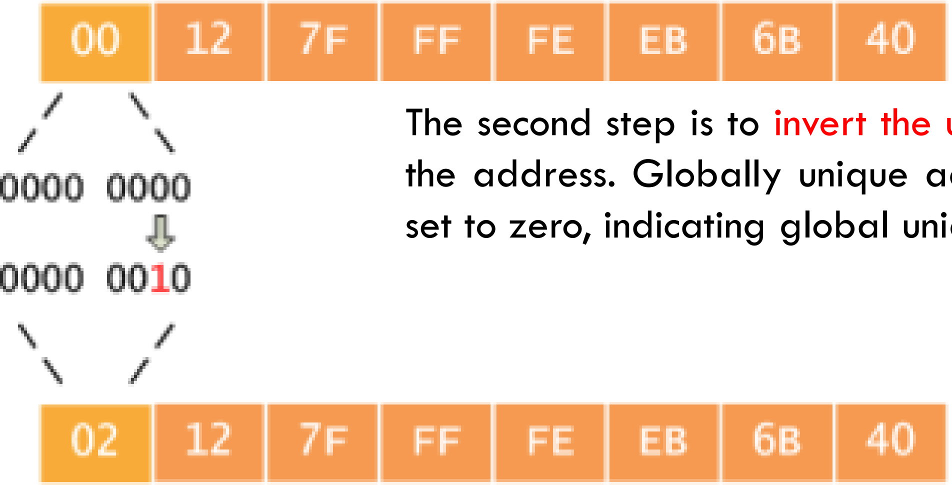

The second step is to invert the universal/local (U/L) flag (bit 7) in the OUI portion of the address. Globally unique addresses assigned by the IEEE originally have this bit set to zero, indicating global uniqueness.

Again, you're probably wondering why this is done. The answer lies buried in section 2.5.1 of RFC 2373

E.G. – EUI-64 ON WINDOWS XP

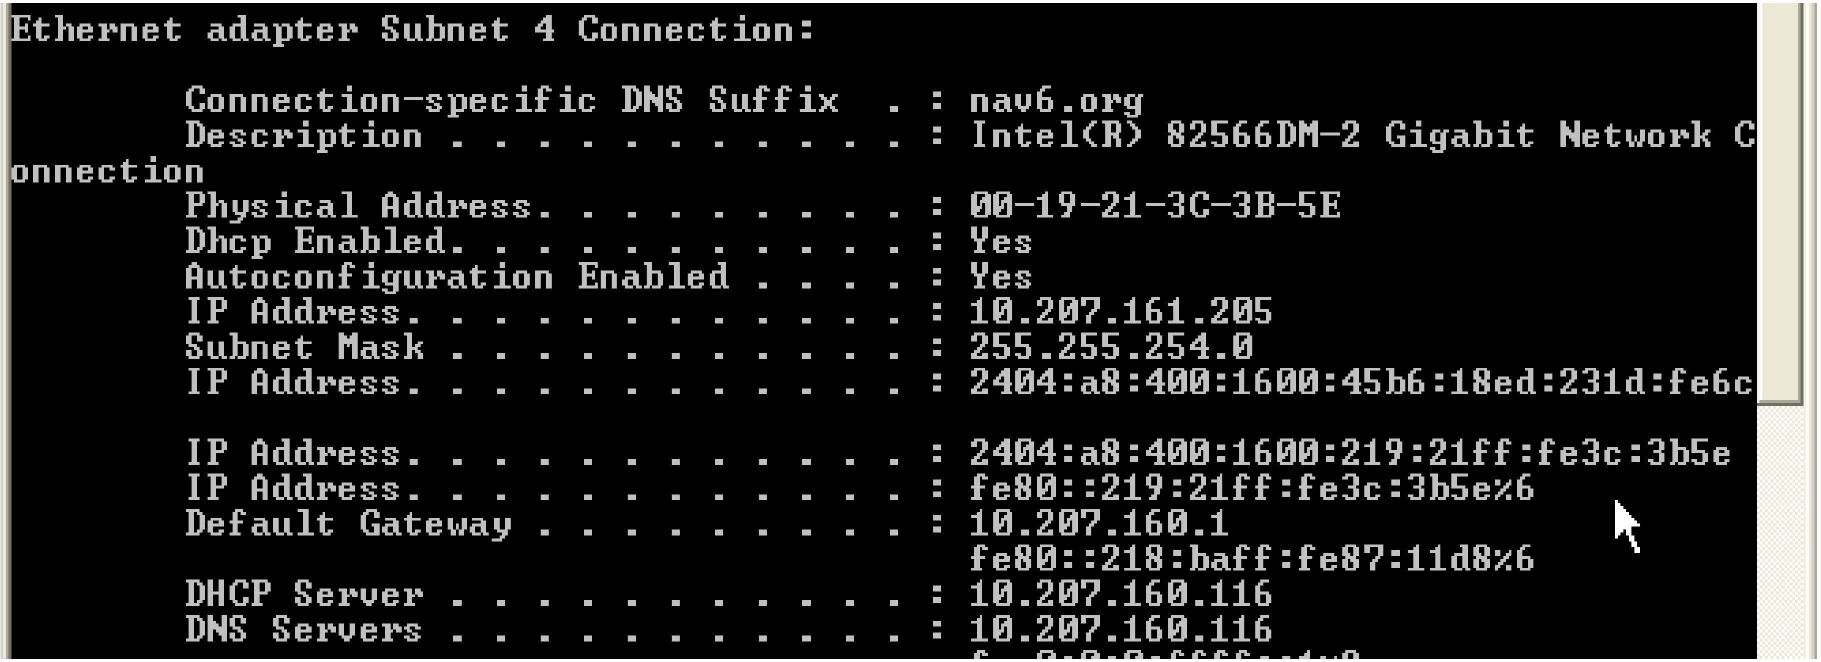

```
Ethernet adapter Subnet 4 Connection:

        Connection-specific DNS Suffix  . : nav6.org
        Description . . . . . . . . . . . : Intel(R) 82566DM-2 Gigabit Network C
onnection
        Physical Address. . . . . . . . . : 00-19-21-3C-3B-5E
        Dhcp Enabled. . . . . . . . . . . : Yes
        Autoconfiguration Enabled . . . . : Yes
        IP Address. . . . . . . . . . . . : 10.207.161.205
        Subnet Mask . . . . . . . . . . . : 255.255.254.0
        IP Address. . . . . . . . . . . . : 2404:a8:400:1600:45b6:18ed:231d:fe6c

        IP Address. . . . . . . . . . . . : 2404:a8:400:1600:219:21ff:fe3c:3b5e
        IP Address. . . . . . . . . . . . : fe80::219:21ff:fe3c:3b5e%6
        Default Gateway . . . . . . . . . : 10.207.160.1
                                            fe80::218:baff:fe87:11d8%6
        DHCP Server . . . . . . . . . . . : 10.207.160.116
        DNS Servers . . . . . . . . . . . : 10.207.160.116
                                            fec0:0:0:ffff::1%2
```

WAYS TO GENERATE INTERFACE - ID

Interface ID can be generated in many different ways:

1. Build one from the layer 2 address in the modified **EUI-64** format. ***Different mechanisms are defined for each media type*** to build the complete interface ID in the modified EUI-64 format.
2. **Autogenerate** a random address as defined in RFC 3041. This assignment mechanism was developed mainly **to limit the exposure of a globally reachable address and to increase privacy**.
3. Acquire the interface ID via **DHCPv6**.
4. **Manual** configuration.
5. **Cryptographically generated addresses (CGAs)** based on RFC 3972 through a **hash that includes a public key**. This method of generating an interface ID provides added security and enables address authentication.

LINK-LOCAL ADDRESS EXPLAINED

- Link-local addresses are utilized by nodes when communicating with neighboring nodes on the same link. **(e.g., in a LAN segment, a virtual LAN (VLAN), etc.)**
- link-local address is required for ***Neighbor Discovery (ND)*** processes and is **always automatically configured**, even in the absence of all other unicast addresses.

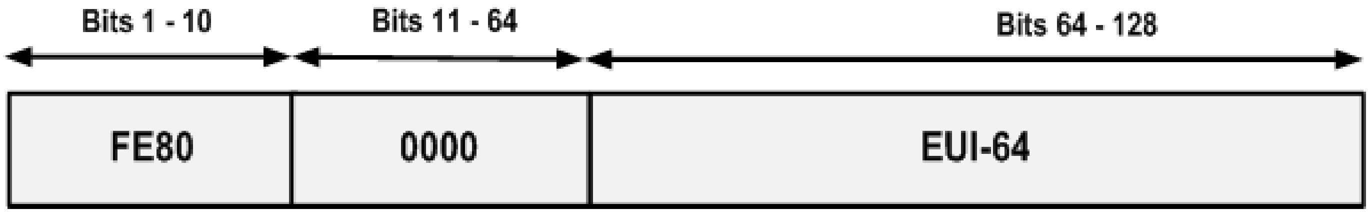

Bits 1 - 10 | Bits 11 - 64 | Bits 64 - 128

FE80 | 0000 | EUI-64

NOTE* Some hosts may use other mechanisms to generate the interface ID.

example

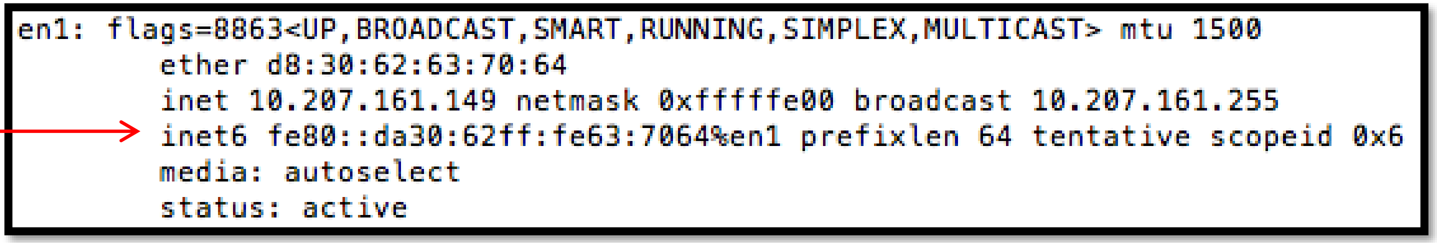

```
en1: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
        ether d8:30:62:63:70:64
        inet 10.207.161.149 netmask 0xfffffe00 broadcast 10.207.161.255
        inet6 fe80::da30:62ff:fe63:7064%en1 prefixlen 64 tentative scopeid 0x6
        media: autoselect
        status: active
```

NOTE* Routers must not forward any packets with Link-Local source or destination addresses to other links.

SITE LOCAL ADDRESS EXPLAINED

- Originally designed to be used for addressing inside of a site <u>without the need for a global prefix.</u>
- Used between nodes that communicate with other nodes in the **same site (organization)**
- The scope of a site-local address is the site, which is the organization intranet (internetwork)

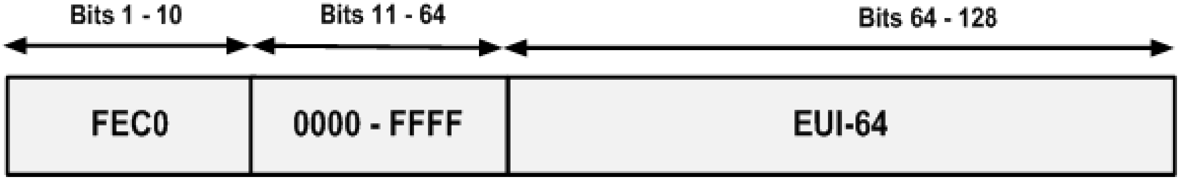

Bits 1 - 10

Bits 11 - 64

Bits 64 - 128

FEC0

0000 - FFFF

EUI-64

NOTE* Site-local addresses are now deprecated.

NOTE* The special behavior of this prefix defined in [RFC3513] must no longer be supported in new implementations.

UNIQUE LOCAL ADDRESS (ULA) EXPLAINED

- An IPv6 address in the block fc00::/7, defined in RFC 4193
- Counterpart of the IPv4 private address
- Available for use in **private networks**.

The address block fc00::/7 is **divided into two /8 groups**:

- fc00::/8 [has not been defined yet]
- fd00::/8 is defined for /48 prefixes

formed by setting the 40 least-significant bits of the prefix to a randomly-generated bit string

Address format = fdxx:xxxx:xxxx::

NOTE* fd00::/8 ULAs are not meant to be routed outside their administrative domain (site or organization)

Abstract from FRC 4193
This document defines an IPv6 unicast address format that is globally unique and is intended for local communications, usually inside of a site. These addresses are not expected to be routable on the global Internet.

UNSPECIFIED ADDRESS EXPLAINED

Special-Purpose Unicast Address - 1

- 0:0:0:0:0:0:0:0 (that is, ::) indicates the absence of an address.
- Typically used as a source address for PDUs that are attempting to verify the *uniqueness of a tentative address*

USAGE RULES:

- The unspecified address must **NOT** be used as the **destination address** of IPv6 packets or in IPv6 Routing headers.
- An IPv6 packet with a **source address** of unspecified must **NEVER** be forwarded by an IPv6 router.

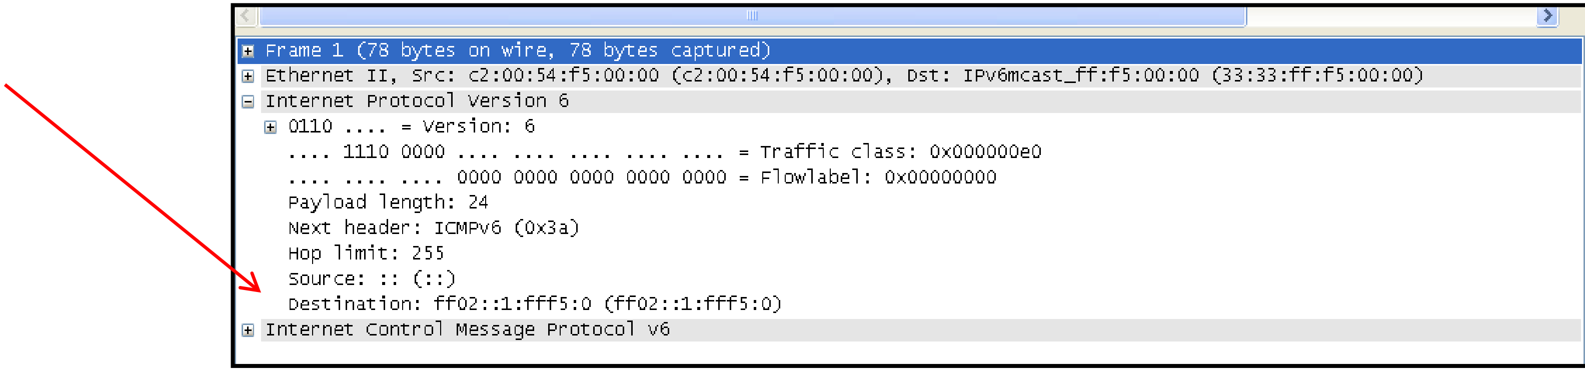

```
⊞ Frame 1 (78 bytes on wire, 78 bytes captured)
⊞ Ethernet II, Src: c2:00:54:f5:00:00 (c2:00:54:f5:00:00), Dst: IPv6mcast_ff:f5:00:00 (33:33:ff:f5:00:00)
⊟ Internet Protocol Version 6
  ⊞ 0110 .... = Version: 6
    .... 1110 0000 .... .... .... .... .... = Traffic class: 0x000000e0
    .... .... .... 0000 0000 0000 0000 0000 = Flowlabel: 0x00000000
    Payload length: 24
    Next header: ICMPv6 (0x3a)
    Hop limit: 255
    Source: :: (::)
    Destination: ff02::1:fff5:0 (ff02::1:fff5:0)
⊞ Internet Control Message Protocol v6
```

LOOPBACK ADDRESS EXPLAINED

Special-Purpose Unicast Address - 2

- 0:0:0:0:0:0:0:1 (that is, ::1) identifies a loopback interface, **enabling a node to send PDUs to itself**.
- PDUs addressed to the loopback address are **never** sent on a link or forwarded by an IPv6 router.

USAGE RULES:

- The unspecified address must **NOT** be used as the **source address** of IPv6 packets sent **outside** of a node.
- Must **NEVER** be forwarded by an IPv6 router.
- A packet received on an interface with a **destination address** of loopback **MUST** be dropped.

IPV4-COMPATIBLE IPV6 ADDRESS EXPLAINED

IPv6 Addresses with Embedded IPv4 Addresses Type - 1

- The "IPv4-Compatible IPv6 address" was defined to assist in the **IPv6 transition**.

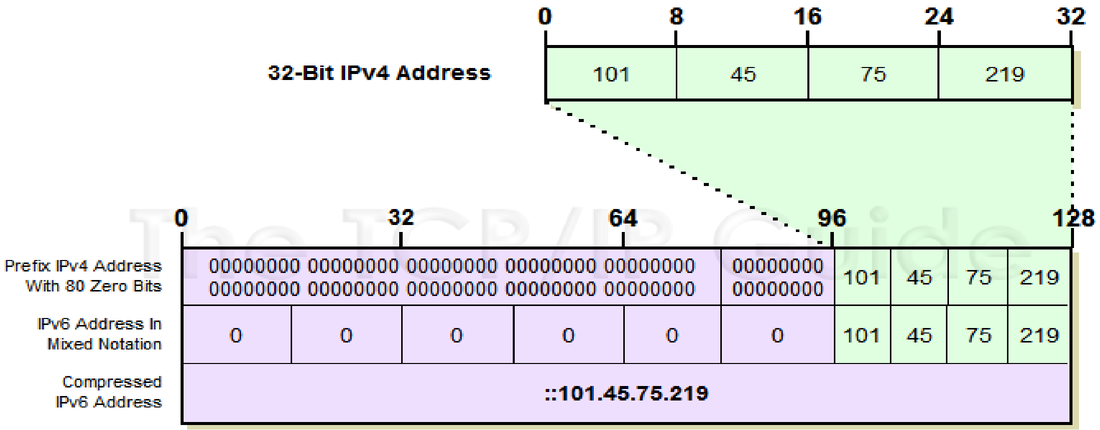

- The "IPv4-Compatible IPv6 address" is now **deprecated** because the current IPv6 transition mechanisms no longer use these addresses.
- New or updated implementations are **NOT** required to support this address type.

NOTE* The IPv4 address used in the "IPv4-Compatible IPv6 address" must be a globally-unique IPv4 unicast address.

IPV4-MAPPED IPV6 ADDRESS EXPLAINED

IPv6 Addresses with Embedded IPv4 Addresses Type - 2

- This address type is used to represent the addresses of IPv4 nodes as IPv6 addresses.
- Most implementations of dual-stack allow **IPv6-only** applications to **interoperate** with **both IPv4 and IPv6 nodes**.
- **IPv4 packets** going to **IPv6 applications** on a dual-stack node reach their destination because their addresses are mapped by using **IPv4-mapped IPv6 addresses**.

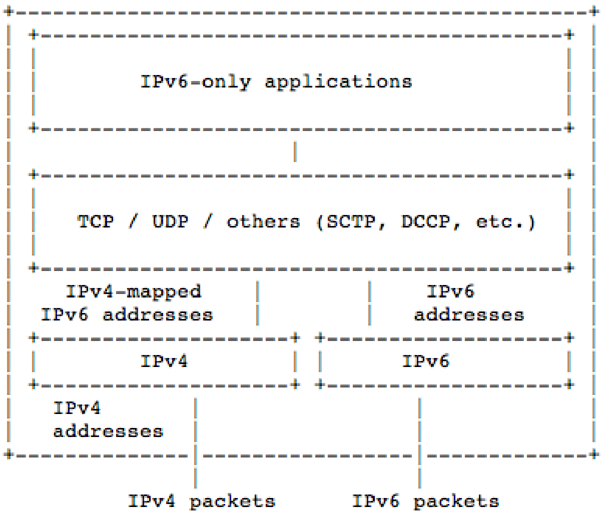

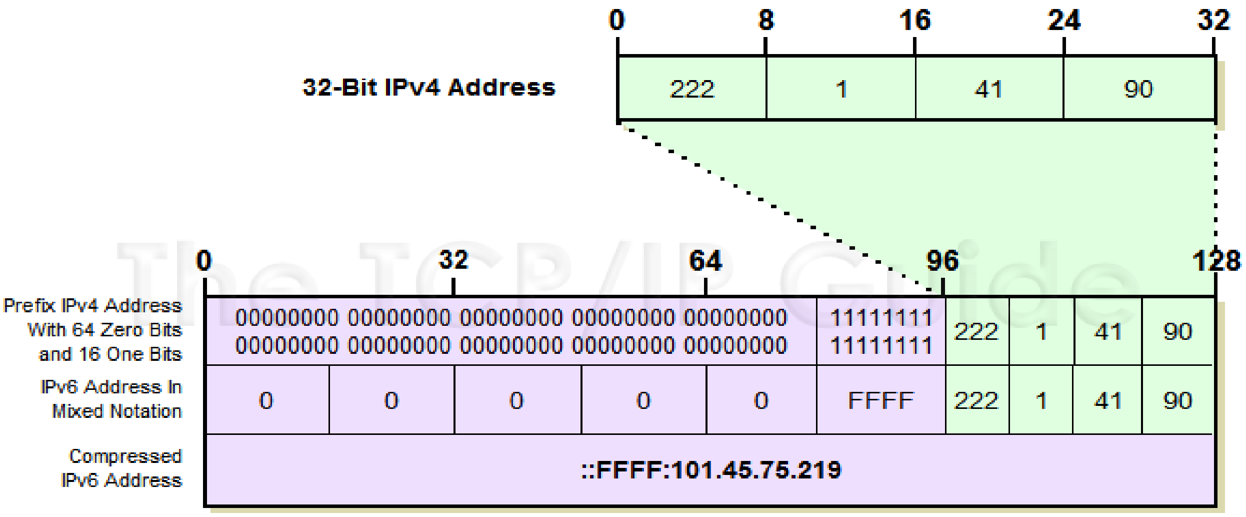

Prefix IPv4 Address With 64 Zero Bits and 16 One Bits	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000					11111111 11111111	222	1	41	90
IPv6 Address In Mixed Notation	0	0	0	0	0	FFFF	222	1	41	90
Compressed IPv6 Address	::FFFF:101.45.75.219									

UNDERSTANDING ADDRESS TYPES

IPv6 Address Type 2

IPv6 Anycast Address

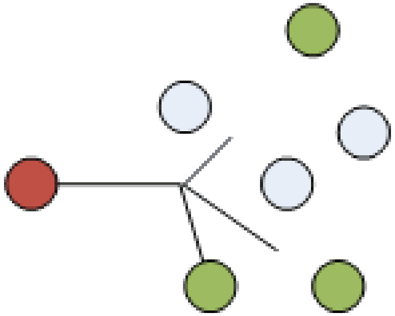

ANYCAST ADDRESS EXPLAINED

- An address that is assigned to **more than one interface** (typically belonging to different nodes)
- Packet sent to an anycast address is **routed** to the **"nearest"** interface having that address.

  *** according to the routing protocols' measure of distance.**
- Anycast addresses are **allocated from the unicast address space**.
- Anycast addresses are **syntactically indistinguishable** from unicast addresses.

Usage:

- to identify the **set of routers belonging** to an organization providing Internet service.
- to identify the **set of routers attached** to a particular subnet.
- the **set of routers providing** entry into a particular routing domain.

ANYCAST ADDRESS EXPLAINED

Currently, the following anycast identifiers for these reserved subnet anycast addresses are defined:

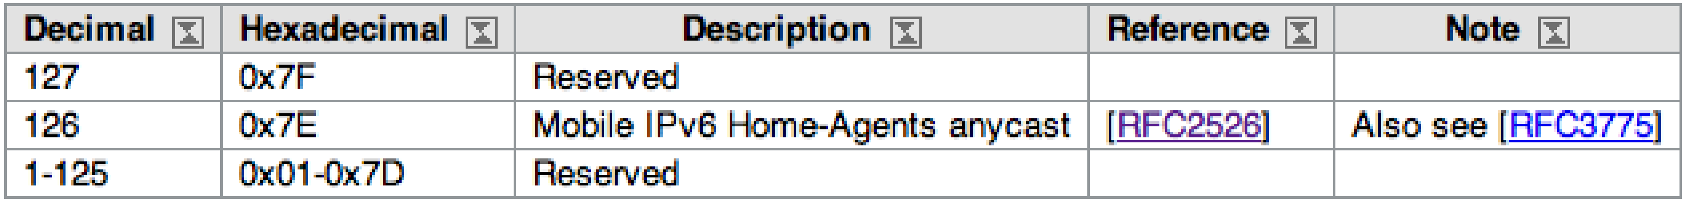

Decimal	Hexadecimal	Description	Reference	Note
127	0x7F	Reserved		
126	0x7E	Mobile IPv6 Home-Agents anycast	[RFC2526]	Also see [RFC3775]
1-125	0x01-0x7D	Reserved		

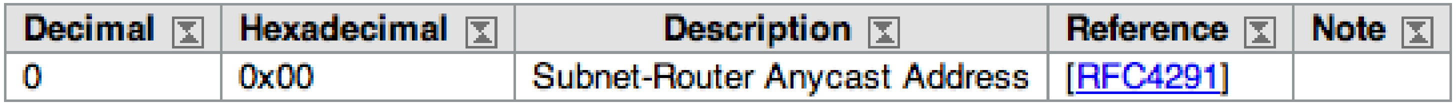

Decimal	Hexadecimal	Description	Reference	Note
0	0x00	Subnet-Router Anycast Address	[RFC4291]	

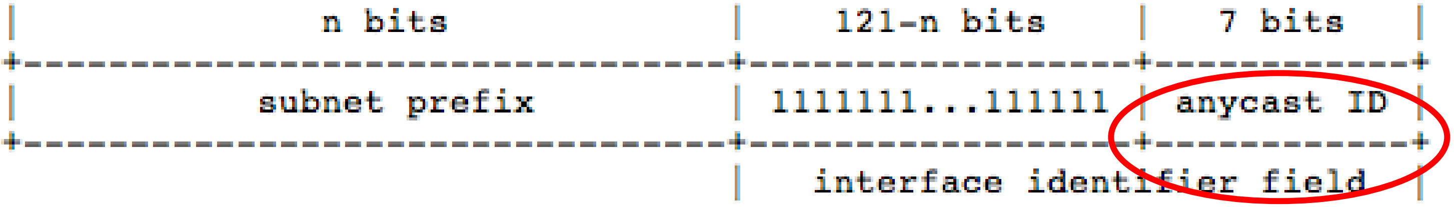

```
|               n bits                |    121-n bits      |   7 bits   |
+-------------------------------------+--------------------+------------+
|           subnet prefix             | 1111111...111111   | anycast ID |
+-------------------------------------+--------------------+------------+
                                      |   interface identifier field    |
```

ANYCAST ADDRESS EXPLAINED

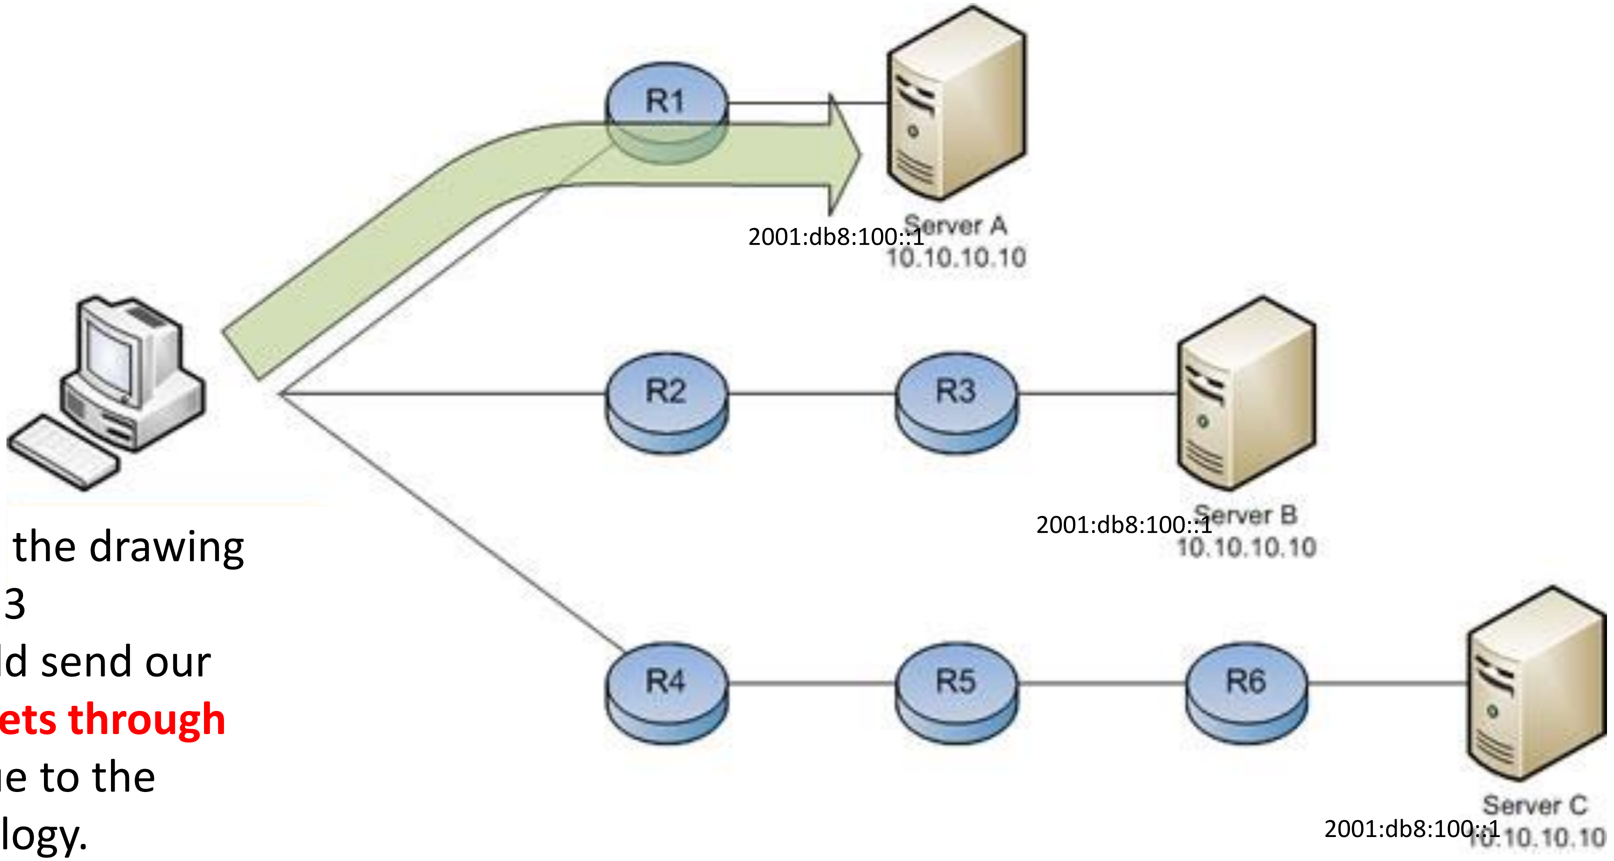

According to the drawing above, layer 3 routing would send our **client's packets through router R1** due to the routing topology.

Should router R1 or Server A fail, our client's packets would **automatically be rerouted** to the next nearest server via routers R2 and R3, and so forth.

Source = http://www.netlinxinc.com

REQUIRED ANYCAST ADDRESS

The **Subnet-Router anycast address** is predefined.

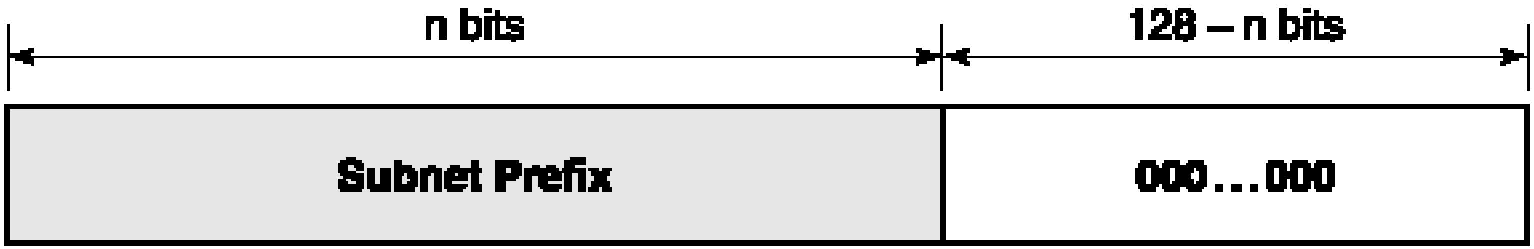

- The "**subnet prefix**" in an anycast address is the **prefix that identifies a specific link**.
- Packets will be delivered to **one router on the subnet.**
- Intended to be used for applications where a node needs to communicate with **any one** of the **set of routers**.

UNDERSTANDING ADDRESS TYPES

IPv6 Address Type 3

IPv6 Multicast Address

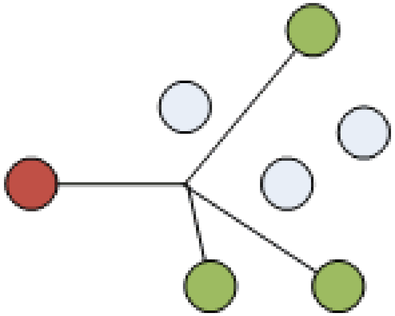

MULTICAST ADDRESS EXPLAINED

❑ An IPv6 multicast address is an identifier for a **group of interfaces** (typically on different nodes).

❑ An interface may belong to **any number of multicast groups**.

❑ Group membership in multicast is **dynamic**, allowing hosts to join and leave the group at any time.

❑ It’s used for **one-to-many** communication.

❑ Packets sent to a multicast address are delivered to **all interfaces** that are identified by the address.

NOTE* Multicast addresses cannot be utilized as **source** addresses

UNDERSTANDING MULTICAST ADDRESS FLAGS

Multicast addresses have the following format:

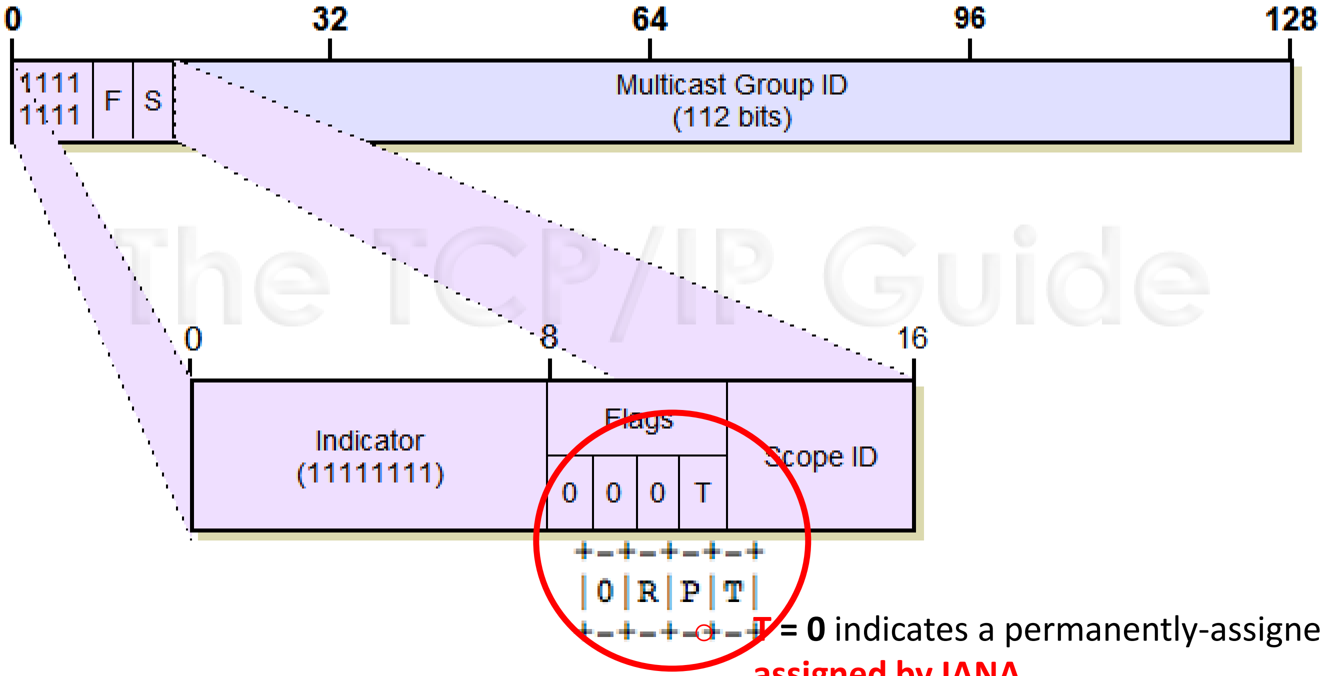

- **T = 0** indicates a permanently-assigned ("**well-known**") multicast address, **assigned by IANA**.
- **T = 1** indicates a **non-permanently-assigned** ("transient" or "dynamically" assigned) multicast address
- **P**; Whether or not assigned based on network prefix (Refer to RFC 3306)
- **R;** Whether Rendezvous Point embedded or not (Refer to RFC 3956

NOTE* Non-permanently)y-assigned multicast addresses are meaningful only within a given scope.

UNDERSTANDING MULTICAST ADDRESS SCOPE

A 4-bit multicast scope value used to limit the scope of the multicast group.
The values are as follows:

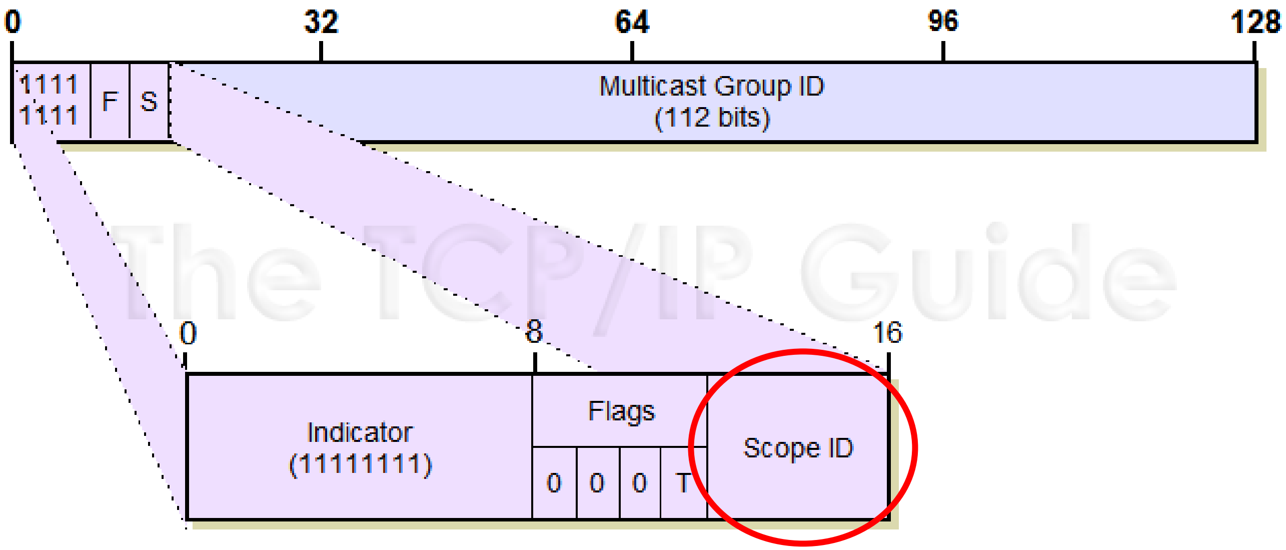

ID	Description
0	Reserved
1	Interface-Local scope
2	**Link-Local scope**
3	Reserved
4	Admin-Local scope
5	Site-Local scope
6	(unassigned)
7	(unassigned)
8	Organization-Local scope
9	(unassigned)
A	(unassigned)
B	(unassigned)
C	(unassigned)
D	(unassigned)
E	Global Scope
F	Reserved

NOTE* Routers must **NOT** forward any multicast packets beyond of the scope indicated by the scope field in the **destination** multicast address.

PRE-DEFINED MULTICAST ADDRESS

Node-Local Scope

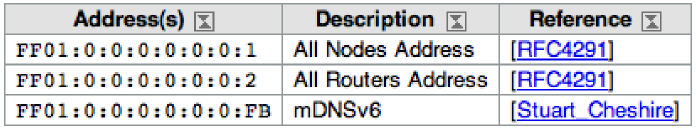

Address(s)	Description	Reference
FF01:0:0:0:0:0:0:1	All Nodes Address	[RFC4291]
FF01:0:0:0:0:0:0:2	All Routers Address	[RFC4291]
FF01:0:0:0:0:0:0:FB	mDNSv6	[Stuart_Cheshire]

Site-Local Scope

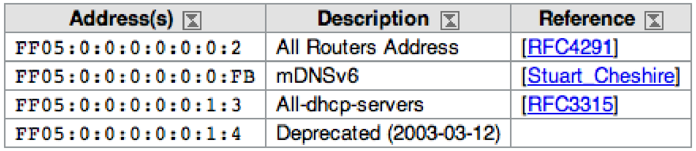

Address(s)	Description	Reference
FF05:0:0:0:0:0:0:2	All Routers Address	[RFC4291]
FF05:0:0:0:0:0:0:FB	mDNSv6	[Stuart_Cheshire]
FF05:0:0:0:0:0:1:3	All-dhcp-servers	[RFC3315]
FF05:0:0:0:0:0:1:4	Deprecated (2003-03-12)	

Link-Local Scope

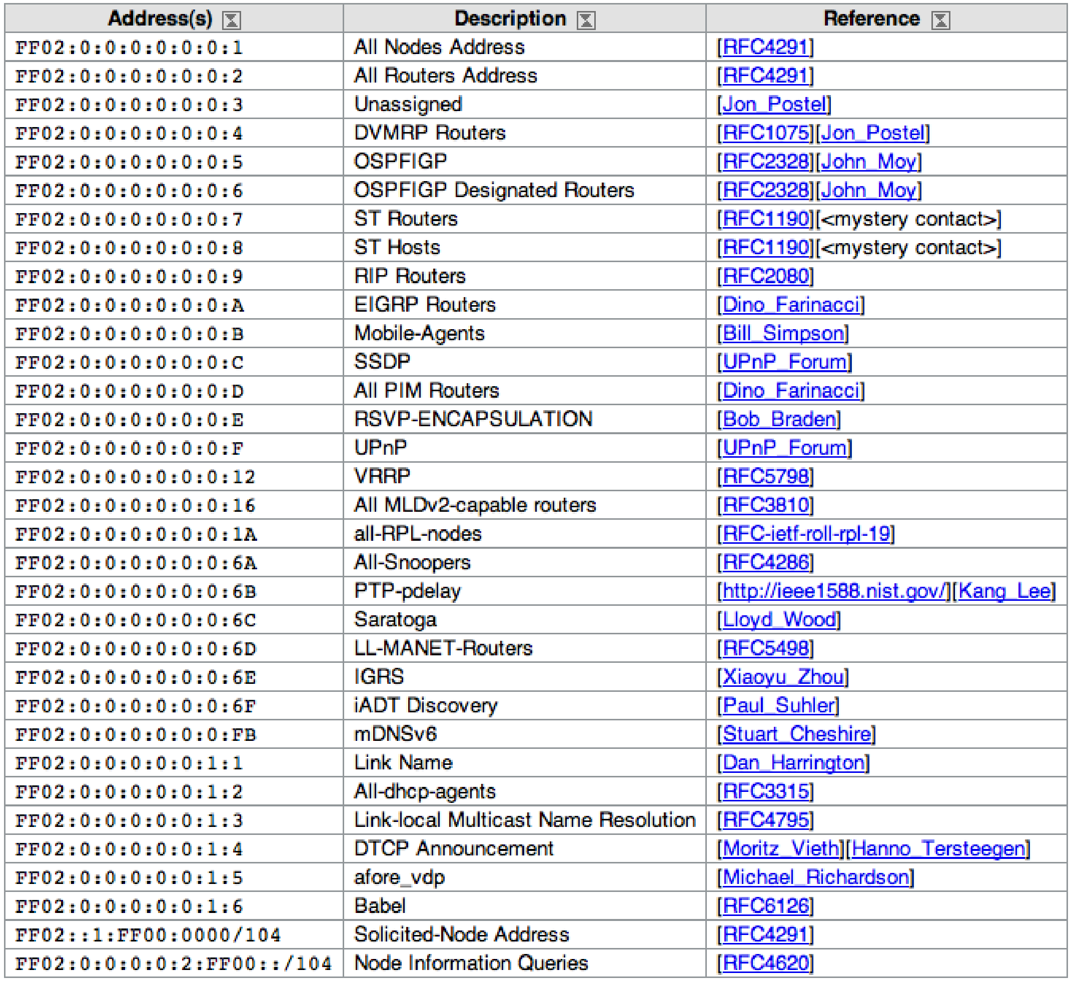

Address(s)	Description	Reference
FF02:0:0:0:0:0:0:1	All Nodes Address	[RFC4291]
FF02:0:0:0:0:0:0:2	All Routers Address	[RFC4291]
FF02:0:0:0:0:0:0:3	Unassigned	[Jon_Postel]
FF02:0:0:0:0:0:0:4	DVMRP Routers	[RFC1075][Jon_Postel]
FF02:0:0:0:0:0:0:5	OSPFIGP	[RFC2328][John_Moy]
FF02:0:0:0:0:0:0:6	OSPFIGP Designated Routers	[RFC2328][John_Moy]
FF02:0:0:0:0:0:0:7	ST Routers	[RFC1190][<mystery contact>]
FF02:0:0:0:0:0:0:8	ST Hosts	[RFC1190][<mystery contact>]
FF02:0:0:0:0:0:0:9	RIP Routers	[RFC2080]
FF02:0:0:0:0:0:0:A	EIGRP Routers	[Dino_Farinacci]
FF02:0:0:0:0:0:0:B	Mobile-Agents	[Bill_Simpson]
FF02:0:0:0:0:0:0:C	SSDP	[UPnP_Forum]
FF02:0:0:0:0:0:0:D	All PIM Routers	[Dino_Farinacci]
FF02:0:0:0:0:0:0:E	RSVP-ENCAPSULATION	[Bob_Braden]
FF02:0:0:0:0:0:0:F	UPnP	[UPnP_Forum]
FF02:0:0:0:0:0:0:12	VRRP	[RFC5798]
FF02:0:0:0:0:0:0:16	All MLDv2-capable routers	[RFC3810]
FF02:0:0:0:0:0:0:1A	all-RPL-nodes	[RFC-ietf-roll-rpl-19]
FF02:0:0:0:0:0:0:6A	All-Snoopers	[RFC4286]
FF02:0:0:0:0:0:0:6B	PTP-pdelay	[http://ieee1588.nist.gov/][Kang_Lee]
FF02:0:0:0:0:0:0:6C	Saratoga	[Lloyd_Wood]
FF02:0:0:0:0:0:0:6D	LL-MANET-Routers	[RFC5498]
FF02:0:0:0:0:0:0:6E	IGRS	[Xiaoyu_Zhou]
FF02:0:0:0:0:0:0:6F	iADT Discovery	[Paul_Suhler]
FF02:0:0:0:0:0:0:FB	mDNSv6	[Stuart_Cheshire]
FF02:0:0:0:0:0:1:1	Link Name	[Dan_Harrington]
FF02:0:0:0:0:0:1:2	All-dhcp-agents	[RFC3315]
FF02:0:0:0:0:0:1:3	Link-local Multicast Name Resolution	[RFC4795]
FF02:0:0:0:0:0:1:4	DTCP Announcement	[Moritz_Vieth][Hanno_Tersteegen]
FF02:0:0:0:0:0:1:5	afore_vdp	[Michael_Richardson]
FF02:0:0:0:0:0:1:6	Babel	[RFC6126]
FF02::1:FF00:0000/104	Solicited-Node Address	[RFC4291]
FF02:0:0:0:0:2:FF00::/104	Node Information Queries	[RFC4620]

Source = http://www.iana.org/assignments/ipv6-multicast-addresses/ipv6-multicast-addresses.xml

SOLICITED-NODE MULTICAST ADDRESS EXPLAINED

- Solicited-Node multicast address (SNMA) are **computed as a function** of a node's **unicast** and **anycast** addresses.
- A SNMA is formed by taking the low-order 24 bits (last 6 characters) of an address (unicast or anycast) and appending those bits to the prefix:

FF02::1:FFXX:XXXX/104

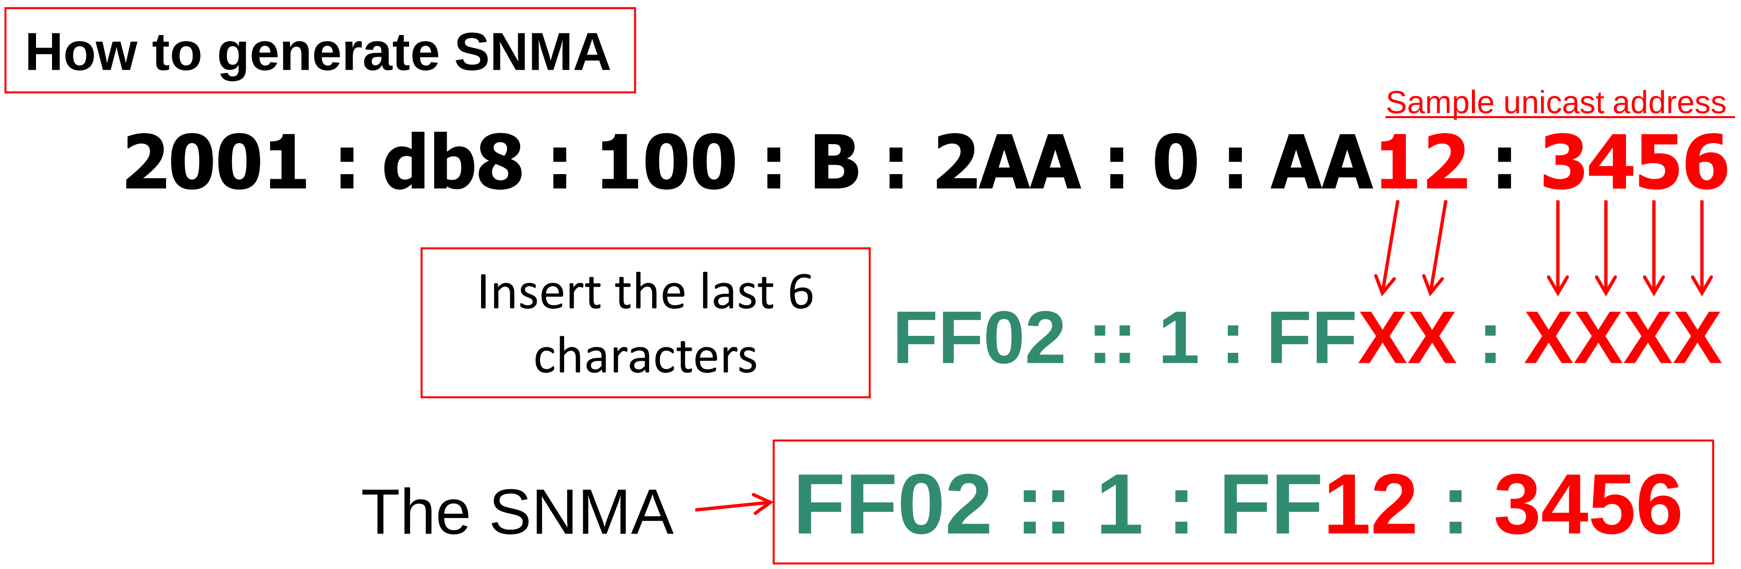

How to generate SNMA

Sample unicast address

2001 : db8 : 100 : B : 2AA : 0 : AA12 : 3456

Insert the last 6 characters

FF02 :: 1 : FFXX : XXXX

The SNMA → FF02 :: 1 : FF12 : 3456

A NODE'S REQUIRED ADDRESSES

HOST

A **host** is required to recognize the following addresses as identifying itself:

- Its required **Link-Local address** for each interface.
- Any additional **Unicast and Anycast addresses** that have been configured for the node's interfaces ***(manually or automatically)***.
- The **loopback address**.
- The **All-Nodes multicast addresses** ***(all scope)***.
- The **Solicited-Node multicast address** for each of its ***unicast and anycast*** addresses.
- **Multicast addresses** of ***all other groups*** to which the node belongs. (optional)

A NODE'S REQUIRED ADDRESSES

ROUTER

A router is required to **recognize all addresses that a host is required** to recognize, **PLUS** the following addresses as identifying itself::

- The **Subnet-Router Anycast addresses** for all interfaces for which it is ***configured to act as a router***.
- **All other Anycast addresses** with which the router has been configured.
- The **All-Routers multicast addresses**. (optional)

MULTIHOMING

Multihoming is a technique used to **increase the reliability** of the Internet connection for an IP network.

- Single Link, Multiple IP address (Spaces)
- Multiple Interfaces, Single IP address per interface
- Multiple Links, Single IP address (Space)
- Multiple Links, Multiple IP address (Spaces)

NOTE* Multihoming in IPv6 is not yet standardized

Current solutions:

- Get a Provider Independent Address Space
- Automated renumbering
- Maintaining multiple simultaneous sets of host addresses, from different upstream

The most effective technique for multihoming in IPv4 is what is known as "Provider Independent" (PI) address space combined with the Border Gateway Protocol (BGP)

OPERATIONS OF IPV6

UNDERSTANDING ICMPV6

Internet Control Message Protocol version 6 (ICMPv6)
is the implementation of the ICMP **for** IPv6

Functions as an integral part of IPv6 and performs:

- error reporting
- diagnostic functions (e.g., ping)
- and as a framework for extensions to implement future changes

TECHNICAL DETAILS OF ICMPV6

ICMPv6 messages may be classified into **two** categories:

- ***error*** *messages*
- ***information*** *messages*

ICMPv6 messages are transported by **IPv6 packets** in which the IPv6 Next Header value for ICMPv6 is set to 58.

58 | 0x3A

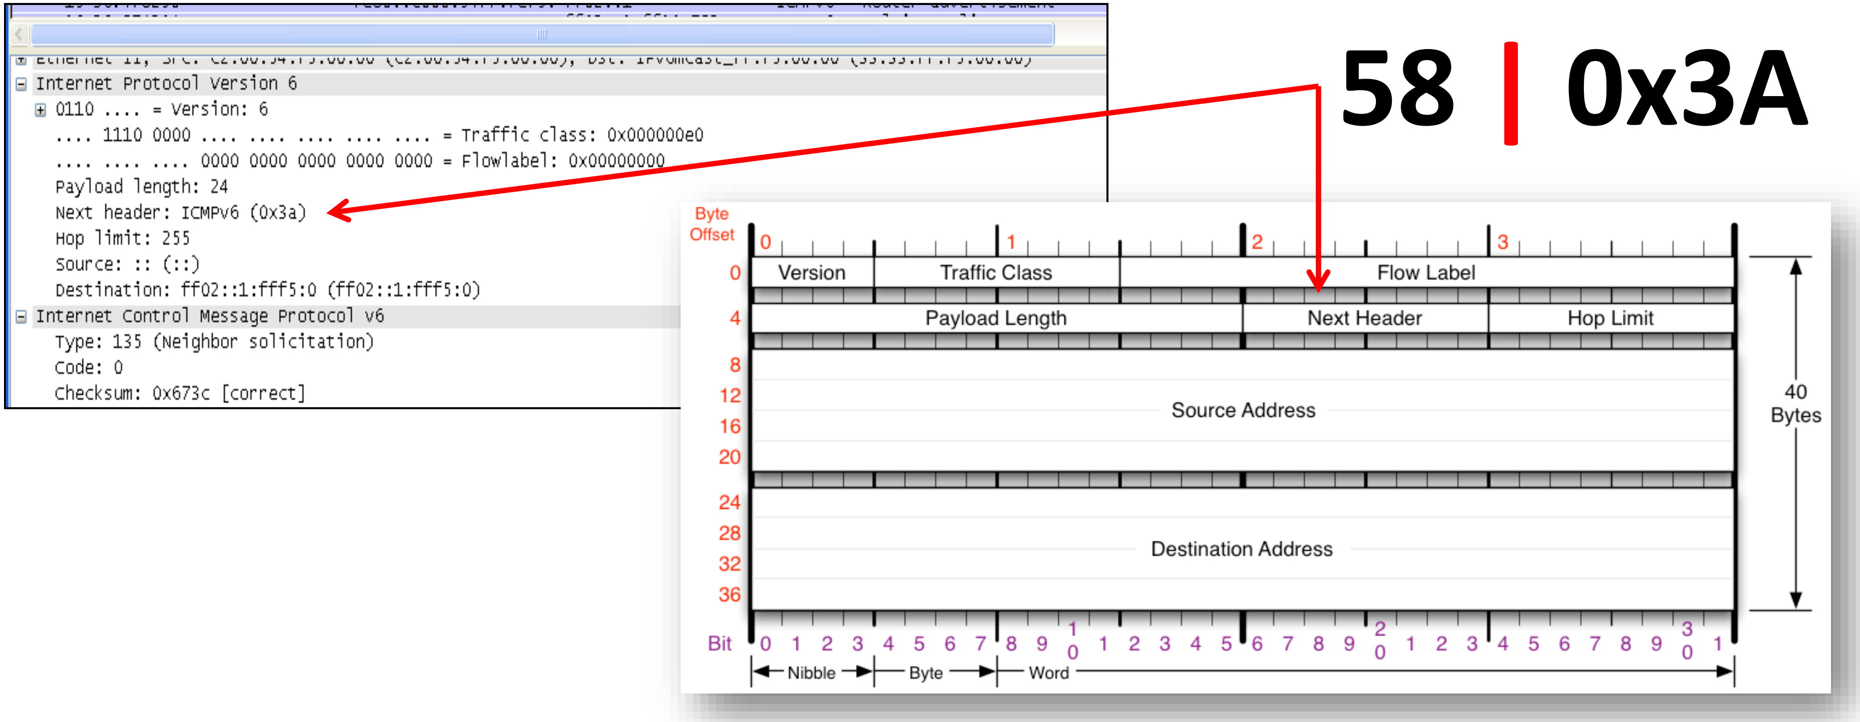

```
⊟ Internet Protocol Version 6
  ⊞ 0110 .... = Version: 6
    .... 1110 0000 .... .... .... .... .... = Traffic class: 0x000000e0
    .... .... .... 0000 0000 0000 0000 0000 = Flowlabel: 0x00000000
    Payload length: 24
    Next header: ICMPv6 (0x3a)
    Hop limit: 255
    Source: :: (::)
    Destination: ff02::1:fff5:0 (ff02::1:fff5:0)
⊟ Internet Control Message Protocol v6
    Type: 135 (Neighbor solicitation)
    Code: 0
    Checksum: 0x673c [correct]
```

ICMPV6 PACKET FORMAT

The ICMPv6 packet consists of a header and the protocol payload

The header contains only **three** fields:

- type (8 bits)
- code (8 bits)
- checksum (16 bits)

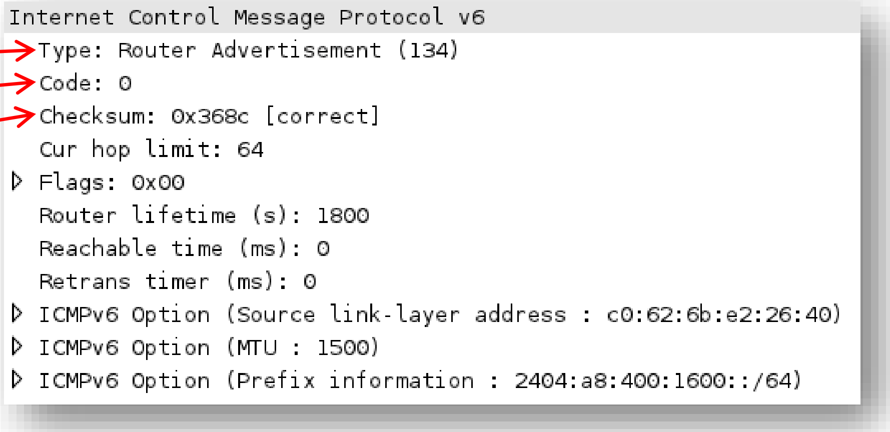

```
Internet Control Message Protocol v6
   Type: Router Advertisement (134)
   Code: 0
   Checksum: 0x368c [correct]
   Cur hop limit: 64
 ▷ Flags: 0x00
   Router lifetime (s): 1800
   Reachable time (ms): 0
   Retrans timer (ms): 0
 ▷ ICMPv6 Option (Source link-layer address : c0:62:6b:e2:26:40)
 ▷ ICMPv6 Option (MTU : 1500)
 ▷ ICMPv6 Option (Prefix information : 2404:a8:400:1600::/64)
```

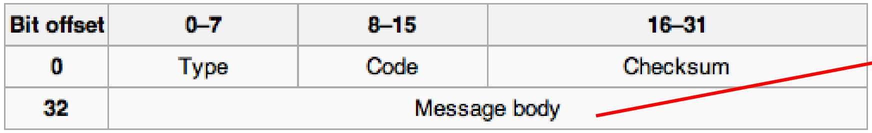

Bit offset	0–7	8–15	16–31
0	Type	Code	Checksum
32	Message body		

Field	Description
Type	Specifies the type of the message. Values in the range from **0 to 127 indicate an error message**, while values in the range from **128 to 255 indicate an information message**.
Code	Value depends on the message type and provides an additional level of message granularity.
Checksum	Provides a minimal level of integrity verification for the ICMP message.

TYPES OF ICMPV6 MESSAGES

Error messages

Informational Messages

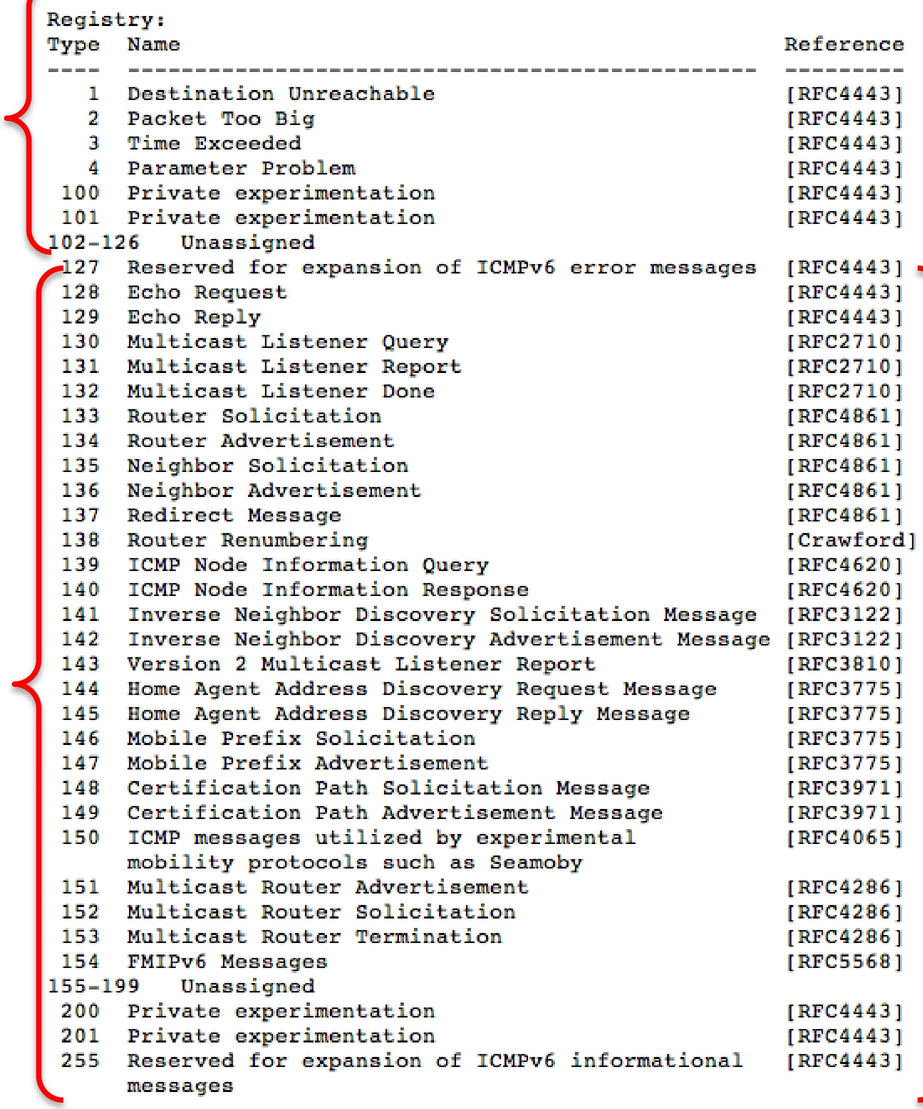

```
Registry:
Type  Name                                               Reference
----  -------------------------------------------------  ---------
   1  Destination Unreachable                            [RFC4443]
   2  Packet Too Big                                     [RFC4443]
   3  Time Exceeded                                      [RFC4443]
   4  Parameter Problem                                  [RFC4443]
 100  Private experimentation                            [RFC4443]
 101  Private experimentation                            [RFC4443]
102-126   Unassigned
 127  Reserved for expansion of ICMPv6 error messages    [RFC4443]
 128  Echo Request                                       [RFC4443]
 129  Echo Reply                                         [RFC4443]
 130  Multicast Listener Query                           [RFC2710]
 131  Multicast Listener Report                          [RFC2710]
 132  Multicast Listener Done                            [RFC2710]
 133  Router Solicitation                                [RFC4861]
 134  Router Advertisement                               [RFC4861]
 135  Neighbor Solicitation                              [RFC4861]
 136  Neighbor Advertisement                             [RFC4861]
 137  Redirect Message                                   [RFC4861]
 138  Router Renumbering                                 [Crawford]
 139  ICMP Node Information Query                        [RFC4620]
 140  ICMP Node Information Response                     [RFC4620]
 141  Inverse Neighbor Discovery Solicitation Message    [RFC3122]
 142  Inverse Neighbor Discovery Advertisement Message  [RFC3122]
 143  Version 2 Multicast Listener Report                [RFC3810]
 144  Home Agent Address Discovery Request Message       [RFC3775]
 145  Home Agent Address Discovery Reply Message         [RFC3775]
 146  Mobile Prefix Solicitation                         [RFC3775]
 147  Mobile Prefix Advertisement                        [RFC3775]
 148  Certification Path Solicitation Message            [RFC3971]
 149  Certification Path Advertisement Message           [RFC3971]
 150  ICMP messages utilized by experimental             [RFC4065]
      mobility protocols such as Seamoby
 151  Multicast Router Advertisement                     [RFC4286]
 152  Multicast Router Solicitation                      [RFC4286]
 153  Multicast Router Termination                       [RFC4286]
 154  FMIPv6 Messages                                    [RFC5568]
155-199   Unassigned
 200  Private experimentation                            [RFC4443]
 201  Private experimentation                            [RFC4443]
 255  Reserved for expansion of ICMPv6 informational     [RFC4443]
      messages
```

SOURCE: http://www.iana.org/assignments/icmpv6-parameters

OPERATIONS OF ICMPV6

Message checksum

ICMPv6 provides a minimal level of message integrity verification by the using a 16-bit checksum in its header.

Message processing

- When an ICMPv6 node receives a packet, it must undertake actions that depend on the type of message.
- The ICMPv6 protocol **must** limit the number of error messages sent to the same destination to avoid network overloading.
- An ICMP error message **must** never be sent in response to another ICMP error message.

UNDERSTANDING NDP

Neighbor Discovery Protocol

The Neighbor Discovery (ND) Protocol is a protocol in the Internet Protocol Suite used with IPv6

TECHNICAL DETAILS

Provides functionality for

❑**Router discovery**: hosts can locate routers residing on attached links.

❑**Prefix discovery**: hosts can discover address prefixes that are on-link for attached links.

❑**Parameter discovery**: hosts can find link parameters (e.g MTU)

❑**Address autoconfiguration**: stateless configuration of addresses of network interfaces.

❑**Address resolution**: mapping between IP addresses and link-layer addresses.

❑**Next-hop determination**: hosts can find next-hop routers for a destination.

❑**Neighbor unreachability detection (NUD)**: determine that a neighbor is no longer reachable on the link.

❑**Duplicate address detection (DAD)**: nodes can check whether an address is already in use.

❑**Redirect**: router can inform a node about better first-hop routers.

NDP MESSAGES

Understanding NDP Messages

The protocol defines **five** different **ICMPv6 packet** types:

- Router Solicitation
- Router Advertisement
- Neighbor Solicitation
- Neighbor Advertisement

ROUTER SOLICITATION (RS)

What Is IPv6 Router Solicitation (RS)?

When a host does not have a configured unicast address, for example at system startup, it sends a router solicitation message.

- RS messages are originated only by the **HOSTS**.
- Used by hosts to find the **Routers** **on the link**.
- Routers **respond** to RS message **by sending** an ***RA***.
- A RS message has a value of **133** in the Type field of the ICMP packet header.
- The **source** address used in a router solicitation messages is **usually** the unspecified IPv6 address **::**
- If the host **has a configured unicast address**, it will be used to send the RS.
- The **destination** is the **all-routers multicast address (FF02::2)**

ROUTER SOLICITATION PACKET

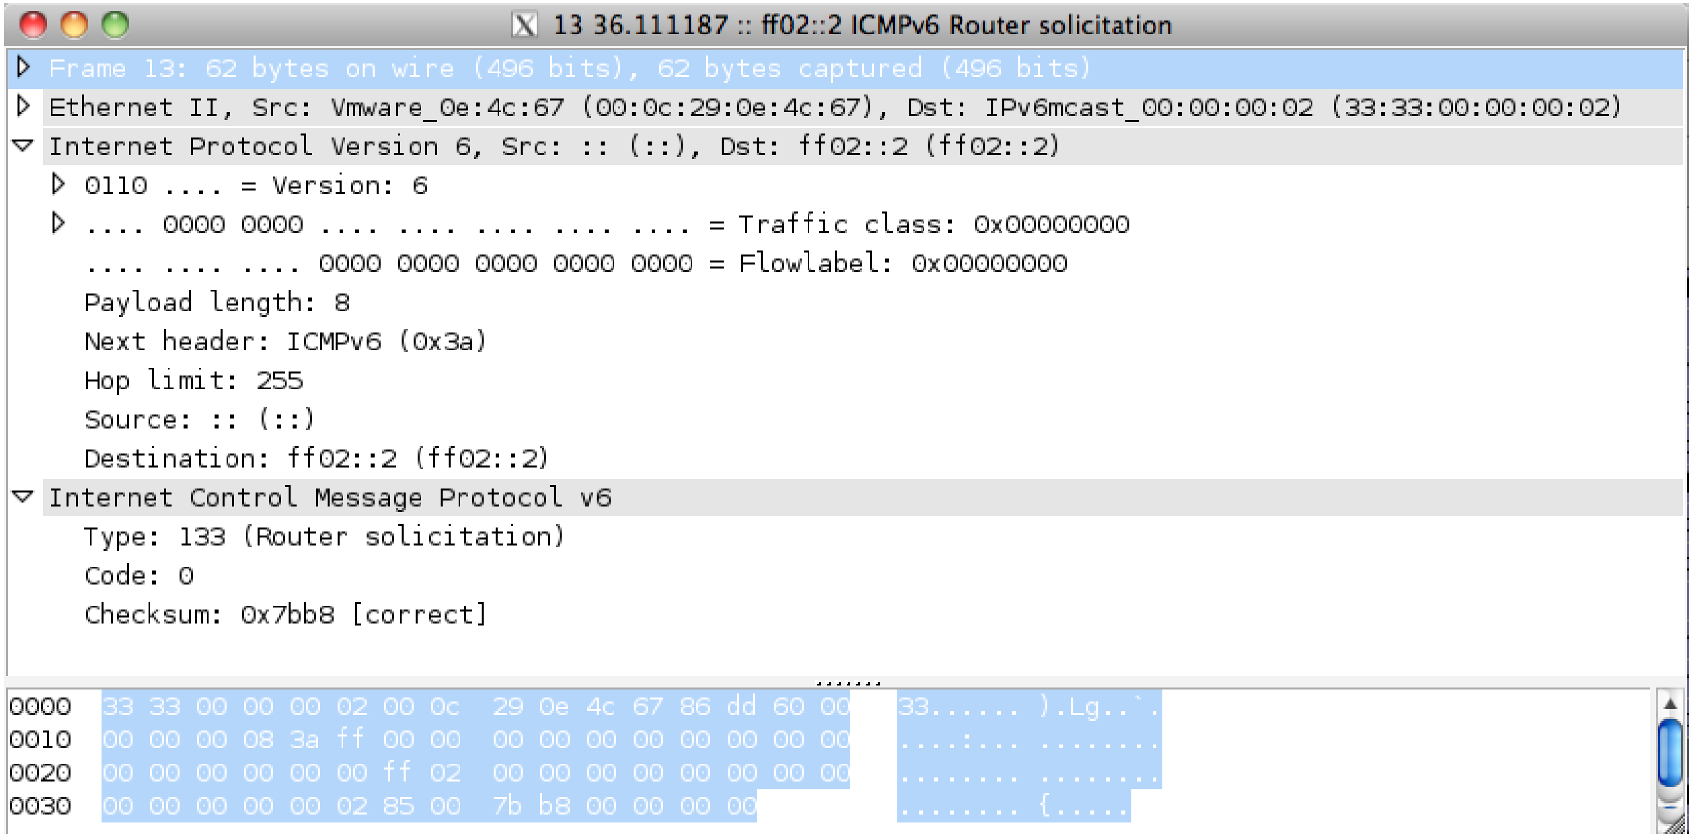

13 36.111187 :: ff02::2 ICMPv6 Router solicitation

```
▷ Frame 13: 62 bytes on wire (496 bits), 62 bytes captured (496 bits)
▷ Ethernet II, Src: Vmware_0e:4c:67 (00:0c:29:0e:4c:67), Dst: IPv6mcast_00:00:00:02 (33:33:00:00:00:02)
▽ Internet Protocol Version 6, Src: :: (::), Dst: ff02::2 (ff02::2)
  ▷ 0110 .... = Version: 6
  ▷ .... 0000 0000 .... .... .... .... .... = Traffic class: 0x00000000
    .... .... .... 0000 0000 0000 0000 0000 = Flowlabel: 0x00000000
    Payload length: 8
    Next header: ICMPv6 (0x3a)
    Hop limit: 255
    Source: :: (::)
    Destination: ff02::2 (ff02::2)
▽ Internet Control Message Protocol v6
    Type: 133 (Router solicitation)
    Code: 0
    Checksum: 0x7bb8 [correct]

0000  33 33 00 00 00 02 00 0c  29 0e 4c 67 86 dd 60 00   33...... ).Lg..`.
0010  00 00 00 08 3a ff 00 00  00 00 00 00 00 00 00 00   ....:... ........
0020  00 00 00 00 00 00 ff 02  00 00 00 00 00 00 00 00   ........ ........
0030  00 00 00 00 00 02 85 00  7b b8 00 00 00 00         ........ {.....
```

ROUTER SOLICITATION FORMAT

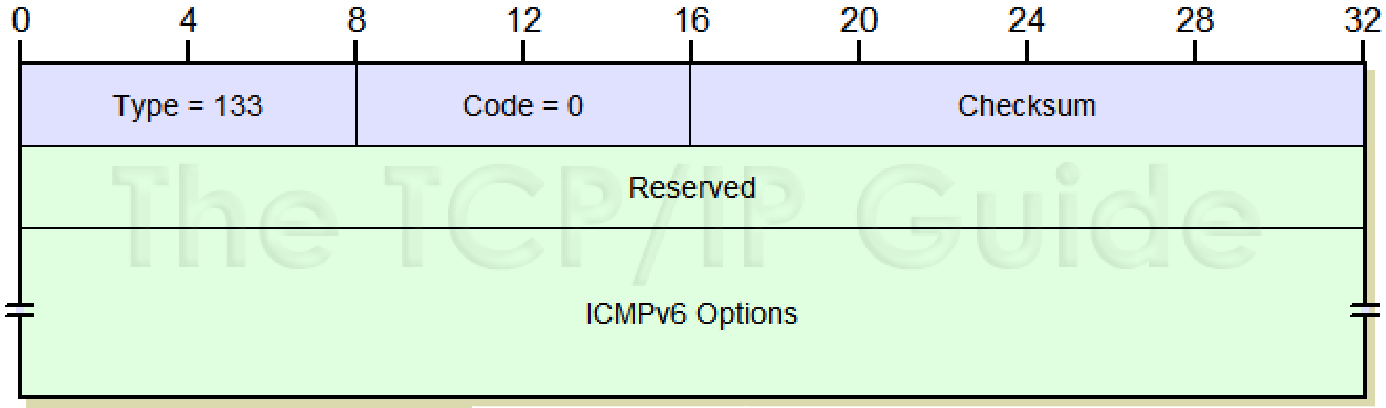

Field Name	Description
Type	Identifies the ICMPv6 message type; for Router Solicitation messages the value is 133.
Code	Not used; set to 0.
Checksum	16-bit checksum field for the ICMP header
Reserved	4 reserved bytes set to 0.
Options	If the device sending the *Router Solicitation* knows its layer two address, it should be included in a *Source Link-Layer Address* option

ROUTER ADVERTISEMENT (RA)

What Is IPv6 Router Advertisement (RA)?

Router advertisements are also sent out **periodically** and also **in response** to router solicitation messages from IPv6 nodes on the link.

- RA messages are **always** originated by **routers**.
- RA messages are used to **indicate the presence** of the Router on a link.
- RA message carry **link-specific parameters** which the hosts on the link should use for their **network parameters configuration**.
- RA messages are sent **periodically on a link** and also **sent in response** to a Router Solicitation message from a host.
- RA messages are sent to the **all-nodes link-local multicast address (FF02 ::1) OR** the **unicast IPv6 address of a node that sent the RS messages**.
- Carries a value of **134** in the Type field of the ICMP packet header

ROUTER ADVERTISEMENT PACKET

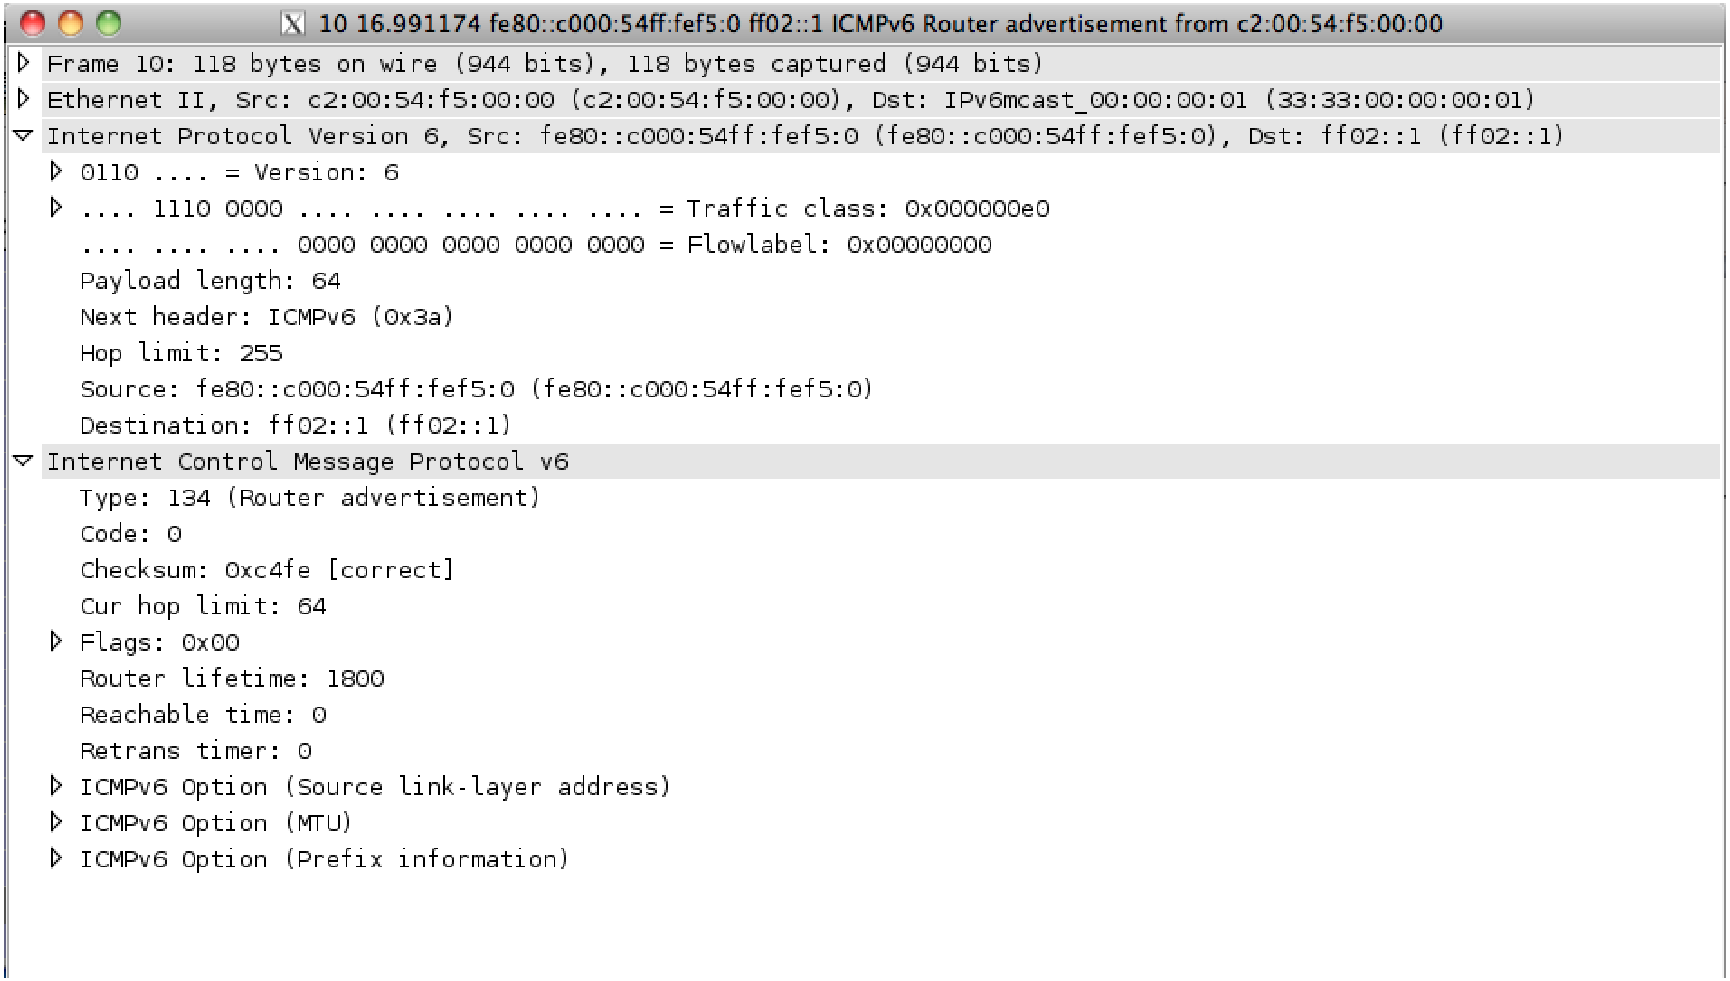

```
10 16.991174 fe80::c000:54ff:fef5:0 ff02::1 ICMPv6 Router advertisement from c2:00:54:f5:00:00

▷ Frame 10: 118 bytes on wire (944 bits), 118 bytes captured (944 bits)
▷ Ethernet II, Src: c2:00:54:f5:00:00 (c2:00:54:f5:00:00), Dst: IPv6mcast_00:00:00:01 (33:33:00:00:00:01)
▽ Internet Protocol Version 6, Src: fe80::c000:54ff:fef5:0 (fe80::c000:54ff:fef5:0), Dst: ff02::1 (ff02::1)
  ▷ 0110 .... = Version: 6
  ▷ .... 1110 0000 .... .... .... .... .... = Traffic class: 0x000000e0
    .... .... .... 0000 0000 0000 0000 0000 = Flowlabel: 0x00000000
    Payload length: 64
    Next header: ICMPv6 (0x3a)
    Hop limit: 255
    Source: fe80::c000:54ff:fef5:0 (fe80::c000:54ff:fef5:0)
    Destination: ff02::1 (ff02::1)
▽ Internet Control Message Protocol v6
    Type: 134 (Router advertisement)
    Code: 0
    Checksum: 0xc4fe [correct]
    Cur hop limit: 64
  ▷ Flags: 0x00
    Router lifetime: 1800
    Reachable time: 0
    Retrans timer: 0
  ▷ ICMPv6 Option (Source link-layer address)
  ▷ ICMPv6 Option (MTU)
  ▷ ICMPv6 Option (Prefix information)
```

ROUTER ADVERTISEMENT FORMAT

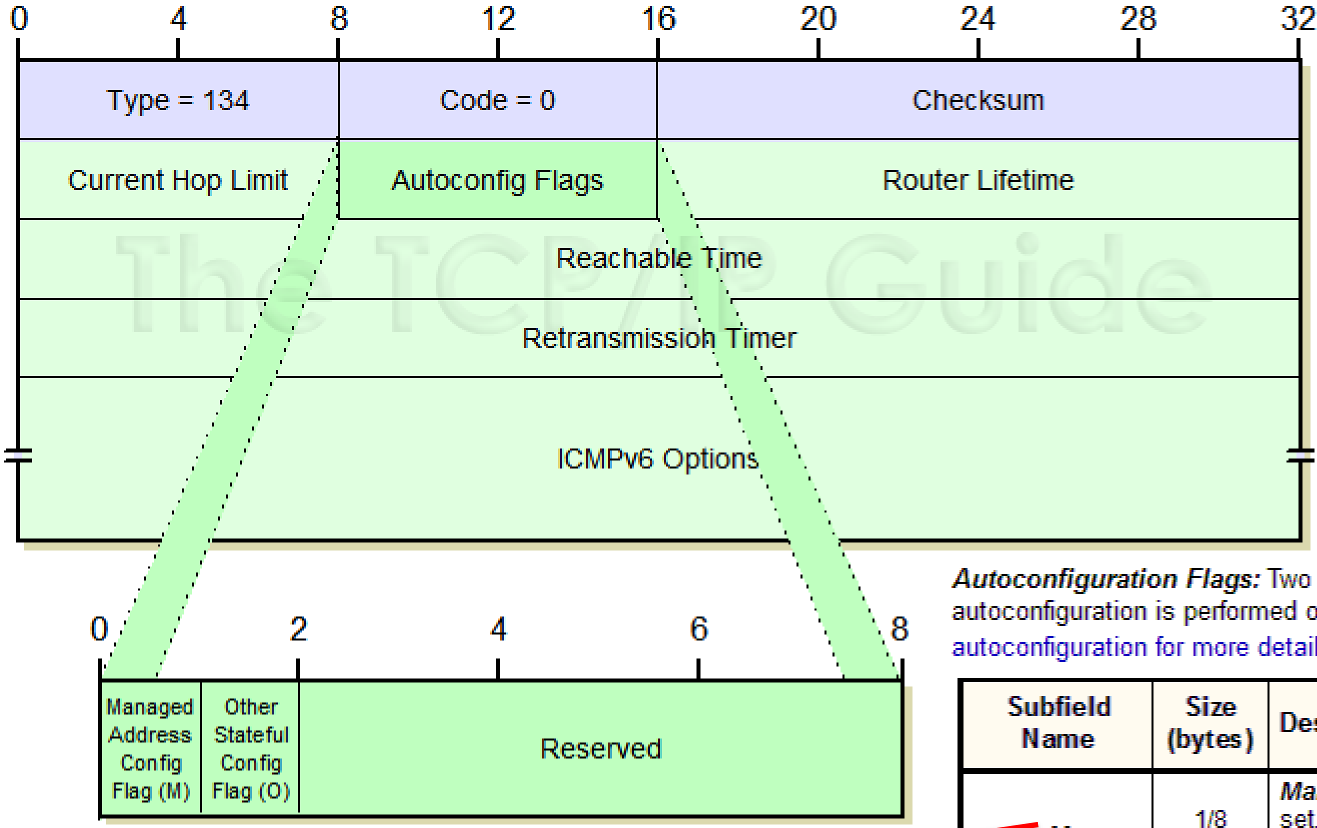

Autoconfiguration Flags: Two flags that let the router tell the host how autoconfiguration is performed on the local network. See the topic on IPv6 autoconfiguration for more details:

Subfield Name	Size (bytes)	Description
M	1/8 (1 bit)	***Managed Address Configuration Flag:*** When set, this flag tells hosts to use an administered or "stateful" method for address autoconfiguration, such as DHCP.
O	1/8 (1 bit)	***Other Stateful Configuration Flag:*** When set, tells hosts to use an administered or "stateful" autoconfiguration method for information other than addresses.
Reserved	6/8 (6 bits)	***Reserved:*** Reserved for future use; sent as zeroes.

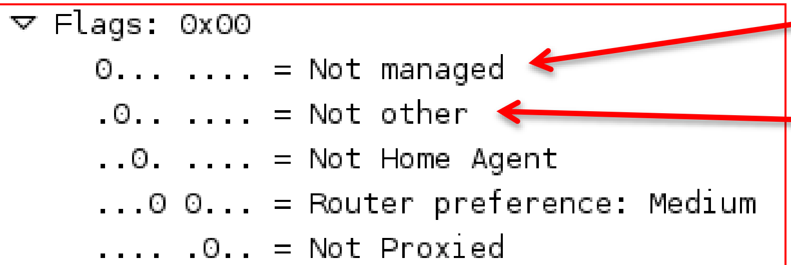

```
▽ Flags: 0x00
    0... .... = Not managed
    .0.. .... = Not other
    ..0. .... = Not Home Agent
    ...0 0... = Router preference: Medium
    .... .0.. = Not Proxied
```

ROUTER ADVERTISEMENT OPTIONS

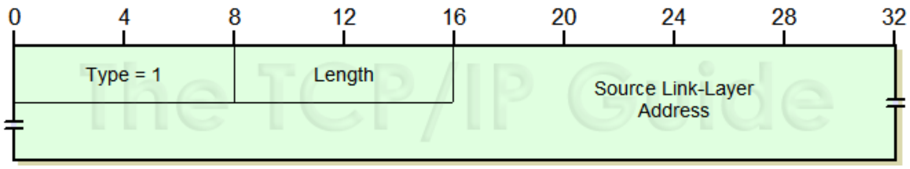

Type 1
ICMPv6 Source Link-Layer Address Option Format

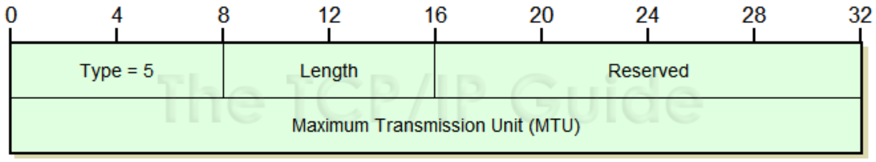

Type 5
ICMPv6 MTU Option Format

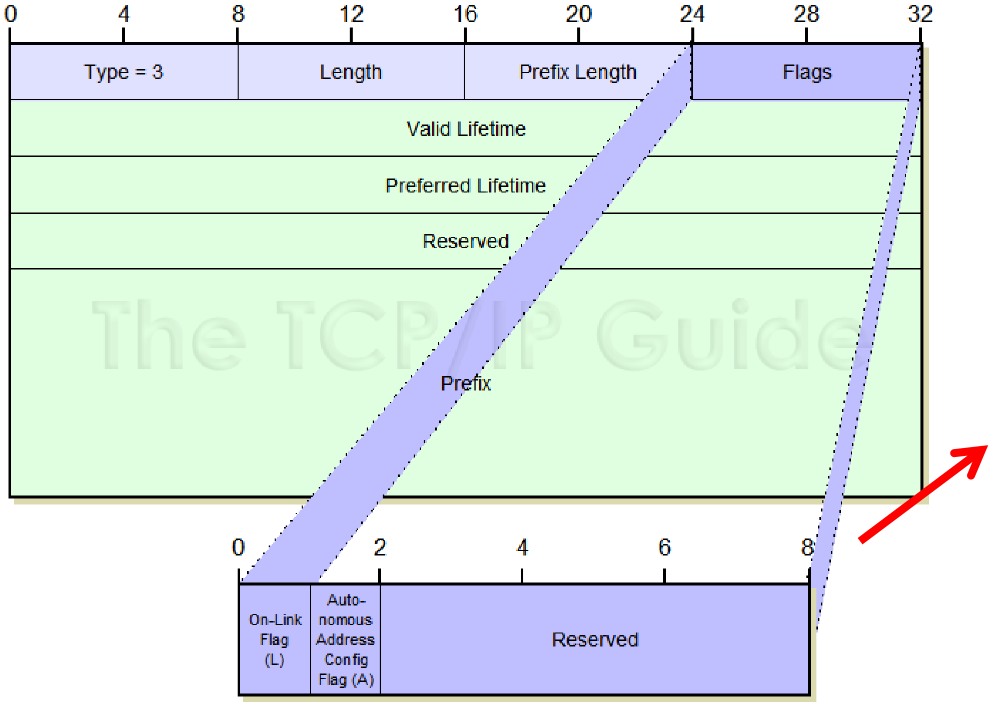

Type 3
ICMPv6 Prefix Information Option Format

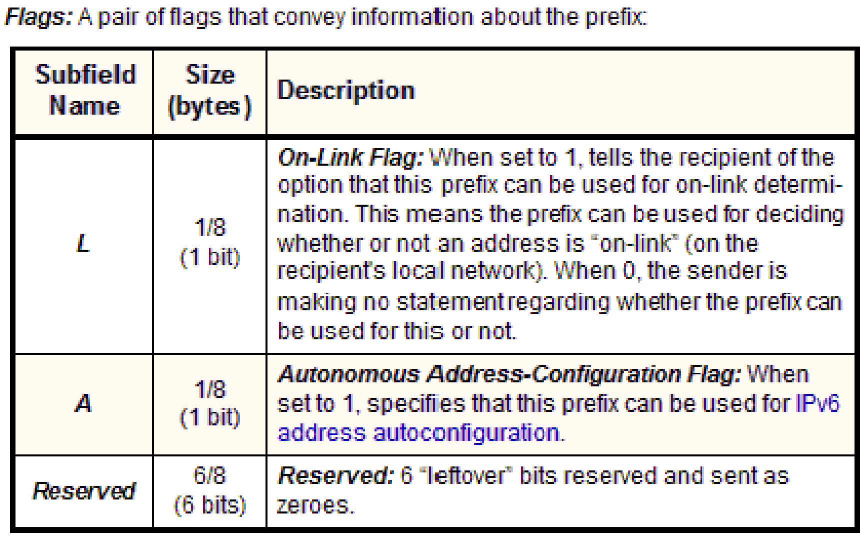

Flags: A pair of flags that convey information about the prefix:

Subfield Name	Size (bytes)	Description
L	1/8 (1 bit)	***On-Link Flag:*** When set to 1, tells the recipient of the option that this prefix can be used for on-link determination. This means the prefix can be used for deciding whether or not an address is "on-link" (on the recipient's local network). When 0, the sender is making no statement regarding whether the prefix can be used for this or not.
A	1/8 (1 bit)	***Autonomous Address-Configuration Flag:*** When set to 1, specifies that this prefix can be used for IPv6 address autoconfiguration.
Reserved	6/8 (6 bits)	***Reserved:*** 6 "leftover" bits reserved and sent as zeroes.

NEIGHBOR SOLICITATION (NS)

What Is IPv6 Neighbor Solicitation?

Neighbor solicitation messages are sent on the local link when a node wants to determine the link-layer address of another node on the same local link.

- NS messages are **originated** by the **nodes**. NS messages is used to **request the link layer address** of another node. NS messages are also used for **duplicate address detection** and **neighbor unreachability detection**.
- This function is similar to the ARP in IPv4, **but avoids broadcasts** used in IPv4 ARP messages.
- Carries a value of **135** in the Type field of the ICMP packet header

NEIGHBOR SOLICITATION PACKET

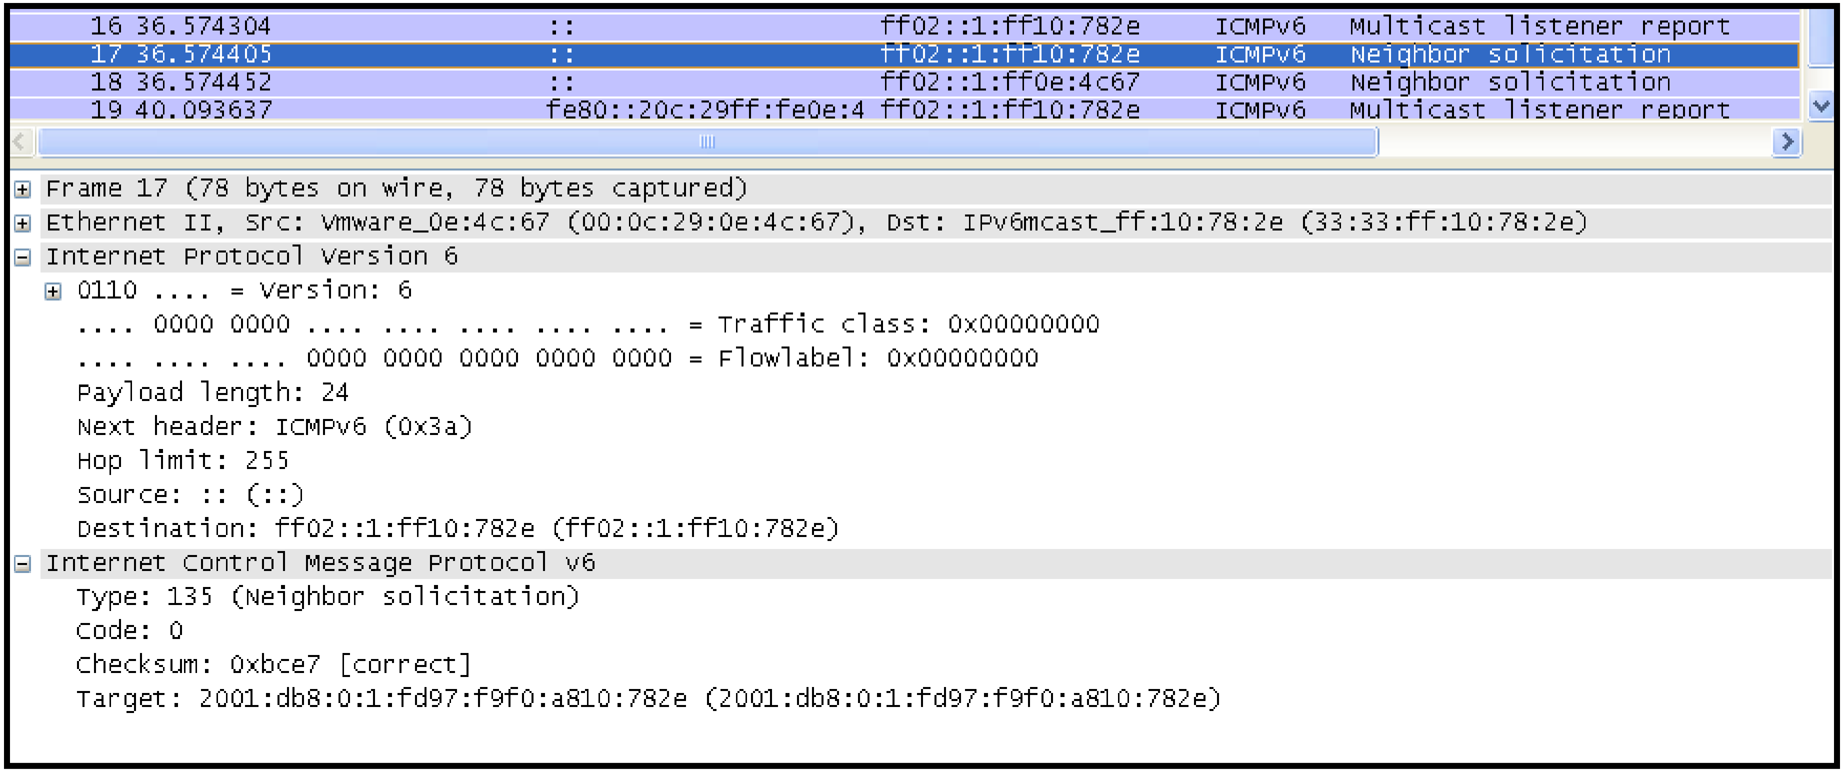

```
16 36.574304  ::                    ff02::1:ff10:782e   ICMPv6   Multicast listener report
17 36.574405  ::                    ff02::1:ff10:782e   ICMPv6   Neighbor solicitation
18 36.574452  ::                    ff02::1:ff0e:4c67   ICMPv6   Neighbor solicitation
19 40.093637  fe80::20c:29ff:fe0e:4 ff02::1:ff10:782e   ICMPv6   Multicast listener report
```

```
⊞ Frame 17 (78 bytes on wire, 78 bytes captured)
⊞ Ethernet II, Src: Vmware_0e:4c:67 (00:0c:29:0e:4c:67), Dst: IPv6mcast_ff:10:78:2e (33:33:ff:10:78:2e)
⊟ Internet Protocol Version 6
  ⊞ 0110 .... = Version: 6
    .... 0000 0000 .... .... .... .... .... = Traffic class: 0x00000000
    .... .... .... 0000 0000 0000 0000 0000 = Flowlabel: 0x00000000
    Payload length: 24
    Next header: ICMPv6 (0x3a)
    Hop limit: 255
    Source: :: (::)
    Destination: ff02::1:ff10:782e (ff02::1:ff10:782e)
⊟ Internet Control Message Protocol v6
    Type: 135 (Neighbor solicitation)
    Code: 0
    Checksum: 0xbce7 [correct]
    Target: 2001:db8:0:1:fd97:f9f0:a810:782e (2001:db8:0:1:fd97:f9f0:a810:782e)
```

NEIGHBOR SOLICITATION OPTIONS

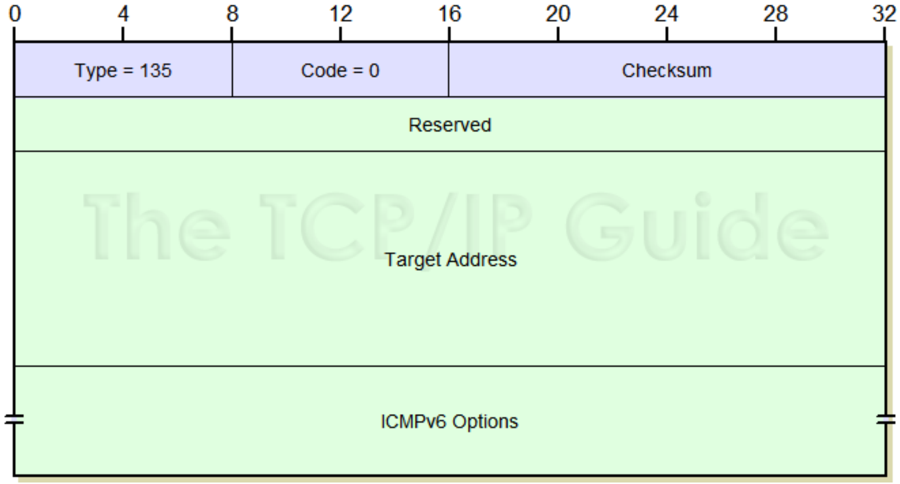

Neighbor solicitation (NS) option formats

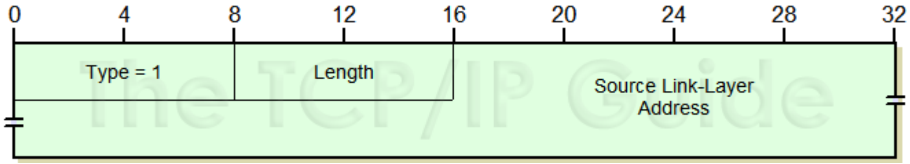

Type 1

ICMPv6 Source Link-Layer Address Option Format

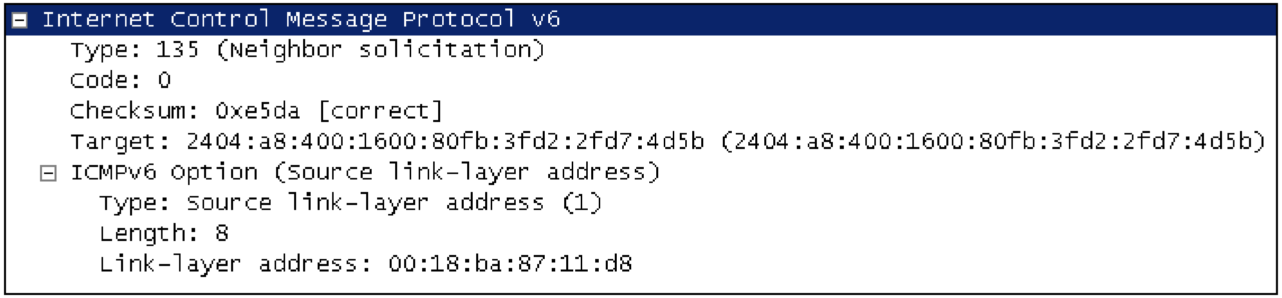

```
⊟ Internet Control Message Protocol v6
    Type: 135 (Neighbor solicitation)
    Code: 0
    Checksum: 0xe5da [correct]
    Target: 2404:a8:400:1600:80fb:3fd2:2fd7:4d5b (2404:a8:400:1600:80fb:3fd2:2fd7:4d5b)
  ⊟ ICMPv6 Option (Source link-layer address)
      Type: Source link-layer address (1)
      Length: 8
      Link-layer address: 00:18:ba:87:11:d8
```

USING NS FOR D.A.D.

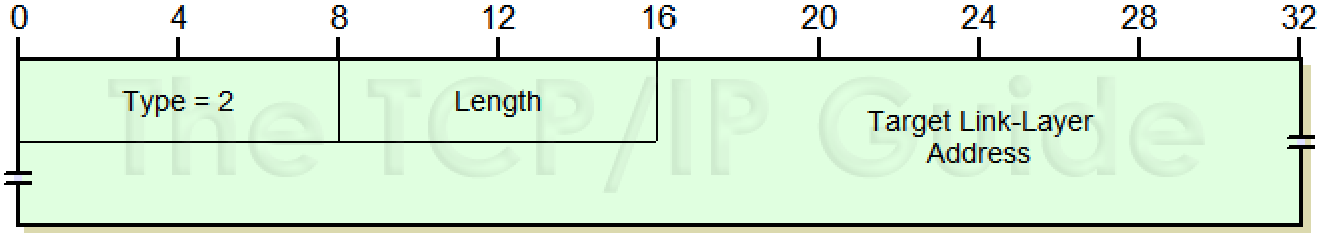

Type 2

ICMPv6 Target Link-Layer Address Option Format

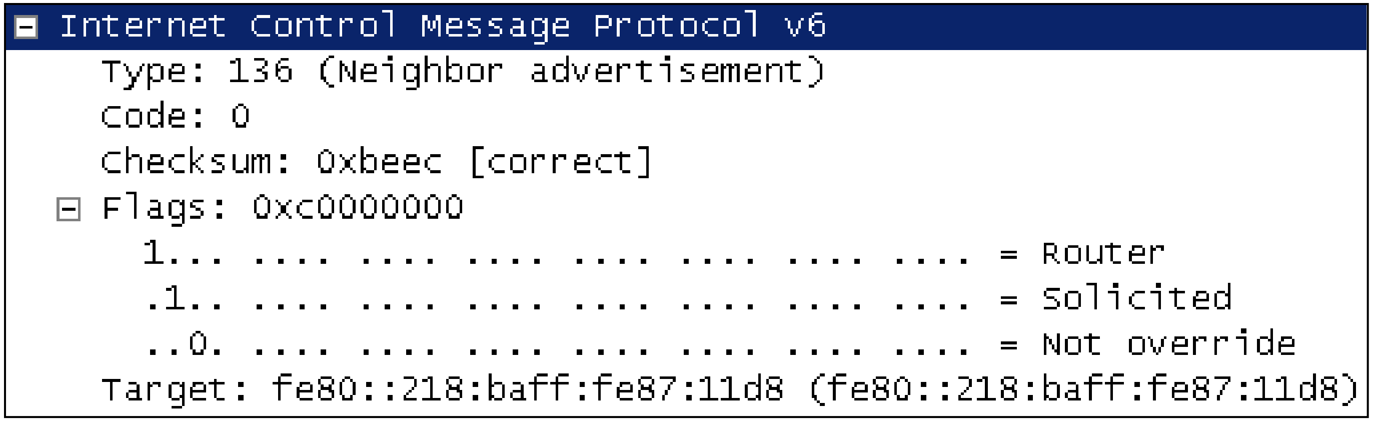

```
⊟ Internet Control Message Protocol v6
    Type: 136 (Neighbor advertisement)
    Code: 0
    Checksum: 0xbeec [correct]
  ⊟ Flags: 0xc0000000
      1... .... .... .... .... .... .... .... = Router
      .1.. .... .... .... .... .... .... .... = Solicited
      ..0. .... .... .... .... .... .... .... = Not override
    Target: fe80::218:baff:fe87:11d8 (fe80::218:baff:fe87:11d8)
```

NEIGHBOR ADVERTISEMENT (NA)

What Is IPv6 Neighbor Advertisement?

The IPv6 neighbor advertisement message is a response to the IPv6 neighbor solicitation message.

- NA messages **are almost always sent** in response to an NS message from a node.
- NA messages can be sent by a node when its **link-layer address is changed**. This NA message is sent as an **unsolicited NA** to advertise its new address.
- After receiving the NA, the **source node** and **destination node** can communicate.
- Carries a value of **136** in the Type field of the ICMP packet header.

NEIGHBOR ADVERTISEMENT PACKET

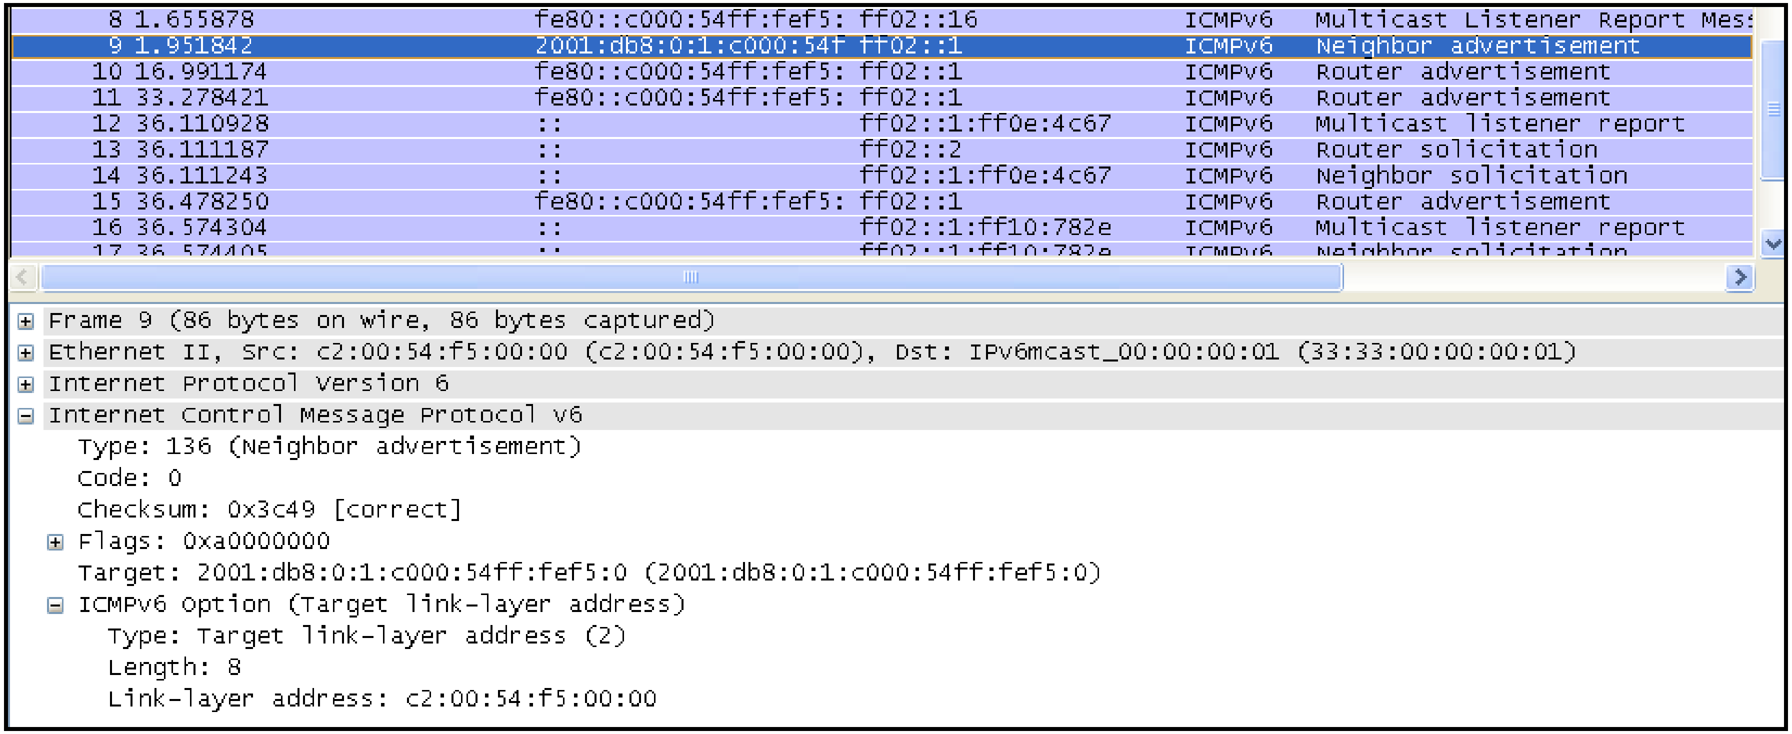

```
     8 1.655878        fe80::c000:54ff:fef5: ff02::16          ICMPv6   Multicast Listener Report Mes:
     9 1.951842        2001:db8:0:1:c000:54f ff02::1           ICMPv6   Neighbor advertisement
    10 16.991174       fe80::c000:54ff:fef5: ff02::1           ICMPv6   Router advertisement
    11 33.278421       fe80::c000:54ff:fef5: ff02::1           ICMPv6   Router advertisement
    12 36.110928       ::                    ff02::1:ff0e:4c67 ICMPv6   Multicast listener report
    13 36.111187       ::                    ff02::2           ICMPv6   Router solicitation
    14 36.111243       ::                    ff02::1:ff0e:4c67 ICMPv6   Neighbor solicitation
    15 36.478250       fe80::c000:54ff:fef5: ff02::1           ICMPv6   Router advertisement
    16 36.574304       ::                    ff02::1:ff10:782e ICMPv6   Multicast listener report
    17 36.574405       ::                    ff02::1:ff10:782e ICMPv6   Neighbor solicitation

⊞ Frame 9 (86 bytes on wire, 86 bytes captured)
⊞ Ethernet II, Src: c2:00:54:f5:00:00 (c2:00:54:f5:00:00), Dst: IPv6mcast_00:00:00:01 (33:33:00:00:00:01)
⊞ Internet Protocol Version 6
⊟ Internet Control Message Protocol v6
    Type: 136 (Neighbor advertisement)
    Code: 0
    Checksum: 0x3c49 [correct]
  ⊞ Flags: 0xa0000000
    Target: 2001:db8:0:1:c000:54ff:fef5:0 (2001:db8:0:1:c000:54ff:fef5:0)
  ⊟ ICMPv6 Option (Target link-layer address)
      Type: Target link-layer address (2)
      Length: 8
      Link-layer address: c2:00:54:f5:00:00
```

NEIGHBOR ADVERTISEMENT OPTIONS

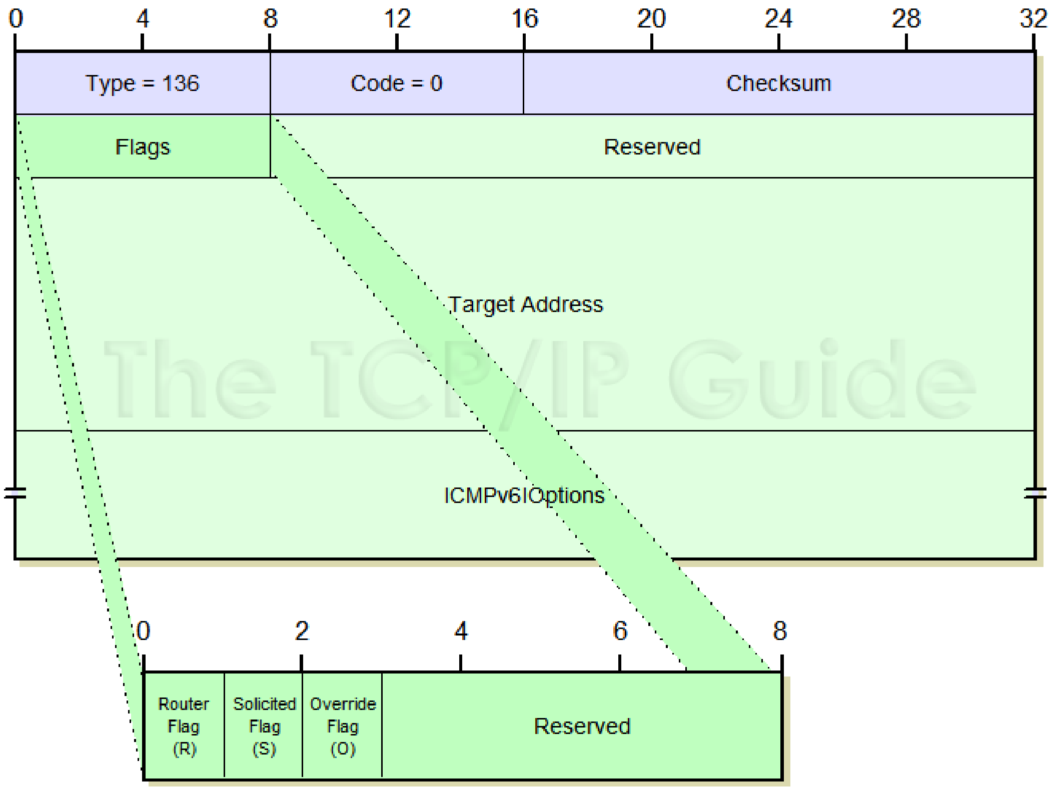

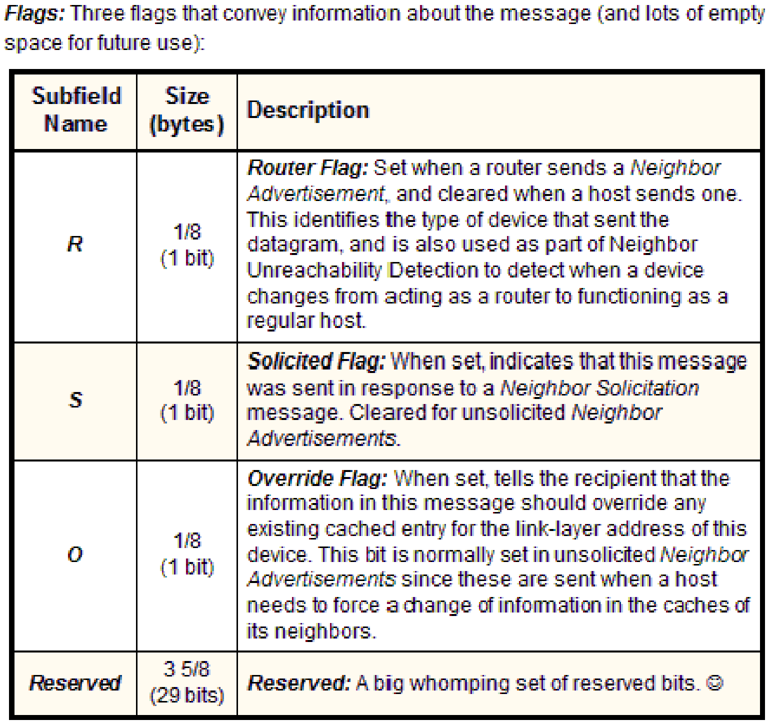

Flags: Three flags that convey information about the message (and lots of empty space for future use):

Subfield Name	Size (bytes)	Description
R	1/8 (1 bit)	***Router Flag:*** Set when a router sends a *Neighbor Advertisement*, and cleared when a host sends one. This identifies the type of device that sent the datagram, and is also used as part of Neighbor Unreachability Detection to detect when a device changes from acting as a router to functioning as a regular host.
S	1/8 (1 bit)	***Solicited Flag:*** When set, indicates that this message was sent in response to a *Neighbor Solicitation* message. Cleared for unsolicited *Neighbor Advertisements*.
O	1/8 (1 bit)	***Override Flag:*** When set, tells the recipient that the information in this message should override any existing cached entry for the link-layer address of this device. This bit is normally set in unsolicited *Neighbor Advertisements* since these are sent when a host needs to force a change of information in the caches of its neighbors.
Reserved	3 5/8 (29 bits)	***Reserved:*** A big whomping set of reserved bits. ☺

REDIRECT

A fifth type of ICMPv6 message, the **Redirect** (type 137), is used by routers to either **point hosts toward a more preferable router**, or to indicate that the destination actually resides on link.

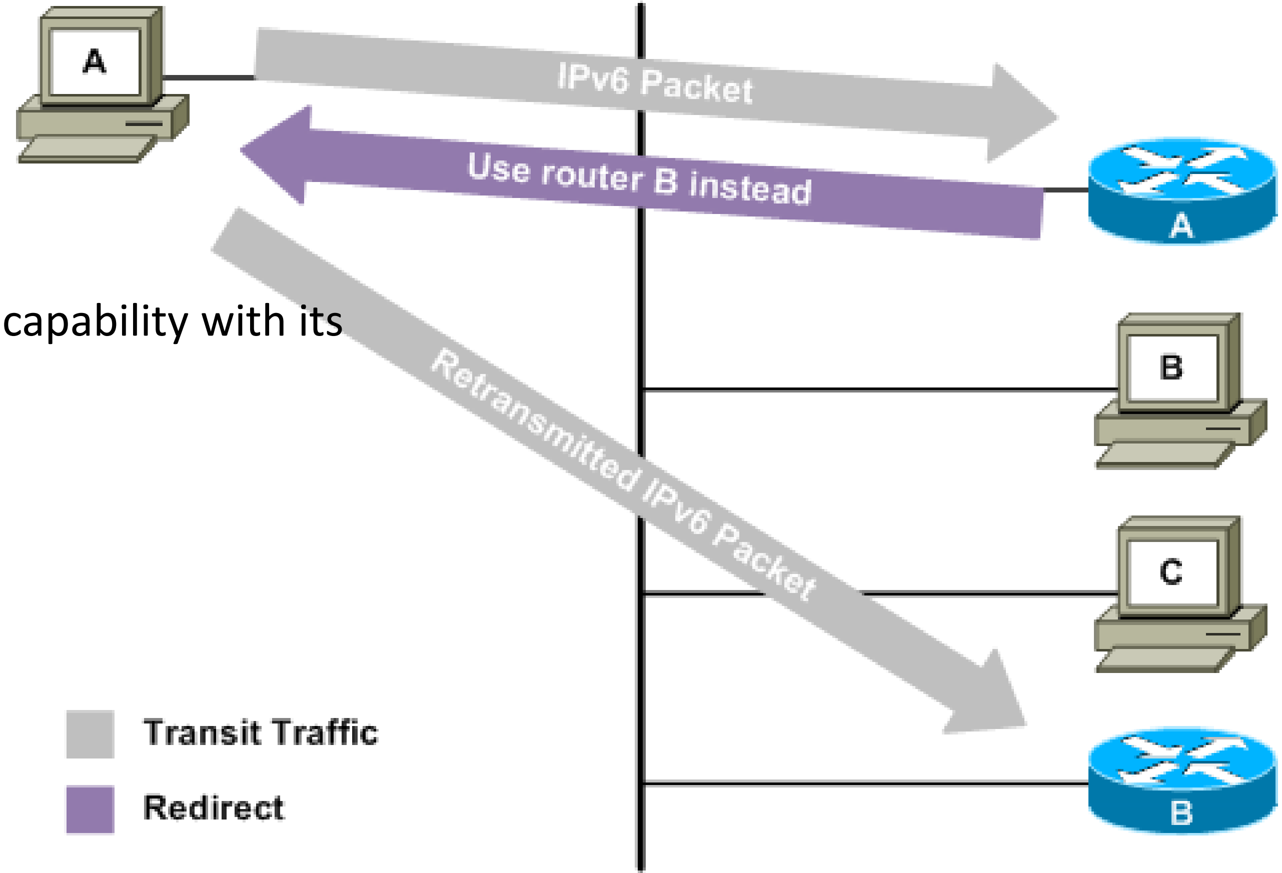

ICMPv4 provides the same capability with its own redirect message

REDIRECT HEADER

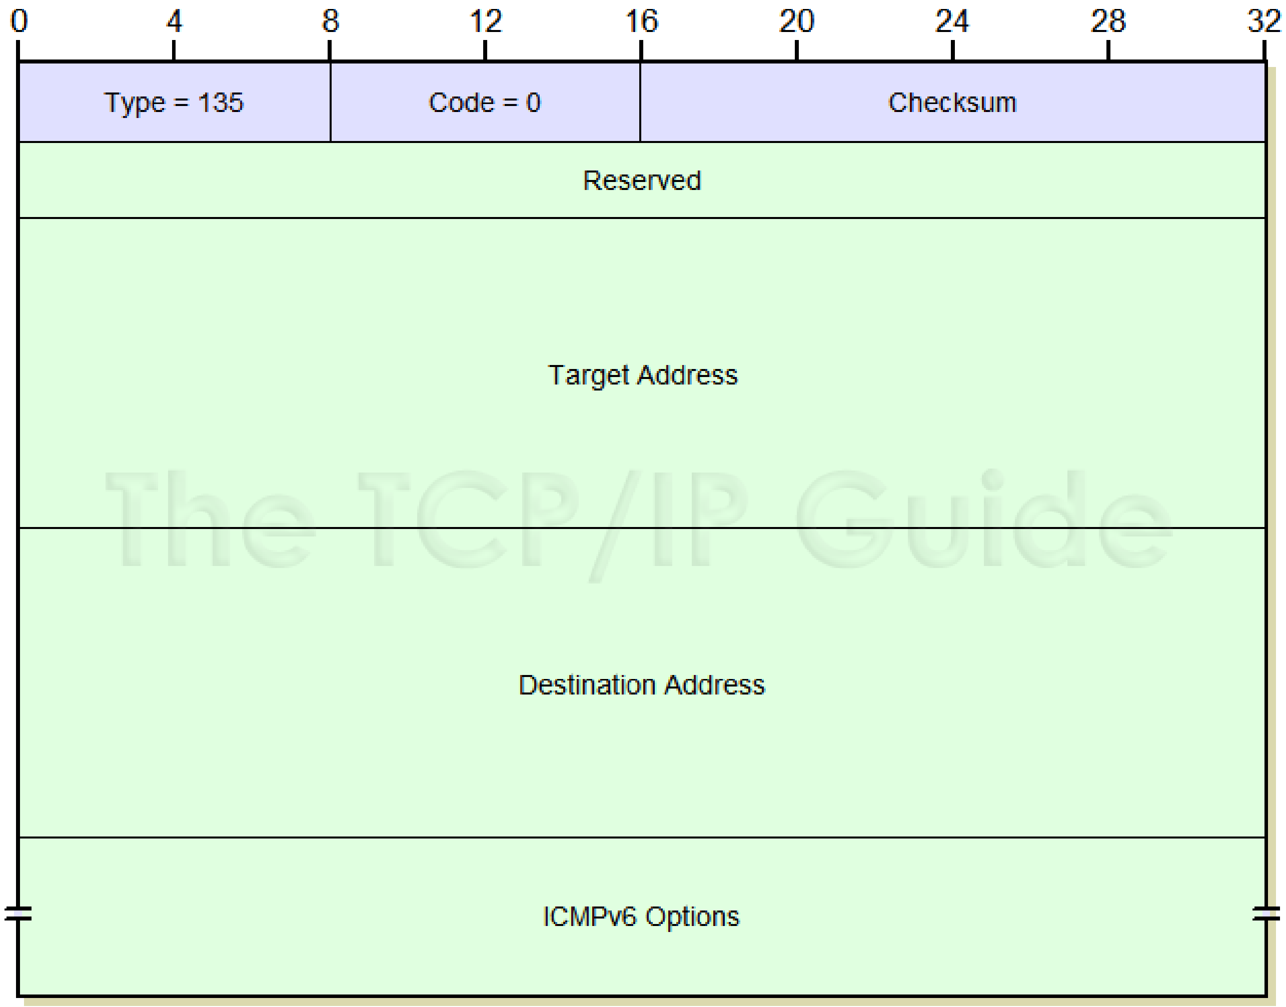

Type = 135

Code = 0

Checksum

Reserved

Target Address

Destination Address

ICMPv6 Options

REDIRECT EXAMPLE

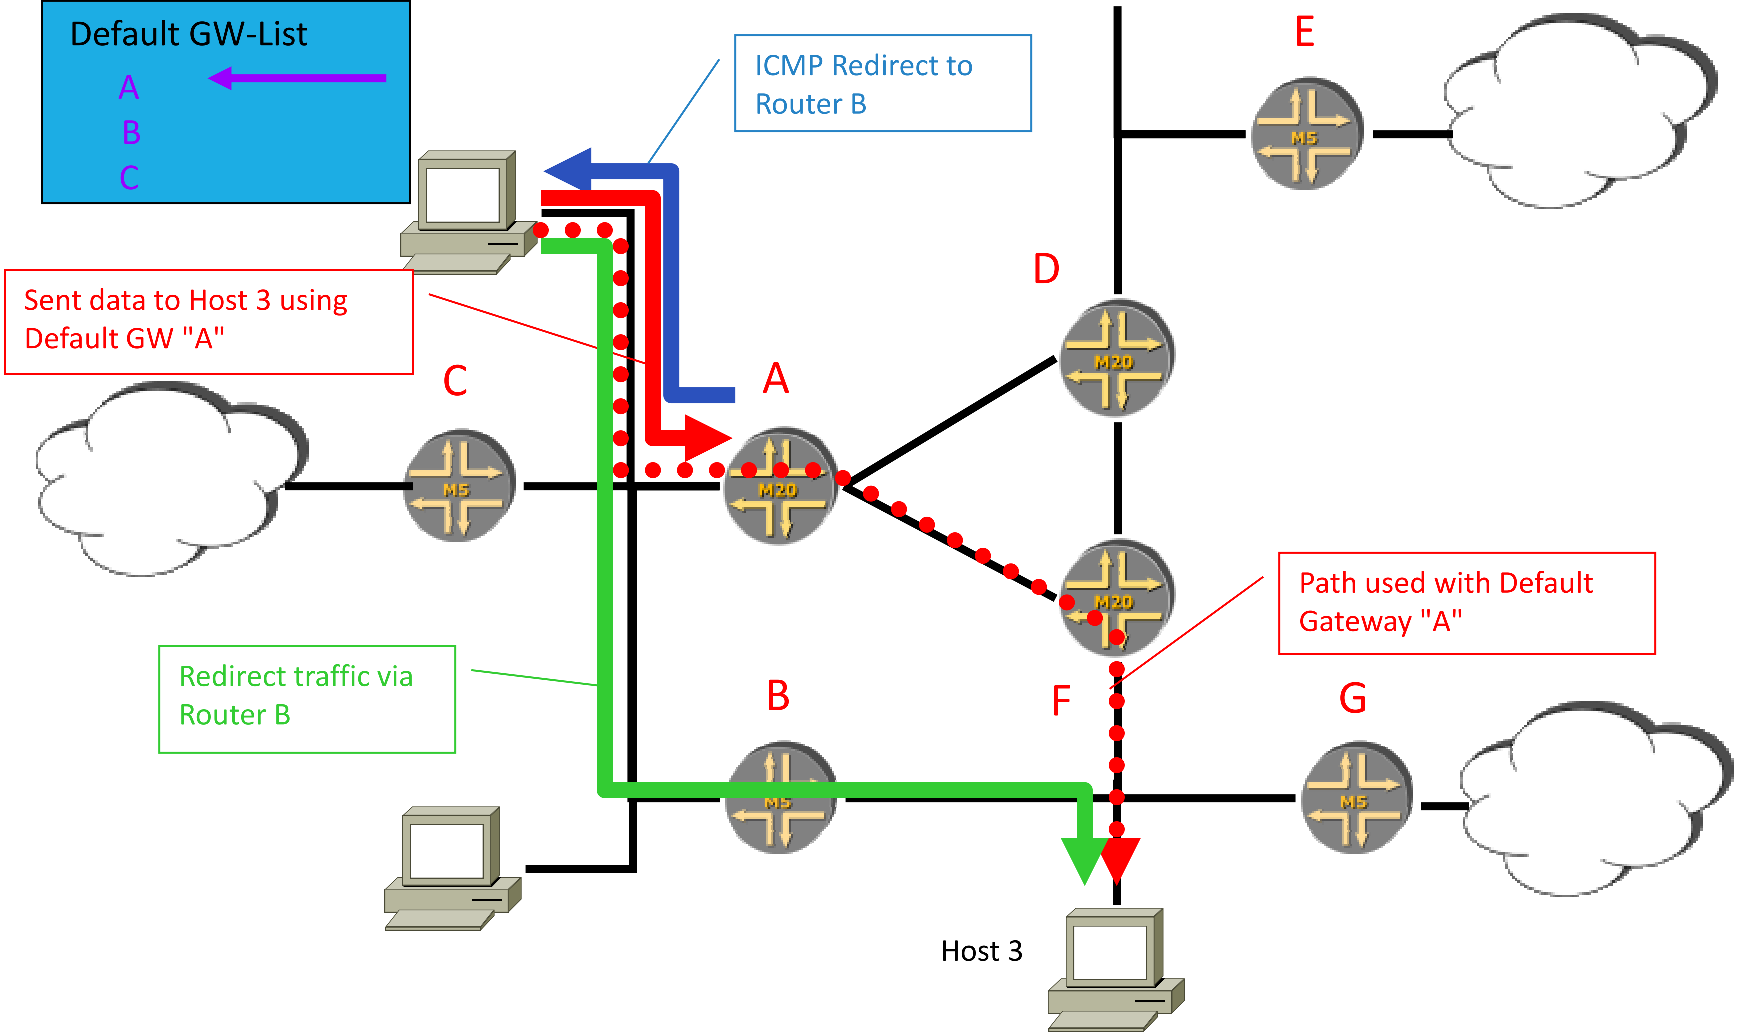

Default GW-List

A

B

C

ICMP Redirect to Router B

Sent data to Host 3 using Default GW "A"

Redirect traffic via Router B

Path used with Default Gateway "A"

UNDERSTANDING NDP PROCESSES

NDP Processes

The ND process perform functions for IPv6 **similar** to:

- Address Resolution Protocol (ARP)
- ICMP Router Discovery and Router Redirect protocols for IPv4

NDP PROCESSES

Router Discovery

Whereas IPv4 hosts must rely on manual configuration or DHCP to provide the address of a default gateway, IPv6 hosts can automatically locate default routers on the link. This is accomplished through the use of **2 ICMPv6 messages**:

Router Solicitation (type 133) and **Router Advertisement (type 134)**. When first joining a link, an IPv6 host multicast a router solicitation to the *all routers* multicast group (FF02::2), and each router active on the link responds by sending a router advertisement with its address to the *all nodes* group (FF02::1).

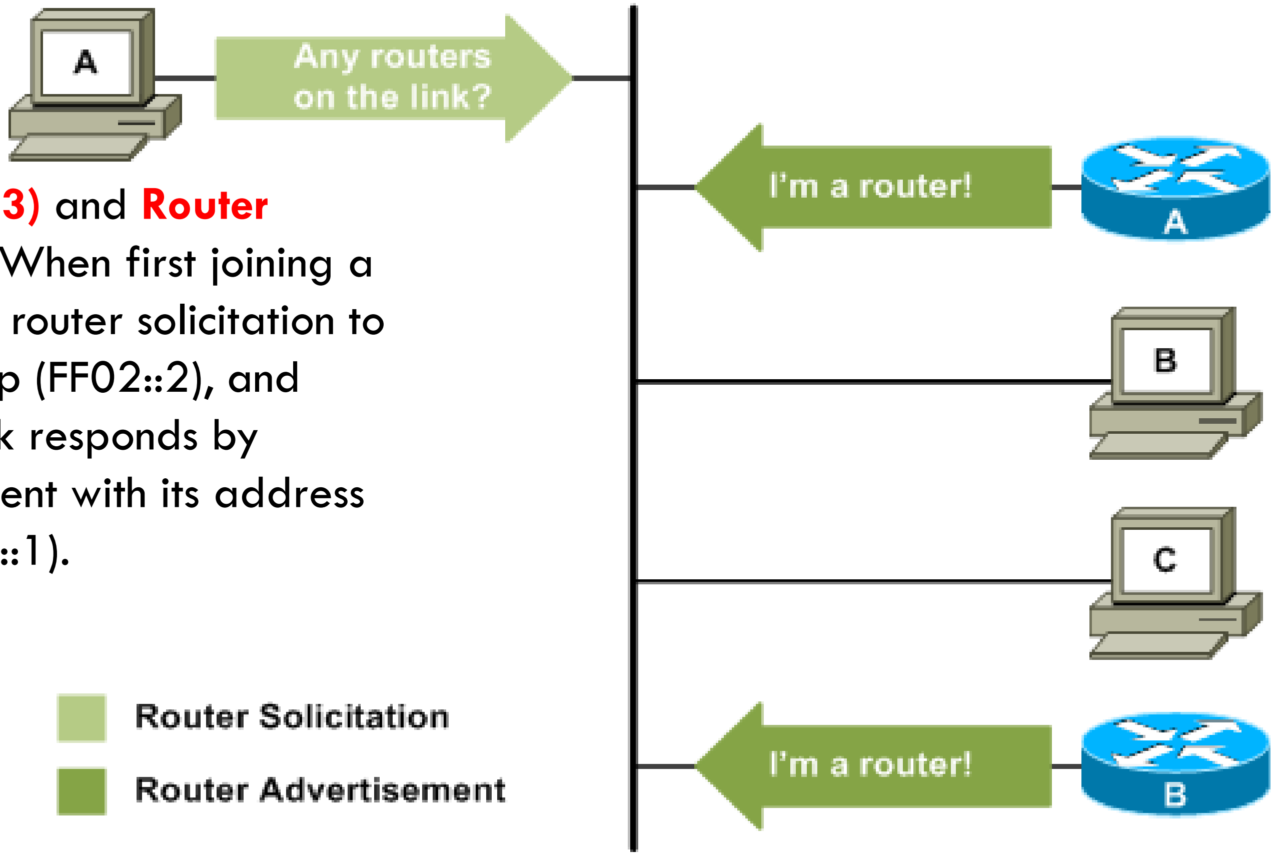

PREFIX DISCOVERY USING RA&RS

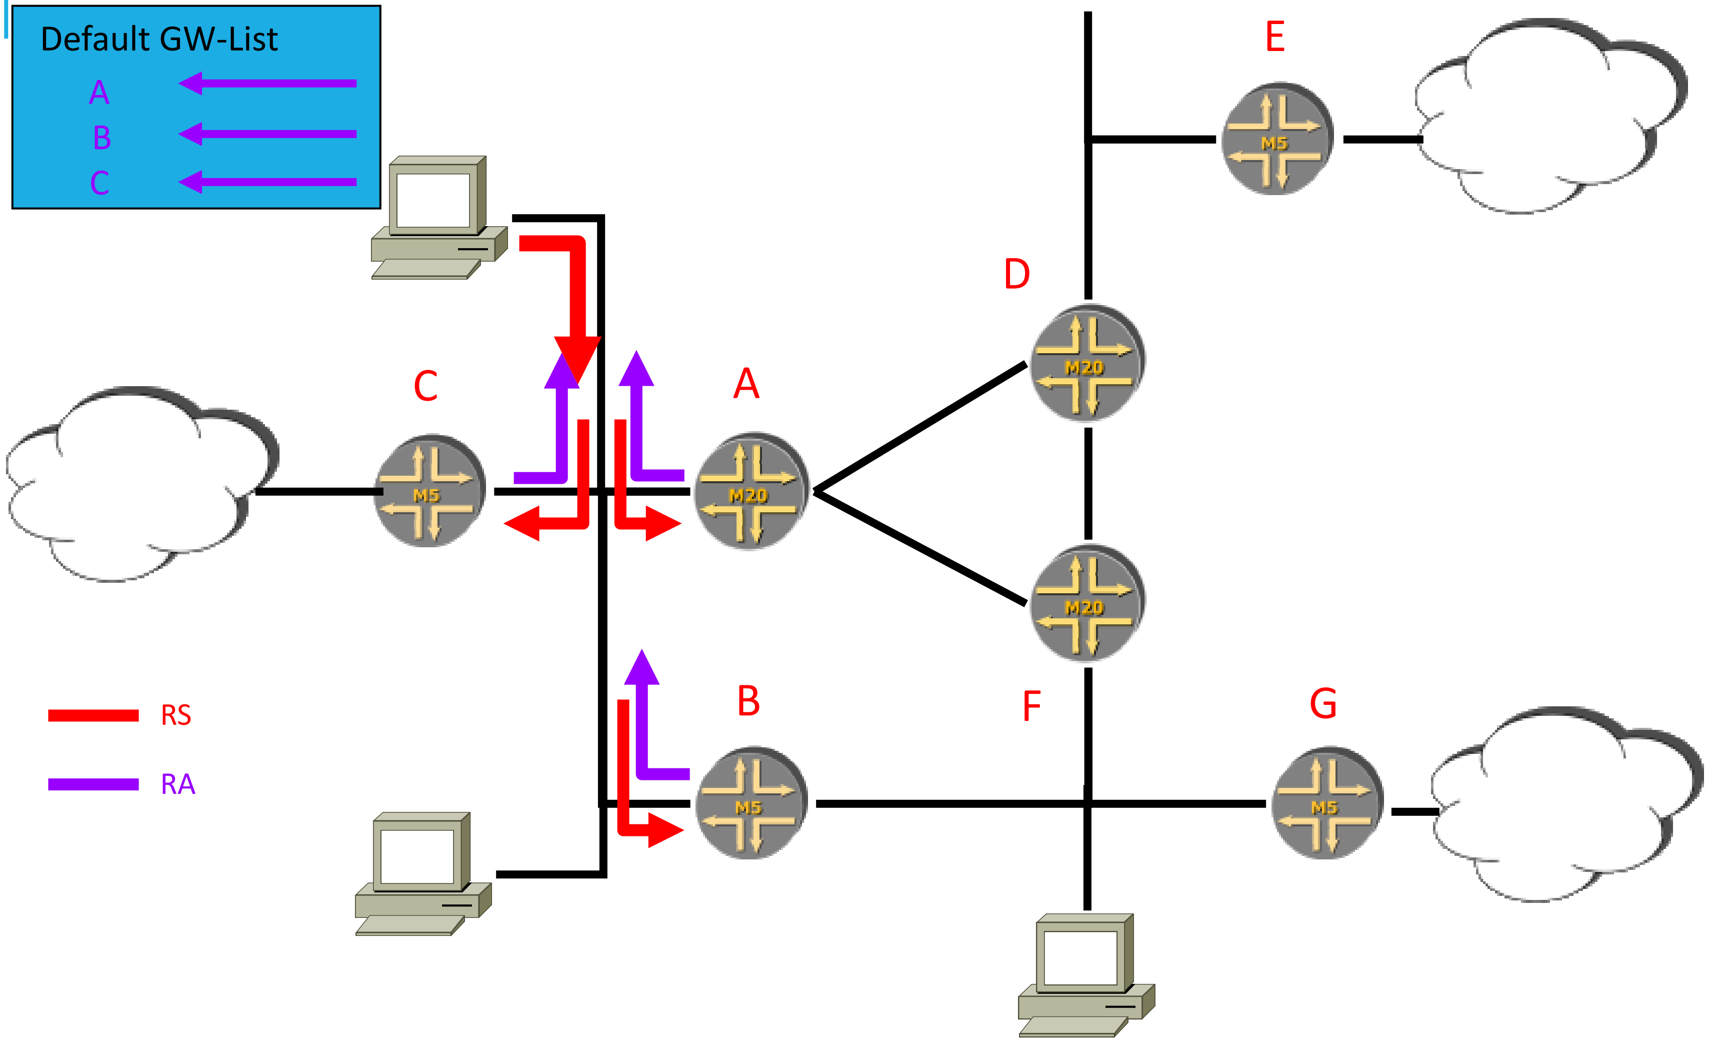

Default GW-List

- A
- B
- C

NDP PROCESSES

Prefix Discovery

- One of the options typically carried by a RA is the **Prefix Information** option (type 3).
- Each prefix information option lists an IPv6 prefix (subnet) reachable on the local link.
- It is **NOT uncommon** for **multiple IPv6 prefixes to reside on the same link**, and routers **may include more than one prefix** in each advertisement.
- A host which knows what prefixes are reachable on the link can communicate directly with destinations in those prefixes without passing its traffic through a router.

NDP PROCESSES

Parameter Discovery

- Another option included in RA is the **MTU** option (type 5), which informs hosts of the IP MTU to use.
- E.g, this value is typically set to 1500 for Ethernet networks.

(However, not all link types have a standardized MTU size. Including this option ensures all hosts know the correct MTU to use.)

- RA also specify the **default value** hosts should use **for the IPv6 hop count**.

(This isn't an option, but a field built into the router advertisement message header.)

NDP PROCESSES

Address Autoconfiguration

- NDP provides mechanisms for a **host to automatically configure itsel**f with an address from a prefix learned from a local router through prefix discovery.
- This is done by concatenating a candidate learned prefix with the **EUI-64** address of the host's interface.
- In this manner, a host can achieve **stateless autoconfiguration**.

NDP PROCESSES

Address Resolution

The function of address resolution was handled by ARP for IPv4, but is handled by ICMPv6 for IPv6. In a process very similar to router discovery, two ICMPv6 messages are used:

Neighbor Solicitation (type 135) and **Neighbor Advertisement** (type 136). A host seeking the link layer address of a neighbor multicasts a neighbor solicitation and the neighbor (if online) responds with its link layer address in a neighbor advertisement.

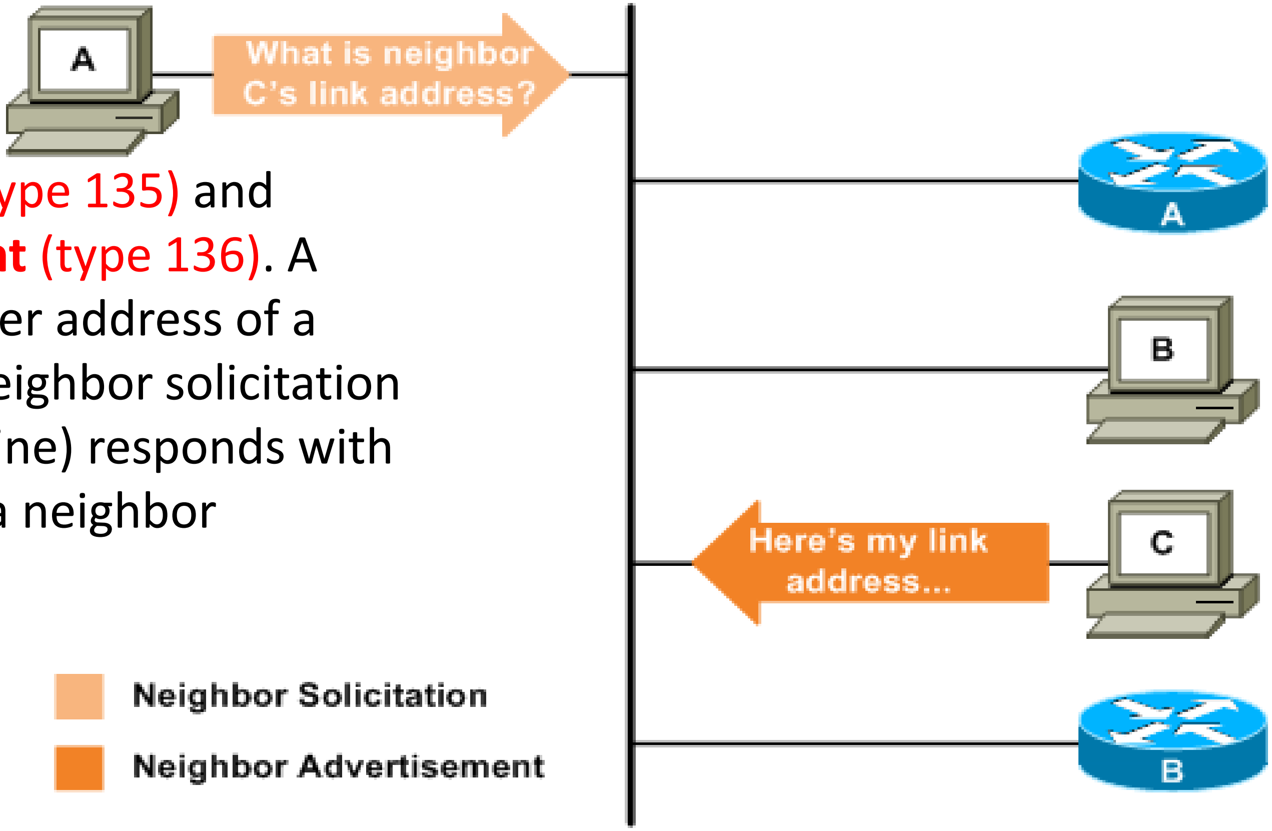

NDP PROCESSES

Next-Hop Determination

- As in IPv4, next-hop determination is simply a procedure for performing longest-match lookups on the host routing table
- And for off-link destinations, the selection of a default router.

Neighbor Unreachability Detection

- NDP is able to determine the reachability of a neighbor by examining clues from upper-layer protocols (for example, received TCP acknowledgments)
- Or by **actively re-performing address resolution** (via ICMPv6) when certain thresholds are reached.
- See Address State in next slide to have further understanding

UNDERSTANDING NEIGHBOR STATES

INCOMPLETE: Address resolution is being performed. NA not received yet.

REACHABLE: Positive confirmation was received. Within *ReachableTimer*.

STALE: *ReachableTimer* has elapsed. Also entered upon receiving unsolicited ND. Does not confirm reachability.

DELAY: *ReachableTimer* elapsed, NS sent but no confirmation (NA not received)

PROBE: Reachability confirmation is actively sought by **retransmiting NS** every *RetransTimer (ms)* until reach. Confirmation is received.

NDP PROCESSES

Duplicate Address Detection

- When **a host first joins a link**, **it send NS for its *own* IPv6 address** for a short period ***before*** attempting to use that address to communicate.
- **IF** it receives a **NA in response**, the host realizes that another neighbor on the link is already using that address.
- The host will **mark the address as a duplicate** and will not use it on the link. (similar to IPv4 gratuitous ARP requests)

DUPLICATE ADDRESS DETECTION (DAD)

MUST be performed by all nodes

Performed **before** assigning a unicast address to an interface.

Performed on interface initialization

NOT performed for anycast addresses

Link must be multicast capable

New address is called "tentative" as long as duplicate address detection takes place

DUPLICATE ADDRESS DETECTION (DAD) EXPLAINED

1. Interface joins **all-nodes** multicast group
2. Interface joins **solicited-node multicast** group
3. Node sends (one) NS with
   - **Target address** = **tentative IP address**
   - **Source address** = **unspecified (::)**
   - **Destination address** = **tentative solicited-node address**

DUPLICATE ADDRESS DETECTION (DAD) EXPLAINED

If address already exists, the particular node sends a **NA** reply with

- **Target address** = **tentative IP address**
- **Destination address** = **tentative solicited-node address**

If soliciting node **receives NA** reply with target address set to the tentative IP address, the **address must be duplicate.**

THANK YOU