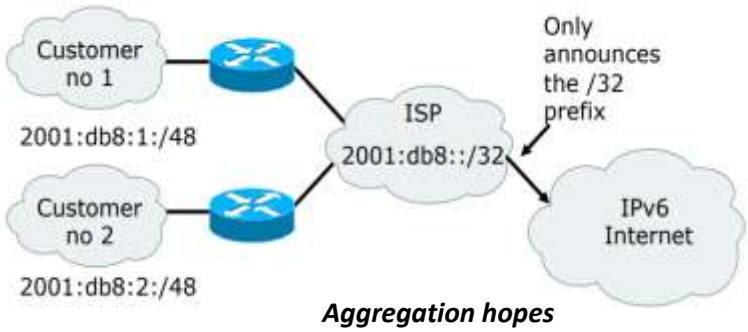

Aggregation hopes

Larger address space enables aggregation of prefixes announced in the global routing table. But current Internet multihoming solution breaks this model

2.1 Internet Protocol Version 6(IPv6)

Internet Protocol Version 6 (IPv6) is a network layer protocol that enables data communications over a packet switched network. Packet switching involves the sending and receiving of data in packets between two nodes in a network.

IPv6 was intended to replace the widely-used Internet Protocol Version 4 (IPv4) that is considered the backbone of the modern Internet. IPv6 is often referred to as the "next generation Internet" because of its expanded capabilities and its growth through recent large scale deployments.

Why is IPv6 necessary?

The most obvious answer is that IPv4 is out of IP addresses. IPv4 has only 4.3 billion addresses, and with PCs, smartphones, tablets, gaming systems, and just about everything else connecting to the Internet we've tapped the system dry. IPv6 uses 128-bit addresses and is capable of 340 undecillion addresses. That is 340 times 10 to the 36th power, or 340 trillion trillion trillion possible IP addresses.

2.1.1 History of IPv6

The current version of the Internet Protocol IPv4 was first developed in the 1970s, and the main protocol standard RFC 791 that governs IPv4 functionality was published in 1981. With the unprecedented expansion of Internet usage in recent years - especially by population dense countries like India and China.

The impending shortage of address space (availability) was recognized by 1992 as a serious limiting factor to the continued usage of the Internet run on IPv4. The following table shows a statistic showing how quickly the address space has been getting consumed over the years after 1981, when IPv4 protocol was published

1985	~ 1/16 of total space
1990	~ 1/8 of total space
1995	~ 1/4of total space
2000	~ 1/2 of total space
Halfway through 2002	~ 2/3 of total space

With admirable foresight, the Internet Engineering Task Force (IETF) initiated as early as in 1994, the design and development of a suite of protocols and standards now known as Internet Protocol Version 6 (IPv6), as a worthy tool to phase out and supplant IPv4 over the coming years. There is an explosion of sorts in the number and range of IP capable devices that are being released in the market and the usage of these by an increasingly tech savvy global population. The new protocol aims to effectively support the ever-expanding Internet usage and functionality, and also address security concerns.

IPv6 uses a128-bit address size compared with the 32-bit system used in IPv4 and will allow for as many as 3.4x1038 possible addresses, enough to cover every inhabitant on planet earth several times over. The 128-bit system also provides for multiple levels of hierarchy and flexibility in hierarchical addressing and routing, a feature that is found wanting on the IPv4-based Internet.

A brief recap of the major events in the development of the new protocol is given below:

- Basic protocol (RFC 2460) published in 1998
- Basic socket API (RFC 2553) and DHCPv6 (RFC 3315) published in 2003.
- Mobile IPv6 (RFC 3775) published in 2004
- Flow label specifications (RFC 3697) added 2004
- Address architecture (RFC 4291) stable, minor revision in 2006
- Node requirements (RFC 4294) published 2006

2.1.4 Feature of IPv6

The massive proliferation of devices, need for newer and more demanding applications on a global level and the increasing role of networks in the way business is conducted are some of the pressing issues the IPv6 protocol seeks to cater to. The following are the features of the IPv6 protocol:

- New header format designed to keep **header overhead to a minimum** - achieved by moving both non-essential fields and optional fields to extension headers that are placed after the IPv6 header. The streamlined IPv6 header is more efficiently processed at intermediate routers.
- **Large address space** - IPv6 has 128-bit (16-byte) source and destination IP addresses. The large address space of IPv6 has been designed to allow for multiple levels of subnetting and address allocation from the Internet backbone to the individual subnets within an organization. Obviates the need for address-conservation techniques such as the deployment of NATs.
- **Efficient and hierarchical addressing and routing infrastructure**- based on the common occurrence of multiple levels of Internet service providers.
- **Stateless and stateful address configuration** both in the absence or presence of a DHCP server. Hosts on a link automatically configure themselves with link-local addresses and communicate without manual configuration.
- **Built-in security**: Compliance with IPSec is mandatory in IPv6, and IPSec is actually a part of the IPv6 protocol. IPv6 provides header extensions that ease the implementation of encryption, authentication, and Virtual Private Networks (VPNs). IPSec functionality is basically identical in IPv6 and IPv4, but one benefit of IPv6 is that IPSec can be utilized along the entire route, from source to destination.
- **Better support for prioritized delivery** thanks to the Flow Label field in the IPv6 header
- **New protocol for neighboring node interaction**- The Neighbor Discovery protocol for IPv6 replaces the broadcast-based Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and ICMPv4 Redirect messages with efficient multicast and unicast Neighbor Discovery messages.
- **Extensibility**- IPv6 can easily be extended for new features by adding extension headers after the IPv6 header.

IPv6 thus holds out the promise of achieving end-to-end security, mobile communications, quality of service (QoS), and simplified system management.

2.1.2 IPv6 Header Format

An Internet Protocol version 6 (IPv6) data packet comprises of two main parts: the header and the payload. The first 40 bytes/octets (40x8 = 320 bits) of an IPv6 packet comprise of the header (see Figure 1) that contains the following fields

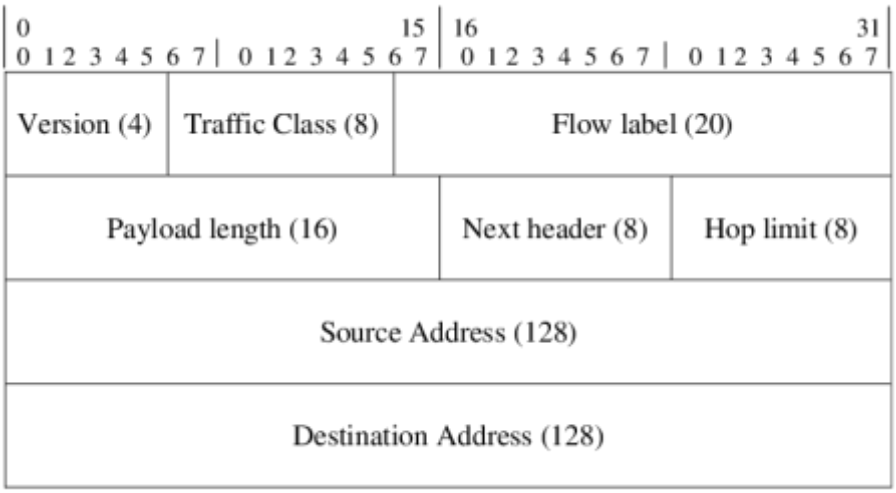

- -Source address (128 bits) The 128-bit source address field contains the IPv6 address of the originating node of the packet. It is the address of the originator of the IPv6 packet.
- -Destination address (128 bits) The 128-bit contains the destination address of the recipient node of the IPv6 packet. It is the address of the intended recipient of the IPv6 packet.
- -Version/IP version (4-bits) The 4-bit version field contains the number 6. It indicates the version of the IPv6 protocol. This field is the same size as the IPv4 version field that contains the number 4. However, this field has a limited use because IPv4 and IPv6 packets are not distinguished based on the value in the version field but by the protocol type present in the layer 2 envelope.
- -Packet priority/Traffic class (8 bits) The 8-bit Priority field in the IPv6 header can assume different values to enable the source node to differentiate between the packets generated by it by associating different delivery priorities to them. This field is subsequently used by the originating node and the routers to identify the data packets that belong to the same traffic class and distinguish between packets with different priorities.
- Flow Label/QoS management (20 bits) The 20-bit flow label field in the IPv6 header can be used by a source to label a set of packets belonging to the same flow. A flow is uniquely identified by the combination of the source address and of a non-zero Flow label. Multiple active flows may exist from a source to a destination as well as traffic that are not associated with any flow (Flow label = 0).

When routers receive the first packet of a new flow, they can process the information carried by the IPv6 header, Routing header, and Hop-by-Hop extension headers, and store the result (e.g. determining the retransmission of specific IPv6 data packets) in a cache memory and use the result to route all other packets belonging to the same flow (having the same source address and the same Flow Label), by using the data stored in the cache memory.

- -Payload length in bytes(16 bits) The 16-bit payload length field contains the length of the data field in octets/bits following the IPv6 packet header. The 16-bit Payload length field puts an upper limit on the maximum packet payload to 64 kilobytes.
- -Next Header (8 bits) The 8-bit Next Header field identifies the type of header immediately following the IPv6 header and located at the beginning of the data field (payload) of the IPv6 packet. This field usually specifies the transport layer protocol used by a packet's payload. The two most common kinds of Next Headers are TCP (6) and UDP (17), but many other headers are also possible. The format adopted for this field is the one proposed for IPv4 by RFC 1700. In case of IPv6 protocol, the Next Header field is similar to the IPv4 Protocol field.
- -Time To Live (TTL)/Hop Limit (8 bits) The 8-bit Hop Limit field is decremented by one, by each node (typically a router) that forwards a packet. If the Hop Limit field is decremented to zero, the packet is discarded. The main function of this field is to identify and to discard packets that are stuck in an indefinite loop due to any routing information errors. The 8-bit field also puts an upper limit on the maximum number of links between two IPv6 nodes. In this way, an IPv6 data packet is allowed a maximum of 255 hops before it is eventually discarded. An IPv6 data packet can pass through a maximum of 254 routers before being discarded. In case of IPv6 protocol, the fields for handling fragmentation do not form a part of the basic header. They are put into a separate extension header. Moreover, fragmentation is exclusively handled by the sending host. Routers are not employed in the Fragmentation process.

Following are the main **comparison between IPv4 header and IPv6 header**.

- IPv6 header is much simpler than IPv4 header.
- The size of IPv6 header is much bigger than that of IPv4 header, because of IPv6 address size.
- In IPv4 header, the source and destination IPv4 addresses are 32 bit binary numbers. In IPv6 header, source and destination IPv6 addresses are 128 bit binary numbers.
- IPv4 header includes space for IPv4 options. In IPv6 header, we have a similar feature known as extension header. IPv4 datagram headers are normally 20-byte in length. IPv6 datagram headers are normally 40-byte in length.
- The fields in the IPv4 header such as IHL (Internet Header Length), identification, flags are not present in IPv6 header.
- Time-to-Live (TTL), a field in IPv4 header, typically used for preventing routing loops, is renamed to it's exact meaning, "Hop Limit".

2.1.5 IPv6 Addressing formats and Types

IPv4 addressing and IPv6 addressing equivalent

IPv4 Address	IPv6 Address
Internet address classes A,B,C,D,E	Not Applicable
Multicast Address(224.0.0.0/4)	FF00::/8
Broadcast Address	Not Applicable
Unspecified Address 0.0.0.0	::
Loopback Address 127.0.0.1	::1
Public IP Address	Global Unicast Address
Private IP Address (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16)	Unique Local(FD00::/8) or Site-local(FEC0::/10)
APIPA Address(169.254.0.0/16)	Link-local address(FE80::/10)
Node ID or Host ID	Interface ID

As defined in RFC 1884 and later revised in RFC 2373, IPv6 addresses are 128-bit identifiers for interfaces and sets of interfaces, not nodes.

IPv6 Addressing Modes

addressing mode refers to the mechanism of hosting an address on the network. IPv6 offers several types of modes by which a single host can be addressed. More than one host can be addressed at once or the host at the closest distance can be addressed.

- ***Unicast***—An identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address. E.g. sending a letter to a friend or phoning them.
- ***Multicast***—An identifier for a set of interfaces that typically belong to different nodes.
- ***Anycast***—An identifier for a set of interfaces that typically belong to nearest nodes provides same service. A packet sent to an anycast address is delivered to the nearest interface *(in terms of routing distance)* in the anycast group
  Example: when you're in Europe, the 8.8.8.8 server will be a close by European server. When you're in Japan, that same IP address(8.8.8.8) will be a close by Asian server. *So, Normally, you want to reach one particular server, anycast wants an answer from "any" server . Anycast can be used for load balancing purposes.*

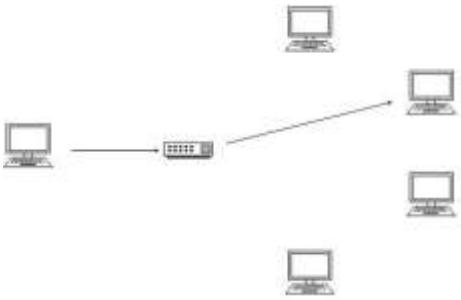

Fig1. Unicast

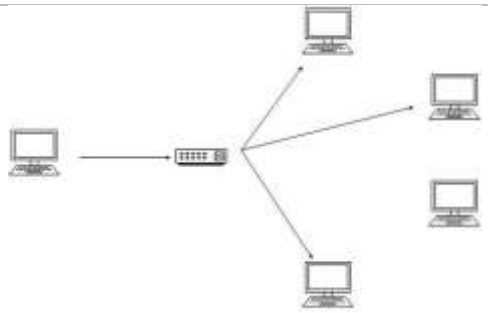

Fig2. Multicast

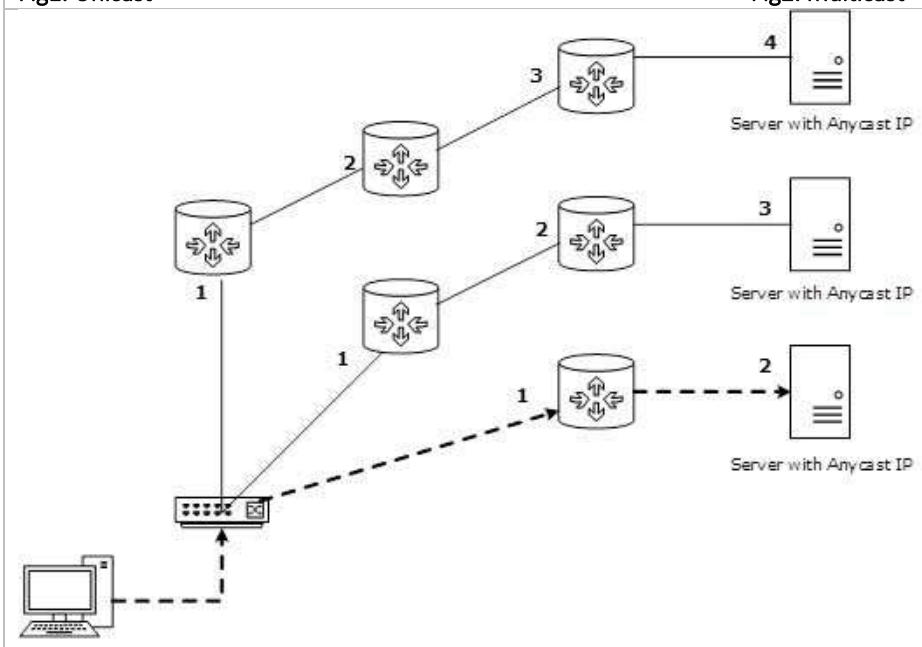

Fig3. Anycast

IPv6 - Address Types & Formats

An IPv6 address can have either of the following two formats:

- Normal - Pure IPv6 format
- Dual - IPv6 plus IPv4 formats

An **IPv6 (Normal)** address has the following format: y : y : y : y : y : y : y : y where *y* is called a *segment* and can be any hexadecimal value between 0 and FFFF. The segments are separated by colons - not periods. An IPv6 normal address must have eight segments, however a short form notation can be used in the Tape Library Specialist Web interface for segments that are zero, or those that have leading zeros. The short form notation can not be used from the operator panel.

he following list shows examples of valid IPv6 (Normal) addresses:

- 2001 : db8: 3333 : 4444 : 5555 : 6666 : 7777 : 8888
- 2001 : db8 : 3333 : 4444 : CCCC : DDDD : EEEE : FFFF
- : : (implies all 8 segments are zero)
- 2001: db8: : (implies that the last six segments are zero)
- : : 1234 : 5678 (implies that the first six segments are zero)
- 2001 : db8: : 1234 : 5678 (implies that the middle four segments are zero)
- 2001:0db8:0001:0000:0000:0ab9:C0A8:0102 (This can be compressed to eliminate leading zeros, as follows: 2001:db8:1::ab9:C0A8:102)

An **IPv6 (Dual)** address combines an IPv6 and an IPv4 address and has the following format: y : y : y : y : y : y : x . x . x . x. The IPv6 portion of the address (indicated with y's) is always at the beginning, followed by the IPv4 portion (indicated with x's).

- In the IPv6 portion of the address, y is called a segment and can be any hexadecimal value between 0 and FFFF. The segments are separated by colons - not periods. The IPv6 portion of the address must have six segments but there is a short form notation for segments that are zero.
- In the IPv4 portion of the address x is called an octet and must be a decimal value between 0 and 255. The octets are separated by periods. The IPv4 portion of the address must contain three periods and four octets.

The following list shows examples of valid IPv6 (Dual) addresses:

- 2001 : db8: 3333 : 4444 : 5555 : 6666 : 1 . 2 . 3 . 4
- : : 11 . 22 . 33 . 44 (implies all six IPv6 segments are zero)
- 2001 : db8: : 123 . 123 . 123 . 123 (implies that the last four IPv6 segments are zero)
- : : 1234 : 5678 : 91 . 123 . 4 . 56 (implies that the first four IPv6 segments are zero)
- : : 1234 : 5678 : 1 . 2 . 3 . 4 (implies that the first four IPv6 segments are zero)
- 2001 : db8: : 1234 : 5678 : 5 . 6 . 7 . 8 (implies that the middle two IPv6 segments are zero)

Address Structure

An IPv6 address is made of 128 bits divided into eight 16-bits blocks. Each block is then converted into 4-digit Hexadecimal numbers separated by colon symbols.

For example, given below is a 128 bit IPv6 address represented in binary format and divided into eight 16-bits blocks:

```
0010000000000001 0000000000000000 0011001000111000 1101111111100001 0000000001100011 0000000000000000
0000000000000000 1111111011111011
```

Each block is then converted into Hexadecimal and separated by ':' symbol:

```
2001:0000:3238:DFE1:0063:0000:0000:FEFB
```

Even after converting into Hexadecimal format, IPv6 address remains long. IPv6 provides some rules to shorten the address. The rules are as follows:

Rule.1: Discard leading Zero(es): In Block 5, 0063, the leading two 0s can be omitted, such as (5th block):

```
2001:0000:3238:DFE1:63:0000:0000:FEFB
```

Rule.2: If two of more blocks contain consecutive zeroes, omit them all and replace with double colon sign ::, such as (6th and 7th block):

```
2001:0000:3238:DFE1:63::FEFB
```

Consecutive blocks of zeroes can be replaced only once by :: so if there are still blocks of zeroes in the address, they can be shrunk down to a single zero, such as (2nd block):

```
2001:0:3238:DFE1:63::FEFB
```

Interface ID – host id in IPv4

IPv6 has three different types of Unicast Address scheme. The second half of the address (last 64 bits) is always used for Interface ID. The MAC address of a system is composed of 48-bits and represented in Hexadecimal. MAC addresses are considered to be uniquely assigned worldwide. Interface ID takes advantage of this uniqueness of MAC addresses. A host can auto-configure its Interface ID by using IEEE's **Extended Unique Identifier (EUI-64) format**. First, a host divides its own MAC address into two 24-bits halves. Then 16-bit Hex value 0xFFFE is sandwiched into those two halves of MAC address, resulting in EUI-64 Interface ID.

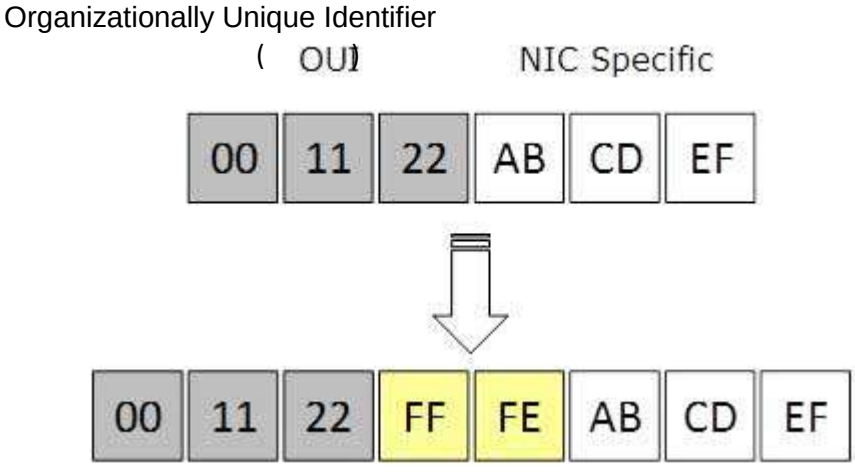

[*Fig: EUI-64 Interface ID*]

Conversion of Extended Unique Identifier(EUI)-64 ID into IPv6 Interface Identifier

To convert EUI-64 ID into IPv6 Interface Identifier, the most significant 7th bit of EUI-64 ID is complemented. For example:

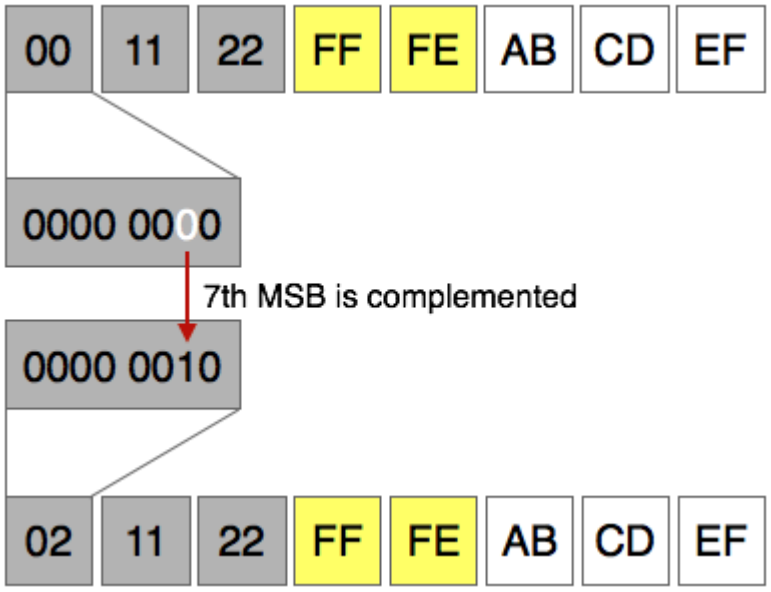

[*Fig: IPV6 Interface ID*]

Now, Link Local IPv6 Address will be generated as

FE80::0211:22FF: FEAB: CDEF

IPv6 Addressing Types

- Unicast addresses. *A* packet is delivered to one interface. In this case there is just one sender, and one receiver.(**one-to-one**) e.g. 3731:54:65fe:2::a7 , 0:0:0:0:0:0:0:1 is called the loopback address. It may be used by a node to send an IPv6 packet to itself. Unicast address are:-
  i. **Interface Identifiers**: used to identify interfaces on a link. They are required to be unique within a subnet prefix. It is recommended that the same interface identifier not be assigned to different nodes on a link. They may also be unique over a broader scope.
  ii. **Unspecified Address**- (**::**) The address 0:0:0:0:0:0:0:0 is called the unspecified address. It must never be assigned to any node. It indicates the absence of an address.
  iii. **Loopback Address**:- (**::0001**) The unicast address 0:0:0:0:0:0:0:1 is called the loopback address. It may be used by a node to send an IPv6 packet to itself. It must not be assigned to any physical interface. It is treated as having Link-Local scope and may be thought of as the Link-Local unicast address of a virtual interface (typically called the "loopback interface") to an imaginary link that goes nowhere.
  iv. **Global unicast addresses**, which are conventional, publicly routable address, just like conventional IPv4 publicly routable addresses.

3 Bits	45 Bits	16 Bits	64 Bits
001	Global Routing Prefix	Subnet ID	Interface ID

- **GRP**(Global Routing Prefix) is used to identify a address type like multicast or an address range assigned to a site.
- **Subnet ID** is used to identify subnets within a site, used within a organization's site.
- **Interface ID** is used to identify an interface on a specific subnet within the site. Its size is 64 bits. It is known Node ID or Host ID in IPv4.

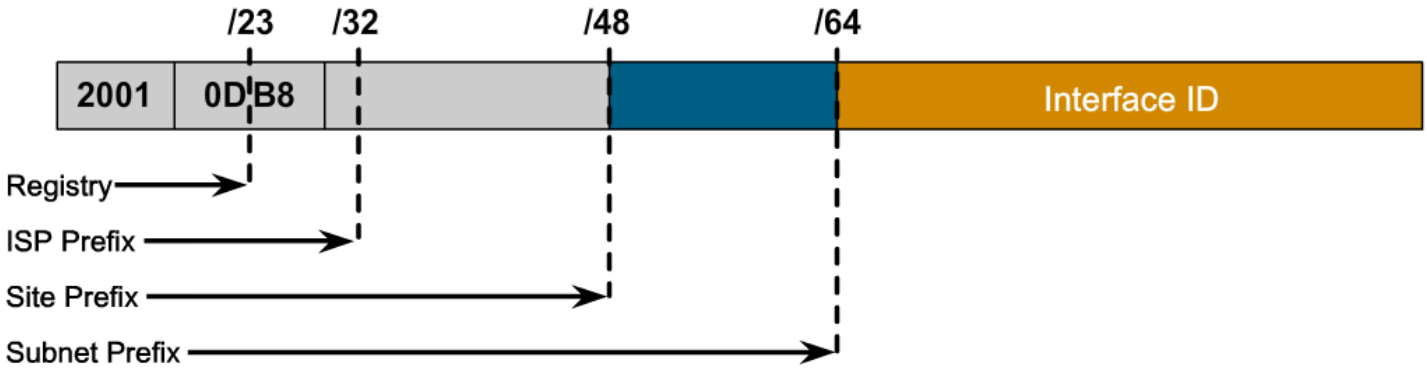

v. **Link-local addresses (FE80::/10 e.g. fe80::200:5aee:feaa:20a2)** are akin to the private, non-routable addresses . They are not meant to be routed, but confined to a single network segment. Link-local addresses mean you can easily throw together a temporary LAN, such as for conferences or meetings, or set up a permanent small LAN the easy way.
Link-local IP addresses are non-routable. By this I mean that no such packet should be received from a router on one IP interface and forwarded out another IP interface, even if it is within an organization.

10 Bits	54 Bits	64 Bits
1111 1110 10	000 … 000	Interface ID

vi. **Site-Local Addresses**: Site-Local addresses were originally designed to be used or addressing inside of a site without the need for a global prefix. **e.g FEC0::/10**
Site-local IP addresses are routable within an organization, but they are not routable in internetworks (such as the Internet). Here a "site" is essentially an organization. IPv4 defines site-local in RFC 1918 and reserves three ranges for this purpose (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16).

10 Bits	54 Bits	64 Bits
1111 1110 11	000 … 000	Interface ID

vii. Transition Address

- **IPv4-mapped address** (**::ffff/96 e.g. ::ffff:192.0.2.47**) Two types of IPv6 addresses are defined that carry an IPv4 address in the low-order 32 bits of the address. Is used to represent an IPv4 as IPv6 address.
- **IPv4 compatible address**- **::192.172.12.3 –** communicating with IPv6 over an IPv4 infrastructure that uses public IPv4 address.
- **6to4 address**- 2002:WWXX:YYZZ:SubnetID:InterfaceID is assigned a node for the 6to4 transition technology.
- **Teredo** address: 2001::/32 is assigned to a node for the Teredo IPv6 transition technology

- **Multicast addresses.** *A* packet is delivered from one or more points to multiple or a set of interfaces.(**one-to-many or many-to-many**). These addresses are used to identify multicast groups. They should only be used as destination addresses, never as source addresses. e.g. addresses fall under the range FF00::/8 , FF01:0:0:0:0:0:0:1

The main goal of multicasting is having an efficient network to save bandwidth on links by optimizing the number of packets exchanged between nodes. Multicast implies the concept of a group:

- Any node can be a member of a multicast group
- A source node may send packets to a multicast group
- All members of a multicast group get packets that are sent to the group

Note : IPv6 does not use broadcast messages.

Multicast Address e.g. ff00::/8

1111 1111	Flags	Scope	Group Identifier
8 bits	4 bits	4 bits	112 bits

For IPv6 multicast addresses, the first eight bits are reserved as 1111 1111. Thus, the prefix of an IPv6 multicast address is ff00::/8. Similar to IPv6 Link Local addresses, it is easy to identify an IPv6 multicast address, because IPv6 multicast addresses have left most hexadecimal digits as "FF"

• After the leftmost 8 bits which are reserved as "1111 1111", the next four bits are known as flags. Only 3 of the 4 flag bits in the flags field are defined currently. The most significant bit in the 4 bits flags field is reserved for future use. The remaining three flags are known as R, P and T.

4 Bits			
0	1	1	0
0	1	2	3

4 Bits inside flags field	Flag name	When "0" set	When "1" set
0 (Most Significant Bit)	Currently not in use	Currently not in use	Currently not in use
1	R (Rendezvous)	When R flag set to 0, the multicast rendezvous point is not embedded with multicast address	When R flag set to 1, the multicast rendezvous point is embedded with multicast address

2	P (Prefix)	When P flag set to 0, the multicast address is not based on network prefix	When P flag set to 1, the multicast address based on network prefix
3 (Least Significant Bit)	T (Transient)	When T flag set to 0, the multicast address is a permanently assigned (well-known) multicast IPv6 address	When T flag set to 1, the multicast address is a transient (Dynamically assigned) multicast address

• After the leftmost 8 bits which are reserved as "1111 1111", and the next four flag bits, the next four bits are defined as the Scope bits. Scope bits (4 bits) are used to indicate the scope of delivery of IPv6 multicast traffic.
The following table lists the values possible currently for the scope field. E.g. **FF02:: – link-local scope**.

Hex Value	Scope	Meaning
0	Reserved	Currently not in use
1	Interface-local scope	The Interface-local scope is limited for a local single interface only. Useful only for loopback delivery of multicasts within a node.
2	Link-local scope	Link-local scope is defined for the local link. The traffic with the multicast address of FF02::2 is limited to local link scope. An IPv6 router will never forward the multicast traffic destined to FF02::2 beyond the local link.
3	Subnet-local scope	Subnet-local scope ranges subnets on multiple links.

Group ID: identifies the multicast group and is unique within the scope. Its size is 112 bits.

- **Anycast addresses.** *A* packet is delivered to the nearest of multiple interfaces (the "nearest" one, according to the routing protocols' measure of distance*).*(**one-to-any**) e.g. FF01:0:0:0:0:0:0:1 (IPv4 - 224.0.0.0/4)

Anycast address are Link-Local (FE80::/10), Site-Local(FEC0::/10), Aggrigatable Global(2001::/16, 2002::/16,3FFE::/16) - Anycast addresses use aggregatable global unicast addresses. They can also use site-local or link-local addresses. Note that it is impossible to distinguish an anycast address from a unicast address.

An IPv6 anycast address is an identifier for a set of interfaces (typically belonging to different nodes) - A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the “nearest” one, according to the routing protocol’s measure of distance).

2.1.3 Problem with IPv4

Initial design of IPv4 did not anticipate the growth of internet and this created many issues, which proved IPv4 need to be changed. The main limitations of IPv4 are listed below.

- **Scarcity of IPv4 Addresses:** The IPv4 addressing system uses 32-bit address space. This 32-bit address space is further classified to usable A, B, and C classes. 32-bit address space allows for 4,294,967,296 IPv4 addresses, but the previous and current IPv4 address allocation practices limit the number of available public IPv4 addresses. Today the Earth’s population stands at around 6.6 billion while the Internet has a population of just 1.3 billion, which is not even 22% of the entire world’s population. Many addresses which are allocated to many companies were not used and this created scarcity of IPv4 addresses.
  Because scarcity of IPv4 addresses, many organizations implemented *NAT (Network Address Translation) to map multiple private IPv4 addresses to a single public IPv4 address.* But NAT (Network Address Translation) do not support network layer security standards and it do not support the mapping of all upper layer protocols. NAT can also create network problems when two organizations which use same private IPv4 address ranges communicate. More servers, workstations and devices which are connected to the internet also demand the need for more addresses and the current statistics prove that public IPv4 address space will be depleted soon. The scarcity of IPv4 address is a major limitation of IPv4 addressing system.
- **Security Related Issues:** As we discussed before, RFC 791 (IPv4) was published in 1981 and the current network security threats were not predicted that time. Internet Protocol Security (IPSec) is a protocol suit which enables network security by protecting the data being sent from being viewed or modified. Internet Protocol Security (IPSec) provides security for IPv4 packets, but Internet Protocol Security (IPSec) is not built-in and optional.
- **Address configuration related issues:** Networks and also internet is expanding and many new computers and devices are using IP. The configuration of IP addresses (static or dynamic) should be simple. But, in IPv4 stateful configuration is happened.
- **Quality of service (QoS):** Quality of Service (QoS) is available in IPv4 and it relies on the 8 bits of the IPv4 Type of Service (TOS) field and the identification of the payload. IPv4 Type of Service (TOS) field has limited functionality and payload identification (uses a TCP or UDP port) is not possible when the IPv4 datagram packet payload is encrypted.

In Summary…

- **Deficiency of address space** - various devices connected to the Internet grows exponentially. The size of address space 2 32 is quickly exhausted.
- **Loss of transparency** - due to the use of mechanisms such as NAT (Network Address Translator).
- **Loss of robustness** - because of the implemented topology that has little room for redundancy.

- **Loss of stable addresses** - i.e. the address of a node changes each time it is connected to the Internet.
- **Weak expansibility of the protocol** - the insufficient size of heading IPv4 doesn't allow placing demanded quantity of additional parameters in it.
- **Problem of safety of communications** - it is not stipulated any means for differentiation of access to the information placed in a network.
- **Absence of support of quality of service (QoS)** - accommodation of the information about throughput, the delays and demanded for normal work of some network appendices is not supported.
- **Absence of the auto-configuration** - IP addresses mechanism. Machine renumbering problem.
- **Loss of application independence** - An example is that many systems are developed with functionality to avoid problems created by NAT.

DIFFERENCES BETWEEN IPV4 AND IPV6

IPv4	IPv6
IPv4 is less address space	IPv6 is increased address space
IPv4 uses a 32 bits address size	IPv6 uses a 128 bits address size
Must support DHCP or be configured manually	Does not require DHCP or manual configuration, it supports stateless auto configuration
IPSec is not compulsory	IPSec is compulsory
Broadcasts sends traffic to all hosts on a subnet	There are no broadcasts instead multicasts is used thereby reducing broadcast floods found in IPv4
The IP header has a variable length of 20-60 bytes depending on the IP header options	IP header has a fixed length of 40 bytes and there are no IP header options available

2.2 ICMPv6

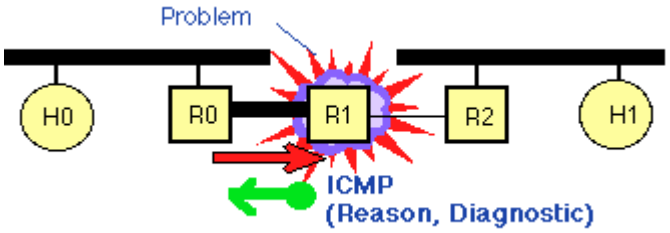

The Internet Control Message Protocol (**ICMP**) is a supporting protocol in the Internet protocol suite. It is used by network devices, including routers, to send error messages and operational information indicating, for example, that a requested service is not available or that a host or router could not be reached.

The **Internet Control Message Protocol Version 6** (ICMPv6) is a new version of the ICM protocol that forms an integral part of the Internet Protocol version 6 (IPv6) architecture and performs error reporting and diagnostic functions (e.g., ping), and has a framework for extensions to implement future changes. ICMPv6 messages are transported within an IPv6 packet that may include IPv6 extension headers.

ICMPv6 offers a comprehensive solution by offering the different functions earlier subdivided among the different protocols such as **ICMP**, **ARP** (Address Resolution Protocol), and **IGMP** (Internet Group Membership Protocol version 3).

- ICMPv6 is used by IPv6 nodes to report errors encountered in processing packets, and to perform other internet-layer functions, such as diagnostics (ICMPv6 "ping").
- ICMPv6 further simplifies the communication process by eliminating obsolete messages.
- ICMPv6 is a multipurpose protocol and is used for a variety of activities including error reporting in packet processing, diagnostic activities, Neighbor Discovery process and IPv6 multicast membership reporting.

To perform these activities, ICMPv6 messages are subdivided into two classes: error messages and information messages.

2.2.1 Features

1. Error Messages RFC 4443- Error messages report errors in the forwarding or delivery of IPv6 packets by either the destination node or an intermediate router. The high-order bit of the 8-bit Type field for all ICMPv6 error messages is set to **0**. Therefore, valid values for the Type field for ICMPv6 error messages are in the range **of 0 - 127**. The ICMPv6 error messages belong to four different categories:

- Destination Unreachable,: A Destination Unreachable message SHOULD be generated by a router, or by the IPv6 layer in the originating node, in response to a packet that cannot be delivered to its destination address for reasons other than congestion.
- Time Exceeded: If a router receives a packet with a Hop Limit of zero, or if a router decrements a packet's Hop Limit to zero, it MUST discard the packet and originate an ICMPv6 Time Exceeded message with Code 0 to the source of the packet. This indicates either a routing loop or too small an initial Hop Limit value.
- Packet Too Big: A Packet Too Big MUST be sent by a router in response to a packet that it cannot forward because the packet is larger than the MTU of the outgoing link. The information in this message is used as part of the Path MTU Discovery process [PMTU].

- Parameter Problems: If an IPv6 node processing a packet finds a problem with a field in the IPv6 header or extension headers such that it cannot complete processing the packet, it MUST discard the packet and SHOULD originate an ICMPv6 Parameter Problem message to the packet's source, indicating the type and location of the problem.

2. Information Messages RFC 4443 - Informational messages provide diagnostic functions and additional host functionality, such as MLD and ND. The high-order bit of the 8-bit Type field for all ICMPv6 informational messages is set to **1**. Therefore, valid values for the Type field for ICMPv6 information messages are in the range of **128 - 255**. The ICMPv6 information messages are subdivided into three groups: diagnostic messages, Neighbor Discovery messages, and messages for the management of multicast groups.

i. **Echo Request Message:** Every node MUST implement an ICMPv6 Echo responder function that receives Echo Requests and originates corresponding Echo Replies. A node SHOULD also implement an application-layer interface for originating Echo Requests and receiving Echo Replies, for diagnostic purposes.

ii. **Echo Reply Message**: Upper Layer Notification - Echo Reply messages MUST be passed to the process that originated an Echo Request message. An Echo Reply message MAY be passed to processes that did not originate the Echo Request message.

2.2.2 General Message Formats

ICMPv6 packets have the format shown in the figure. The 8-bit Type field indicates the type of the message. If the high-order bit has value zero (values in the range from 0 to 127), it indicates an error message; if the high-order bit has value 1 (values in the range from 128 to 255),it indicates an information message. The 8-bitCodefield content depends on the message type. The Checksum field helps in the detection of errors in the ICMP message and in part of the IPv6 message.

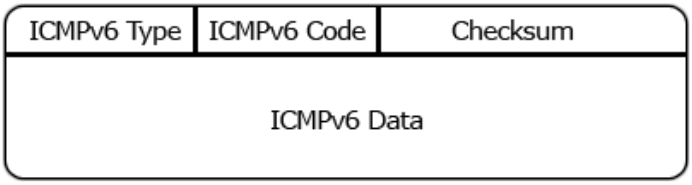

ICMPv6 Message Types

ICMPv6 is a multipurpose protocol as it is used for a flood of activities such as reporting errors encountered in processing data packets, reporting multicast memberships, performing Neighbor Discovery, and performing diagnostics. An ICMP message is identified by a value of 58 in the Next Header field of the IPv6 header or of the preceding Header.

ICMPv6 error messages report forwarding or delivery errors by either a router or the destination host and they consist of the following messages: -

- ICMPv6 Type 1 (Destination Unreachable)
- ICMPv6 Type 2 (Packet Too Big)
- ICMPv6 Type 3 (Time Exceeded)
- ICMPv6 Type 4 (Parameter Problem)

ICMPv6 Advantages

- If a wrong IP address is used for configuring a client to the DNS server, an ICMP message is sent by the destination device to indicate the error.
- If a program does not allow fragmentation of its communications but it is required to communicate with a destination device, the router undertaking the fragmentation of the packet sends an ICMP message to the source device to indicate the error.
- If a client sends all communications to a particular router despite another router offering a best route, the particular router responds with the IP address of the router that provides a better route in the form of an ICMP message.
- All IP headers contain a Time to Live (TTL) value. This value is decremented as the IP packet is forwarded through each router. If a packet arrives at a router with a Time To Live (TTL) value of 1, the router cannot decrement the value any further and forward it. Instead, the router discards the packet and sends an ICMP message to indicate the expiry of the packet's TTL value.
- The Internet Control Message Protocol Version 6 (ICMPv6) also provides testing and diagnostics services for many utilities. In order to test the communication process, an ICMP echo is used by the Internet Protocol Packet Internet Gopher (PING) utility.

2.2.3 ICMP Error and International Message Types

Every ICMP message can have a slightly different header depending on the kind of error report or information it carries. The following sections outline the structure of each type of ICMPv6 message. If the destination is unreachable due to congestion, no ICMP message is generated. A host that receives a Destination Unreachable message must inform the upper-layer process.

ICMP Error Message Types

1. Destination Unreachable

A Destination Unreachable message is generated if an IP datagram cannot be delivered. A Type field with the value 1 identifies this message. The ICMP message is sent to the source address of the invoking packet. The format of the Destination Unreachable message is shown in Figure 2.2.3.1

The **Type** field is set to one, which is the value for the Destination Unreachable message. The **Code** field supplies more information about the reason why the datagram was not delivered. The possible codes are listed in Table 2.2.3.1. The data portion of the ICMP message contains parts of the original message—as much as will fit into the ICMP message.

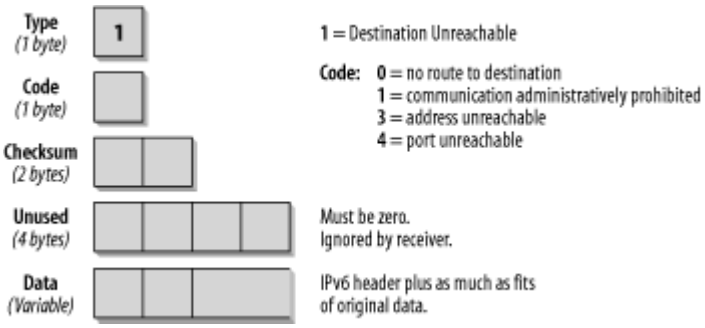

Figure 2.2.3.1. Format of the Destination Unreachable message

Table 2.2.3.1. Code values of the Destination Unreachable message (type 1)

Code	Description
0	**No route to destination** This message is generated if a router cannot forward a packet because it does not have a route in its table for a destination network. This can only happen if the router does not have an entry for a default route.
1	**Communication with destination administratively prohibited** The communication is prohibited by administrative policy. For example, be sent by a firewall that cannot forward a packet to a host inside the firewall because of a packet filter. It might also be sent if a node is configured not to accept unauthenticated Echo Requests.
2	**Beyond/ Outside scope of source address** The destination is beyond the scope of the source address. A router sends this when the packet must be forwarded using an interface that is not within the scoped zone of the source address.
3	**Address unreachable** This code is used if a destination address cannot be resolved into a corresponding network address or if there is a data-link layer problem preventing the node from reaching the destination network.
4	**Port unreachable** This code is used if the transport protocol (e.g., UDP) has no listener and if there is no other means to inform the sender. For example, if a Domain Name System (DNS) query is sent to a host and the DNS server is not running, this type of message is generated.
5	**Source Address Failed Ingress/ Egress Policy** The packet with this source is not allowed because of inbound(ingress) or outbound(egress) packet filtering.
6	**Reject Route to Destination** The packet matched a reject route and was discarded. A reject route is an address prefix configured on a router for traffic that the router must immediately discard.

2. Packet Too Big

If a router cannot forward a packet because it is larger than the MTU of the outgoing link, it will generate a Packet Too Big message (shown in Figure 2.2.3.2). This ICMPv6 message type is used as part of the Path MTU discovery process discussed later in this chapter. The ICMP message is sent to the source address of the invoking packet.

The **Type** field has the value 2, which identifies the Packet Too Big message. In this case, the **Code** field is not used and is set to zero. The important information for this type of message is the MTU field, which contains the MTU size of the next hop link.

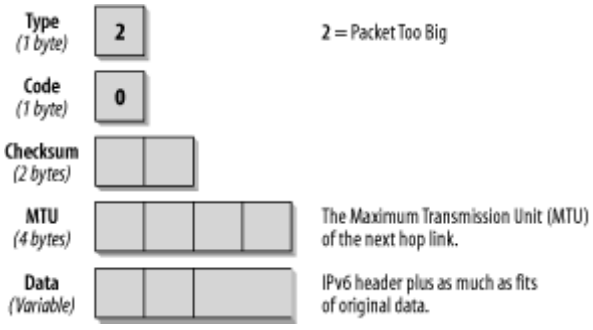

Figure 2.2.3.2. Format of the Packet Too Big message

3. Time Exceeded

When a router forwards a packet, it always decrements the hop limit by one. Remember, the hop limit makes sure that a packet does not endlessly travel through a network. If a router receives a packet with a hop limit of one and decrements the limit to zero, it discards the packet, generates a Time Exceeded message with a code value of zero, and sends this message back to the source host. This error

can indicate a routing loop or the fact that the sender's initial hop limit is too low. It can also tell you that someone used the *traceroute* utility, which is described later in the chapter. Figure 4-4 shows the format of the Time Exceeded message.

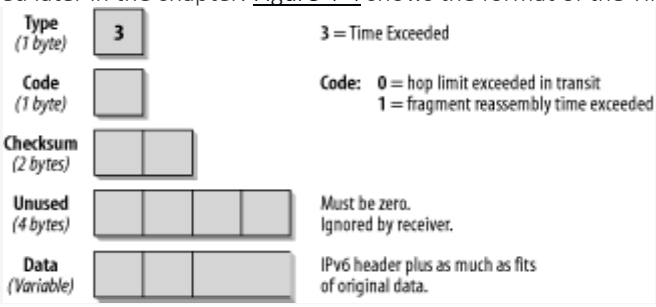

Figure 4-4. Format of the Time Exceeded message

The **Type** field carries the value 3, specifying the Time Exceeded message. The **Code** field can be set to 0, which means hop limit exceeded in transit, or to 1, which means that the fragment reassembly time is exceeded. The **data** portion of the ICMP message contains parts of the original message—as much as fits into the ICMP message, depending on the MTU used.

An incoming Time Exceeded message must be passed to the upper-layer process. Table 4-4 shows the Code fields for the Time Exceeded message.

Table 4-4. Code values for Time Exceeded message (type 3)

Code	Description
0	**Hop limit exceeded in transit.** Possible causes: the initial hop limit value is too low, or there are routing loops.
1	**Fragment reassembly time exceeded.** If a fragmented packet is sent by using a fragment header and the receiving host cannot reassemble all packets within a certain time, it notifies the sender by issuing this ICMP message.

4. Parameter Problem

If an IPv6 node cannot complete the processing of a packet because it has a problem identifying a field in the IPv6 header or in an Extension header, it must discard the packet and it should send an ICMP Parameter Problem message back to the source of the problem packet. This type of message is often used when an error that does not fit into any of the other categories is encountered. The format of this ICMP message is shown in Figure 4-5.

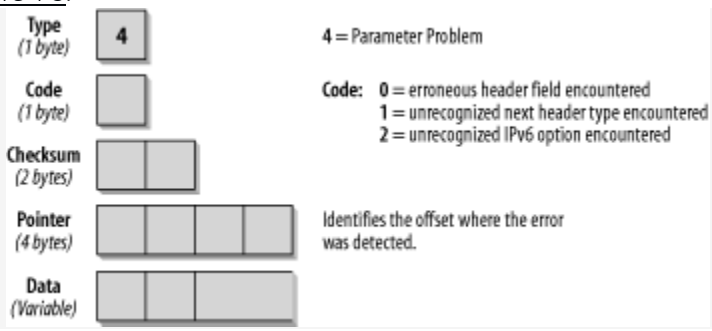

Figure 4-5. Format of the Parameter Problem message

The **Type** field has the value 4, which specifies the Parameter Problem message. The **Code** field can contain any of the three values described in Table 4-5. The **Pointer** field identifies at which byte in the original packet the error was detected. The ICMP message includes as much of the original **data** as fits up to the minimum IPv6 MTU. It is possible that the pointer points beyond the ICMPv6 message. This would be the case if the field in error was beyond what can fit in the maximum size of an ICMPv6 error message. Table 4-5 shows the Code fields for the Parameter Problem message.

Table 4-5. Code values for Parameter Problem (type 4)

Code	Description
0	**Erroneous header field encountered** An error in a field within the IPv6 header or an extension header was encountered
1	**Unrecognized next header type encountered** An enrecognized Next Header field value was encountered.
2	**Unrecognized IPv6 option encountered** An unrecognized IPv6 option was encountered.

For example, if you see an ICMPv6 message of type 4 with a code value of 1 and a pointer set to 40, this indicates that the next header type in the header following the IPv6 header was unrecognized

ICMP International Message Types/ Formats

1. Echo Request

An IPv6 node sends an ICMPv6 Echo Message to a destination to solicit an immediate Echo Reply message. The Echo Request/ Echo Reply message facility provides a simple diagnostic function to aid in the troubleshooting of a variety of reachability and routing problems.
In the Echo Request message, the Type field is set to 128 and the Code field is set to 0. Following Checksum field are the 16 bit(2byes) Identifier and 2 Bytes Sequence Number fields. The Identifier and Sequence Number fields are set by the sending host so that they can be used to match an incoming Echo Reply message with a sent Echo Request message. The Data field is zero or more bytes of optional data that is also set by sending host.

2. Echo Reply

An IPv6 node sends an ICMPv6 Echo Reply message in response to the receipt of an ICMPv6Echo Request message.
In the Echo Reply message, the Type field is set to 129 and the Code field is set to 0. Following Checksum field are 16 bit Identifier and 16 bit Sequence Number fields. The Identifier, Sequence Number and Data fields are set with the same values as those tin the Echo Request message that prompted the Echo Reply.

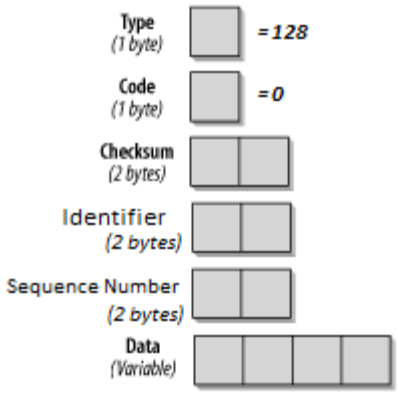

Fig. Structure of Echo Request Message

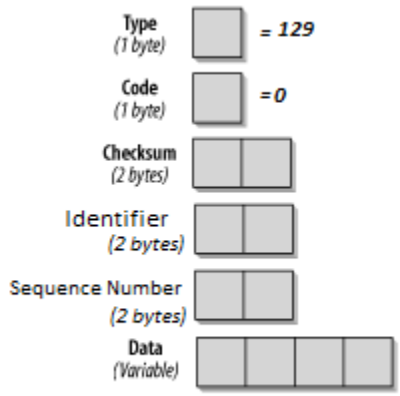

Fig. Structure of Echo Reply Message

Comparison between ICMPv4 and ICMPv6 Messages

ICMPv4	ICMPv6
Destination Unreachable-Network Unreachable (Type 3, Code 0)	**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)	**Destination Unreachable**-Unrecognized Next Header Type Encountered (Type 4, Code 1)
Destination Unreachable-Port Unreachable (Type 3, Code 3)	**Destination Unreachable**-Port Unreachable (Type 1, Code 4)
Destination Unreachable-Fragmentation Needed (Type 3, Code 4)	Packet Too Big (Type 2, Code 0)
Destination Unreachable-Communication with Destination Hos Administratively Prohibited (Type 3, Code 10)	**Destination Unreachable**-Communication with Destination Hos Administratively Prohibited (Type 1, Code 1)
Source Quench (Type 4, Code 0)	Not present
Redirect (Type 5, Code 0)	Neighbor Discover Redirect message(Type 137,Code 0)
Time Exceeded-TTL exceeded in Transit (Type 11, Code 0)	**Time Exceeded**-Hop Limit exceeded in Transit (Type 3, Code 0)
Time Exceeded-Fragment Reassembly Time Exceeded (Type 11, Code 1)	**Time Exceeded**-Fragment Reassembly Time Exceeded (Type 3, Code 1)
Parameter Problem (Type 12, Code 0)	Parameter Problem (Type 4, Code 0 or 2)

2.2.4 Neighbor Discovery (ND) RFC4861: *that determine relationships between neighboring nodes*

IPv6 ND is a set of message and process that determine relationships between neighboring nodes. The **Neighbor Discovery Protocol** (**NDP**, **ND**) is a protocol in the Internet protocol suite used with Internet Protocol Version 6 (IPv6). It operates in the Link Layer of the Internet model, and is responsible for address autoconfiguration of nodes, discovery of other nodes on the link, determining the addresses of other nodes, duplicate address detection, finding available routers and Domain Name System (DNS) servers, address prefix discovery, and maintaining reachability information of other active neighbor nodes.
The protocol defines five different ICMPv6 packet types to perform functions for IPv6 similar to the Address Resolution Protocol (ARP) and Internet Control Message Protocol (ICMP) Router Discovery and Router Redirect protocols for IPv4, which is replaced by ND.
IPv6 nodes use Neighbor Discovery for the following purposes:

- For Stateless Autoconfiguration of IPv6 addresses
- To determine network prefixes, routes, and other configuration information
- For Duplicate IP Address Detection (**DAD**)
- To determine Layer 2 addresses of nodes on the same link

- To find neighboring routers that can forward their packets
- To keep track of which neighbors are reachable and which are not (NUD)
- To detect changed link-layer addresses

ND is used by nodes **(e.g. switch, workstation, pc as clients)** to do the following:

- Resolve the link-layer address of a neighboring node to which an IPv6 packet is being forwarded.
- Determine when the link-layer address of a neighboring node has changed.
- Determine whether a neighbor is still reachable.

ND is used by hosts **(e.g. server, printer)** to do the following:

- Discover neighboring routers.
- Auto configure addresses, address prefixes, routes, and other configuration parameters.

ND is used by routers to do the following:

- Advertise their presence, host configuration parameters, routes, and on-link prefixes.
- Inform hosts of a better next-hop address to forward packets for a specific destination.

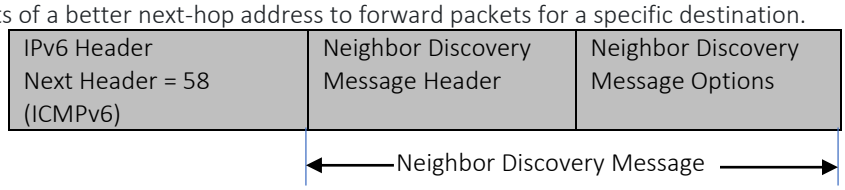

Fig. ND Message Format

The Neighbor Discovery protocol provides a multitude of **improvements over the IPv4** set of protocols:

- **Router Discovery** is part of the base protocol set; there is no need for hosts to "snoop" the routing protocols.
- **Router Advertisements** carry link-layer addresses; no additional packet exchange is needed to resolve the router's link-layer address.
- **Router Advertisements** carry prefixes for a link; there is no need to have a separate mechanism to configure the "netmask".
- **Router Advertisements** enable Address Autoconfiguration.
- **Routers can advertise an MTU** for hosts to use on the link, ensuring that all nodes use the same MTU value on links lacking a well-defined MTU.
- **Address resolution** multicasts are "spread" over 16 million (2^24) multicast addresses, greatly reducing address-resolution-related interrupts on nodes other than the target. Moreover, non-IPv6 machines should not be interrupted at all.
- **Redirects** contain the link-layer address of the new first hop; separate address resolution is not needed upon receiving a redirect.
- **Neighbor Unreachability Detection** is part of the base, which significantly improves the robustness of packet delivery in the presence of failing routers, partially failing or partitioned links, or nodes that change their link-layer addresses. For instance, mobile nodes can move off-link without losing any connectivity due to stale ARP caches.
- Unlike ARP, **Neighbor Discovery detects** half-link failures (using Neighbor Unreachability Detection) and avoids sending traffic to neighbors with which two-way connectivity is absent.
- Unlike in IPv4 Router Discovery, the **Router Advertisement** messages do not contain a preference field. The preference field is not needed to handle routers of different "stability"; the Neighbor Unreachability Detection will detect dead routers and switch to a working one.
- The use of link-local addresses to uniquely identify routers (for **Router Advertisement and Redirect messages**) makes it possible for hosts to maintain the router associations in the event of the site renumbering to use new global prefixes.
- By setting the Hop Limit to 255, **Neighbor Discovery** is immune to off-link senders that accidentally or intentionally send ND messages. In IPv4, off-link senders can send both ICMP Redirects and Router Advertisement messages.
- Placing address resolution at the **ICMP layer** makes the protocol more media-independent than ARP and makes it possible to use generic IP-layer authentication and security mechanisms as appropriate.

There are five different ND messages:

ND message options provide additional information indicating MAC address, on-link network prefixes, on-link MTU information, redirection data , mobility information and specific routes. IPv6 NDP uses 5 ICMPv6 messages for the neighbor discovery mechanism

1. Router Solicitation (RS - 133) – Router – to - Device Messaging- ***Sent when device needs IPv6 addressing information***

The Router Solicitation message is sent by IPv6 hosts to discover the presence of IPv6 routers on the link. A host sends a multicast Router Solicitation message to prompt IPv6 routers to respond immediately, rather than waiting for an unsolicited Router Advertisement message.

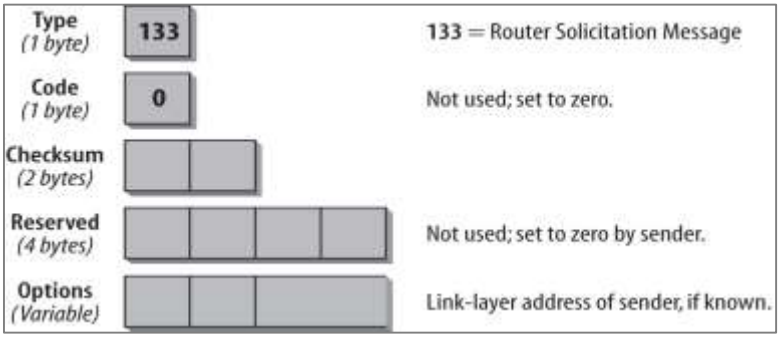

Working Principle

- 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.

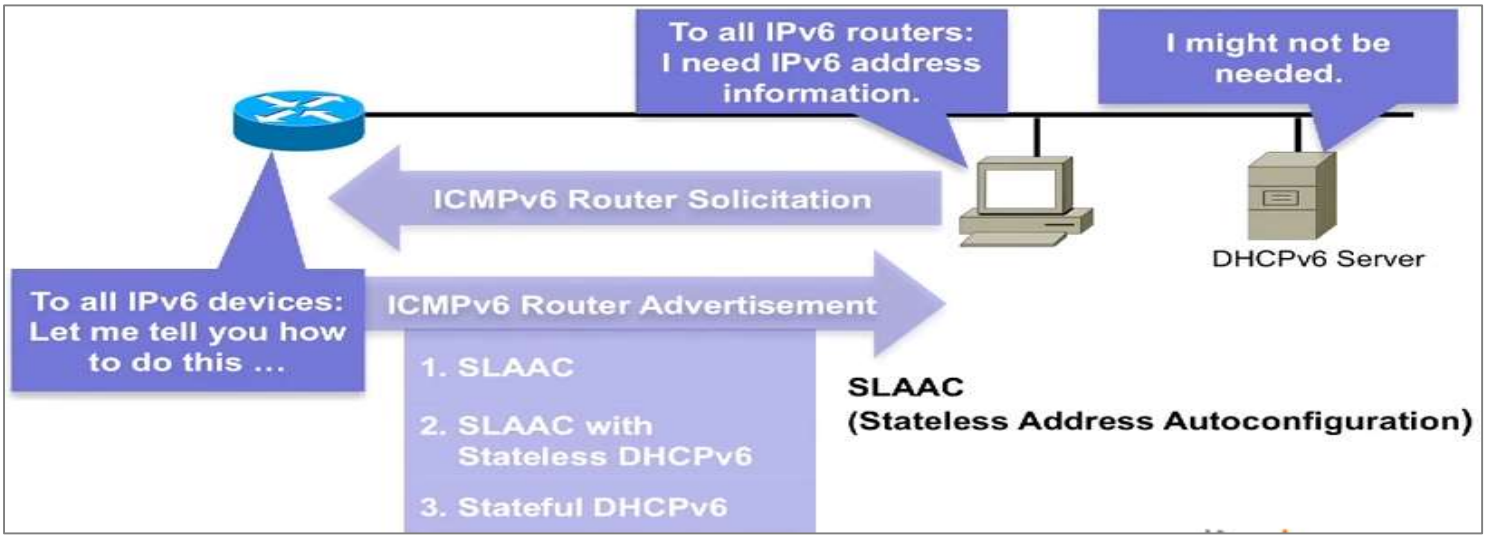

Fig. Dynamic Address Allocation

2. Router Advertisement (RA - 134) – Router – to - Device Messaging – *sent every time or in response to RS*

IPv6 routers send unsolicited Router Advertisement messages pseudo-periodically—that is, the interval between unsolicited advertisements is randomized to reduce synchronization issues when there are multiple advertising routers on a link—and solicited Router Advertisement messages in response to the receipt of a Router Solicitation message. The Router Advertisement message contains the information required by hosts to determine the link prefixes, the link MTU, specific routes, whether or not to use address autoconfiguration, and the duration for which addresses created through address autoconfiguration are valid and preferred.

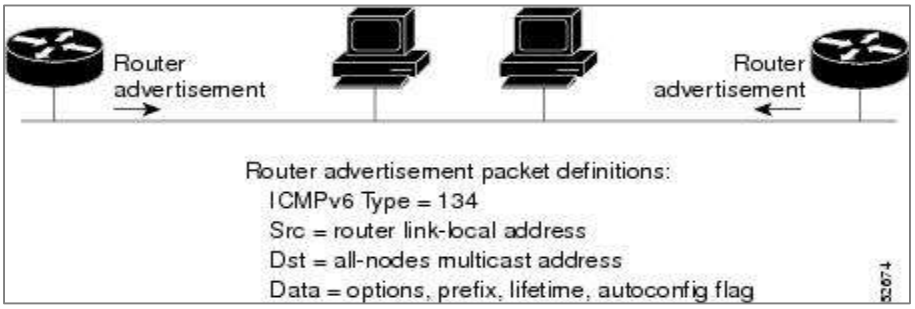

Working Principle

- RS messages are originated only by the hosts.
- RS messages are originated by hosts to find the Routers on the link.
- Routers respond to RS message by sending an RA.

The RA messages are sent to the all-nodes multicast address

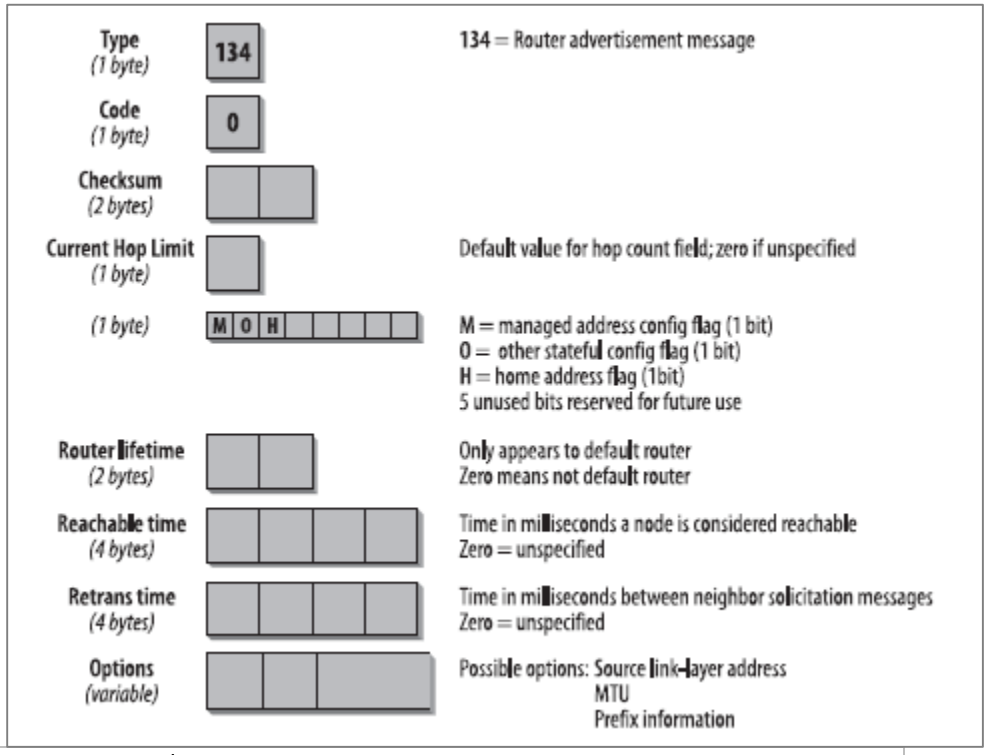

Fig. Structure of the Router Advertisement message

Table 134: ICMPv6 *Router Advertisement* Message Format

Field Name	Size (bytes)	Description
Type	1	***Type:*** Identifies the ICMPv6 message type; for *Router Advertisement* messages the value is 134.
Code	1	***Code:*** Not used; set to 0.
Checksum	2	***Checksum:*** 16-bit checksum field for the ICMP header,
Cur Hop Limit	1	***Current Hop Limit:*** This is a default value the router is telling hosts on the local network they should put in the *Hop Limit* field of datagrams they send. If 0, the router is not recommending a *Hop Limit* value in this *Router Advertisement*.

<table>
<tr><td>Autoconfig Flags</td><td>1</td><td>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:
<table>
<tr><th>Subfield Name</th><th>Size (bytes)</th><th>Description</th></tr>
<tr><td>M</td><td>1/8 (1 bit)</td><td>Managed Address Configuration Flag: When set, this flag tells hosts to use an administered or "stateful" method for address autoconfiguration, such as DHCP.</td></tr>
<tr><td>O</td><td>1/8 (1 bit)</td><td>Other Stateful Configuration Flag: When set, tells hosts to use an administered or "stateful" autoconfiguration method for information other than addresses.</td></tr>
<tr><td>Reserved</td><td>6/8 (6 bits)</td><td>Reserved: Reserved for future use; sent as zeroes.</td></tr>
</table></td></tr>
<tr><td>Router Lifetime</td><td>2</td><td>Router Lifetime: Tells the host receiving this message how long, in seconds, this router should be used as a default router. If 0, tells the host this router should not be used as a default router.</td></tr>
<tr><td>Reachable Time</td><td>4</td><td>Reachable Time: Tells hosts how long, in milliseconds, they should consider a neighbor to be reachable after they have received reachability confirmation.</td></tr>
<tr><td>Retrans Timer</td><td>4</td><td>Retransmission Timer: The amount of time, in milliseconds, that a host should wait before retransmitting Neighbor Solicitation messages.</td></tr>
<tr><td>Options</td><td>Variable</td><td>Options: Router Advertisement messages may contain three possible options
▯ Source Link-Layer Address: Included when the router sending the Advertisement knows its link-layer (layer two) address.
▯ MTU: Used to tell local hosts the MTU of the local network when this information may not be known by hosts on the network.
▯ Prefix Information: Tells local hosts what prefix or prefixes to use for the local network.</td></tr>
</table>

3. Neighbor Solicitation (NS - 135) - similar as IPv4 ARP Request **– Device – tot - Device Messaging**

IPv6 nodes send the Neighbor Solicitation message to discover the link-layer address of an on-link IPv6 node or to confirm a previously determined link-layer address. It typically includes the link-layer address of the sender. Typical Neighbor Solicitation messages are multicast for address resolution and unicast when the reach ability of a neighboring node is being verified.

Working Principal

- NS messages are originated by the nodes.
- NS messages are originated by nodes to request the link layer address of another node.
- NS messages are also used for duplicate address detection and neighbor unreachability detection.

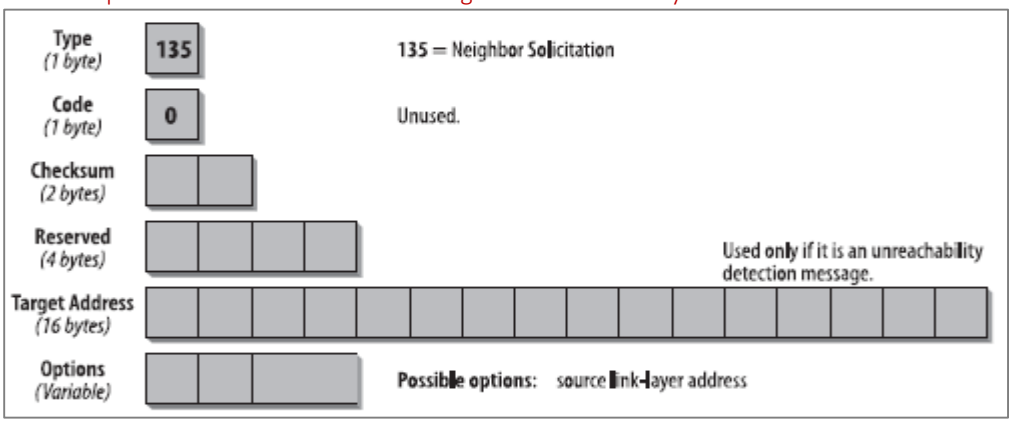

Fig. Structure of Neighbor Solicitation Message

Table 135: ICMPv6 *Neighbor Solicitation* Message Format		
Field Name	Size (bytes)	Description

Type	1	***Type:*** Identifies the ICMPv6 message type; for *Neighbor Solicitation* messages the value is 135.
Code	1	***Code:*** Not used; set to 0.
Checksum	2	***Checksum:*** 16-bit checksum field for the ICMP header
Reserved	4	***Reserved:*** 4 reserved bytes set to 0.
Target Address	16	***Target Address:*** The IPv6 address of the target of the solicitation. For IPv6 address resolution, this is the actual unicast IP address of the device whose layer two (link-layer) address we are trying to resolve.
Options	Variable	***Options:*** If the device sending the *Neighbor Solicitation* knows both its own IP address and layer two address, it should include the layer two address in a *Source Link-Layer Address* option. The inclusion of this option will allow the destination of the *Neighbor Solicitation* to enter the layer two and layer three addresses of the source of this message into its own address cache.

Example: Description of Figure

When a node wants to determine the link-layer address of another node, the source address in a neighbor solicitation message is the IPv6 address of the node sending the neighbor solicitation message. The destination address in the neighbor solicitation message is the solicited-node multicast address that corresponds to the IPv6 address of the destination node. The neighbor

After receiving the neighbor solicitation message, the destination node replies by sending a neighbor advertisement message, which has a value of 136 in the Type field of the ICMP packet header, on the local link. The source address in the neighbor advertisement message is the IPv6 address of the node (more specifically, the IPv6 address of the node interface) sending the neighbor advertisement message. The destination address in the neighbor advertisement message is the IPv6 address of the node that sent the neighbor solicitation message. The data portion of the neighbor advertisement message includes the link-layer address of the node sending the neighbor advertisement message.

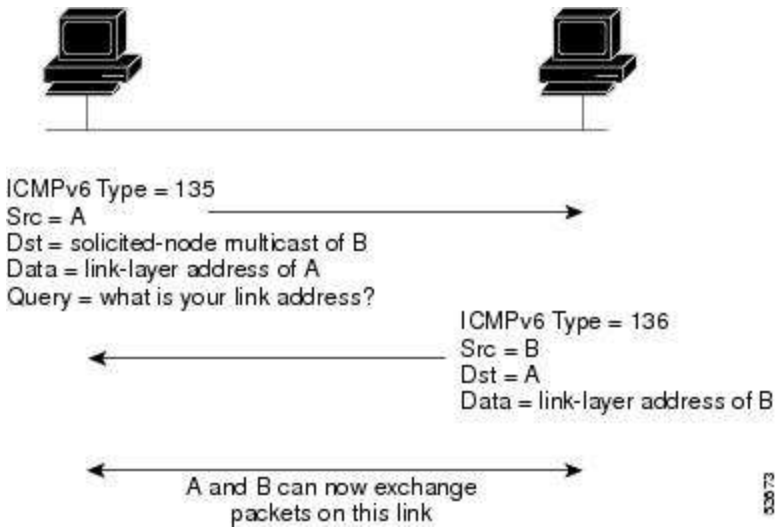

After the source node receives the neighbor advertisement, the source node and destination node can communicate.

4. Neighbor Advertisement (NA - 136) – similar as IPv4 ARP Reply **Device – to - Device Messaging**

An IPv6 node sends the Neighbor Advertisement message in response to a Neighbor Solicitation message. An IPv6 node also sends unsolicited Neighbor Advertisements to inform neighboring nodes of changes in link-layer addresses or the node's role. The Neighbor Advertisement contains information required by nodes to determine the type of Neighbor Advertisement message, the sender's role on the network, and typically the link-layer address of the sender.

Working Principal

- 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.

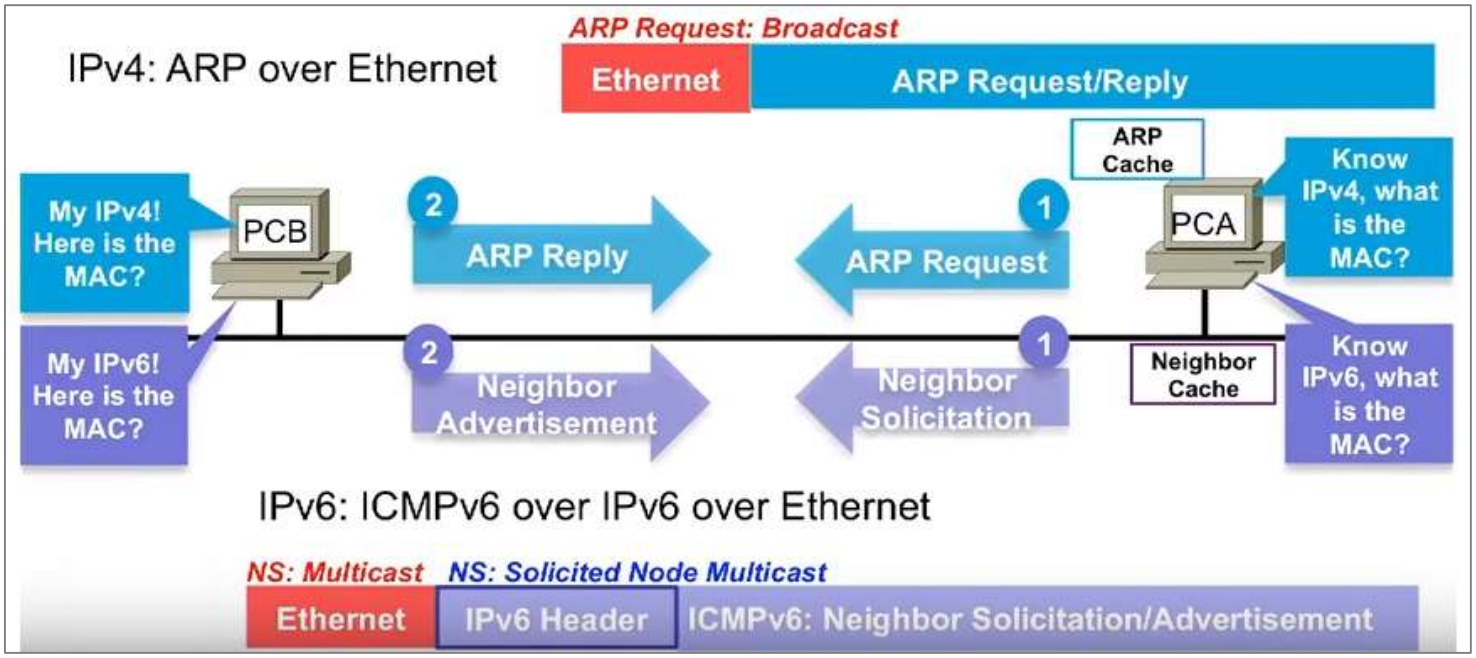

Fig. Address Resolution of IPv4 and IPv6

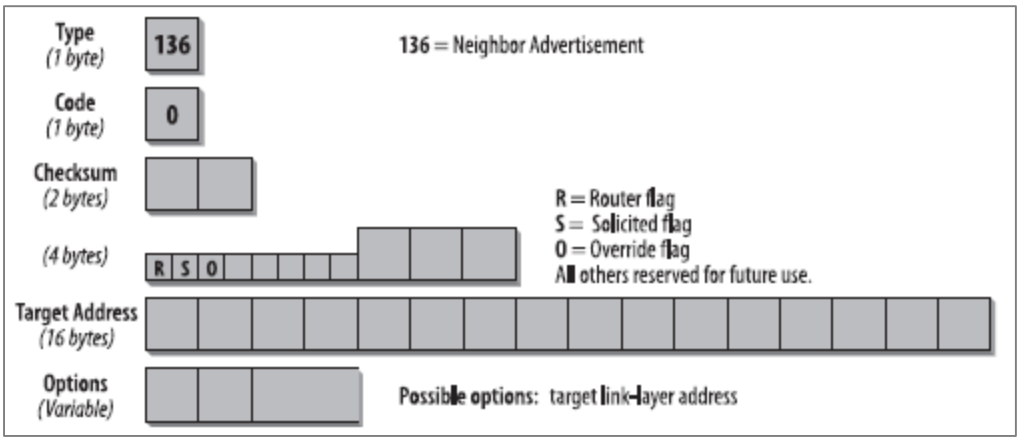

Fig. Structure of **Neighbor Advertisement Message**

Table 136: ICMPv6 *Neighbor Advertisement* Message Format		
Field Name	**Size (bytes)**	**Description**
Type	1	***Type:*** Identifies the ICMPv6 message type; for *Neighbor Advertisement* messages the value is 136.
Code	1	***Code:*** Not used; set to 0.
Checksum	2	***Checksum:*** 16-bit checksum field for the ICMP header

<table>
<tr><td>Flags</td><td>4</td><td>Flags: Three flags that convey information about the message (and lots of empty space for future use):
<table>
<tr><th>Subfield Name</th><th>Size (bytes)</th><th>Description</th></tr>
<tr><td>R</td><td>1/8 (1 bit)</td><td>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.</td></tr>
<tr><td>S</td><td>1/8 (1 bit)</td><td>Solicited Flag: When set, indicates that this message was sent in response to a Neighbor Solicitation message. Cleared for unsolicited Neighbor Advertisements.</td></tr>
<tr><td>O</td><td>1/8 (1 bit)</td><td>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.</td></tr>
<tr><td>Reserved</td><td>3 5/8 (29 bits)</td><td>Reserved: A big whomping set of reserved bits. ☺</td></tr>
</table></td></tr>
<tr><td>Target Address</td><td>16</td><td>Target Address: If the Neighbor Advertisement is being sent in response to a Neighbor Solicitation, this is the same value as in the Target Address field of the Solicitation. This field will commonly contain the IPv6 address of the device sending the Neighbor Advertisement, but not in all cases. For example, if a device responds as a proxy for the target of the Neighbor Solicitation, the Target Address field contains the address of the target, not the device sending the response.

If the Neighbor Advertisement is being sent unsolicited, then this is the IPv6 address of the device sending it.</td></tr>
<tr><td>Options</td><td>Variable</td><td>Options: When sent in response to a multicast Neighbor Solicitation, a Neighbor Advertisement message must contain a Target Link-Layer Address option, which carries the link-layer address of the device sending the message. This is a good example of an “option” that's not really “optional”. J

When the Neighbor Advertisement is sent in response to a unicast Neighbor Solicitation, this option is technically not required (since the sender of the Solicitation must already have the target's link-layer address to have sent it unicast.) Despite this, it is still normally included, to ensure that the link-layer address of the target is refreshed in the cache of the device that sent the Neighbor Solicitation.</td></tr>
</table>

5. Redirect - 137

The Redirect message is sent by an IPv6 router to inform an originating host of a better first hop address for a specific destination. Redirect messages are sent only by routers for unicast traffic, are unicast only to originating hosts, and are processed only by hosts.

If a local network has only a single router, then it will send all such non-local traffic to that router. If it has more than one local router, the host then must decide which router to use for which traffic. In general terms, a host will not know the most efficient choice of router for every type of datagram it may need to send.

When a router receives datagrams destined for certain networks, it may realize that it would be more efficient if such traffic was sent by a host to a different router on the local network. If so, it will invoke the Redirect function by sending an ICMPv6 Redirect message to the device that sent the original datagram

Working Principal

- Redirect messages are used in the same way as IPv4 ICMP redirect messages.
- Redirect messages are always sent by the router to a host asking the host to update its routing information.
- Upon receiving a packet from a host, the router can send Redirect message back to the host only when a router knows that the best path for that host to reach the destination is another router and not itself. On receiving the Redirect message, the host can update its routing information, and send subsequent packets directly to the other router.

Structure of Redirect Message

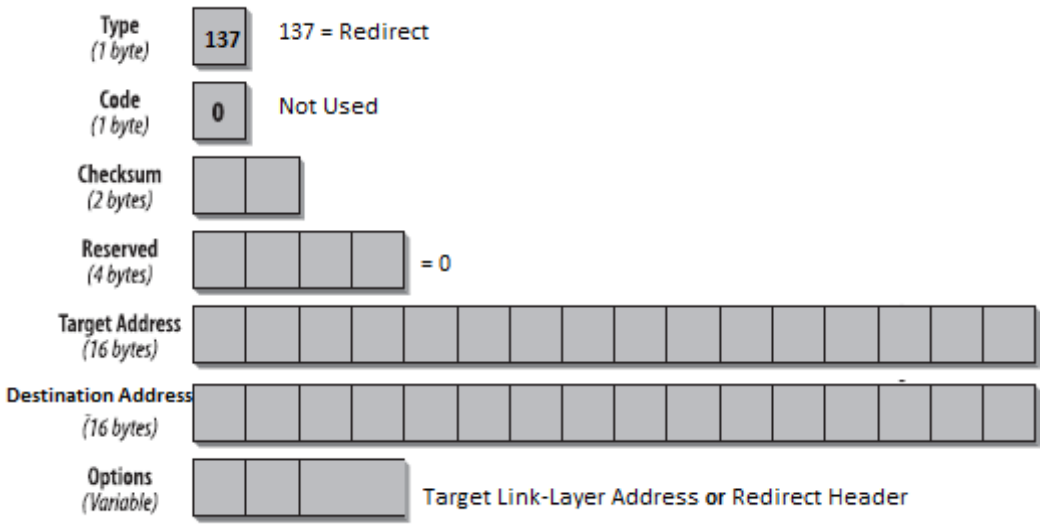

Fig. Structure of Redirect Message

Table 136: ICMPv6 *Redirect* Message Format		
Field Name	Size (bytes)	Description
Type	1	***Type:*** Identifies the ICMPv6 message type; for *Redirect* messages the value is **137**.
Code	1	***Code:*** Not used; set to 0.
Checksum	2	***Checksum:*** 16-bit checksum field for the ICMP header
Reserved	4	***Reserved:*** 4 bytes sent as zeroes.
Target Address	16	***Target Address:*** The address of the router that the router creating the *Redirect* is telling the recipient of the *Redirect* to use as a first hop for future transmissions to the destination. Phew. Example time: if router *R2*generated a *Redirect* telling host *A* that in the future transmissions to host *B* should be sent first to router *R1*, then *R1*'s IPv6 address would be in this field.
Destination Address	16	***Destination Address:*** The address of the device whose future transmissions are being redirected; this is the destination of the datagram that originally led to the *Redirect* being generated. Repeating the example above: if router *R2* generated a *Redirect* telling host *A* that in the future transmissions to host *B* should be sent first to router *R1*, then host *B*'s IPv6 address would be in this field.
Options	Variable	***Options:*** Redirect messages normally include two ICMPv6 option fields: [***Target Link-Layer Address:*** The layer-two address of the *Target Address*, if known. This saves the recipient of the *Redirect* message from needing to perform an address resolution on the target. [***Redirected Header:*** As much of the IPv6 datagram that spawned this *Redirect* as will fit without causing the size of the ICMPv6 error message (including its own IP header) to exceed the minimum IPv6 maximum transmission unit (MTU) of 1280 bytes.

Neighbor Discovery Processes

The ND protocol provides message exchanges for the following processes:

1. Address resolution (including duplicate address detection)

The address resolution process for IPv6 nodes consists of an exchange of Neighbor Solicitation and Neighbor Advertisement messages to resolve the link-layer address of the on-link next-hop address for a given destination. The sending host sends a multicast Neighbor Solicitation message on the appropriate interface. The multicast address of the Neighbor Solicitation message is the solicited-node multicast address derived from the target IP address. The Neighbor Solicitation message includes the link-layer address of the sending host in the Source Link-Layer Address option. For information about how a host determines the next-hop address for a destination, see "Host Sending Algorithm" in this chapter.

When the target host receives the Neighbor Solicitation message, it updates its own neighbor cache based on the source address of the Neighbor Solicitation message and the link-layer address in the Source Link-Layer Address option. Next, the target node sends a unicast Neighbor Advertisement to the Neighbor Solicitation sender. The Neighbor Advertisement includes the Target Link-Layer Address option.

After receiving the Neighbor Advertisement from the target, the sending host updates its neighbor cache with an entry for the target based on the information in the Target Link-Layer Address option. At this point, unicast IPv6 traffic between the sending host and the target of the Neighbor Solicitation can be sent.

Address Resolution Example

Host A has an Ethernet MAC address of 00-10-5A-AA-20-A2 and a corresponding link-local address of FE80::210:5AFF:FEAA:20A2. Host B has an Ethernet MAC address of 00-60-97-02-6E-A5 and a corresponding link-local address of FE80::260:97FF:FE02:6EA5. To send a packet to Host B, Host A must first use address resolution to resolve Host B's link-layer address.

2. Router discovery (includes prefix and parameter discovery)

Router discovery is the process through which nodes attempt to discover the set of routers on the local link. Router discovery in IPv6 is similar to ICMP router discovery for IPv4 described in RFC 1256. ICMP router discovery is a set of ICMP messages that allow IPv4 hosts to determine the presence of local routers, determine which local router is automatically configured as a default gateway, and to automatically switch to a different router as their default gateway when the current default gateway becomes unavailable.

An important difference between ICMPv4 router discovery and IPv6 router discovery is the mechanism through which a new default router is selected when the current one becomes unavailable. In ICMPv4 router discovery, the Router Advertisement message includes an Advertisement Lifetime field. Advertisement Lifetime is the time after which the router can be considered unavailable. In the worst case, a router can become unavailable and hosts will not attempt to discover a new default router until the Router Advertisement time has elapsed.

IPv6 has a Router Lifetime field in the Router Advertisement message. This field indicates the length of time that the router can be considered a default router. However, if the current default router becomes unavailable, the condition is detected through neighbor unreachability detection instead of the Router Lifetime field in the Router Advertisement message. Because neighbor unreachability detection determines that the router is no longer reachable, a new router is chosen immediately from the default router list or the host sends a Router Solicitation message to determine if additional default routers are present on the link. For more information, see the "Neighbor Unreachability Detection" section in this chapter.

In addition to configuring a default router, IPv6 router discovery also configures the following:

- The default setting for the Hop Limit field in the IPv6 header.
- A determination of whether the node should use a stateful address protocol, such as Dynamic Host Configuration Protocol for IPv6 (DHCPv6), for addresses and other configuration parameters.
- The timers used in neighbor unreachability detection and the retransmission of Neighbor Solicitations.
- The list of network prefixes defined for the link. Each network prefix contains both the IPv6 network prefix and its valid and preferred lifetimes. If indicated, a network prefix combined with the interface identifier creates a stateless IP address configuration for the receiving interface. A network prefix also defines the range of addresses for nodes on the local link.
- The MTU of the local link.
- Specific routes to add to the routing table.

The IPv6 router discovery processes are the following:

- IPv6 routers pseudo-periodically send a Router Advertisement message on the local link advertising their existence as routers. They also provide configuration parameters such as default hop limit, MTU, prefixes, and routes. For more information about how often routers send pseudo-periodic router advertisements, see section 6.2.4 of RFC 2461.
- Active IPv6 hosts on the local link receive the Router Advertisement m essages and use the contents to maintain the default router and prefix lists, autoconfigure addresses, add routes, and configure other parameters.
- A host that is starting sends a Router Solicitation message to the link-local scope all-routers multicast address (FF02::2). If the starting host is already configured with a unicast address, the Router Solicitation is sent with a unicast source address. Otherwise, the Router Solicitation is sent with an unspecified source address (::). Upon receipt of a Router Solicitation message, all routers on the local link send a Router Advertisement message to either the unicast address of the host that sent the Router Solicitation (if the source address of the Router Solicitation is a unicast address), or to the link-local scope all-nodes multicast address (FF02::1) (if the source address of the Router Solicitation message is unspecified). The host receives the Router Advertisement messages and uses their contents to build the default router and prefix lists and set other configuration parameters. The number of Router Solicitations sent before abandoning the router discovery process is set by a configurable variable. RFC 2461 uses the variable name of MAX_RTR_SOLICITATIONS and recommends a value of 3.

Router Discovery Example

Host A has the Ethernet MAC address of 00-B0-D0-E9-41-43. The router has an Ethernet MAC address of 00-10-FF-D6-58-C0 and a corresponding link-local address of FE80::210:FFFF:FED6:58C0. To forward packets to off-link destinations, Host A must discover the presence of the router.

3. Neighbor unreachability detection

A neighboring node is reachable if there has been a recent confirmation that IPv6 packets sent to the neighboring node were received and processed by the neighboring node. Neighbor unreachability does not necessarily verify the end-to-end reachability of the destination.

Because a neighboring node can be a host or router, the neighboring node might not be the final destination of the packet. Neighbor unreachability verifies only the reachability of the first hop to the destination.

One of the ways that reachability is confirmed is through the sending of a unicast Neighbor Solicitation message and the receipt of a solicited Neighbor Advertisement message. A solicited Neighbor Advertisement message, which has its Solicited flag set to 1, is sent only in response to a Neighbor Solicitation message. Unsolicited Neighbor Advertisement or Router Advertisement messages are not considered proof of reachability. The exchange of Neighbor Solicitation and Neighbor Advertisement messages confirms only the reachability of the node that sent the Neighbor Advertisement from the node that sent the Neighbor Solicitation. It does not confirm the reachability of the node that sent the Neighbor Solicitation from the node that sent the Neighbor Advertisement.

For example, if Host A sends a unicast Neighbor Solicitation to Host B and Host B sends a solicited unicast Neighbor Advertisement to Host A, Host A considers Host B reachable. Because there is no confirmation in this exchange that Host A actually received the Neighbor Advertisement, Host B does not consider Host A reachable. To confirm reachability of Host A from Host B, Host B must send its own unicast Neighbor Solicitation to Host A and receive a solicited unicast Neighbor Advertisement from Host A.

4. Redirect function

Routers use the redirect function to inform originating hosts of a better first-hop neighbor to which traffic should be forwarded for a specific destination. There are two instances where redirect is used:

1. A router informs an originating host of the IP address of a router available on the local link that is "closer" to the destination. "Closer" is a routing metric function used to reach the destination network segment. This condition can occur when there are multiple routers on a network segment, and the originating host chooses a default router and it is not the better ("closer") one to use to reach the destination.
2. A router informs an originating host that the destination is a neighbor (it is on the same link as the originating host). This condition can occur when the prefix list of a host does not include the prefix of the destination. Because the destination does not match a prefix in the list, the originating host forwards the packet to its default router.

The following steps occur in the IPv6 redirect process:

1. The originating host forwards a unicast packet to its default router.
2. The router processes the packet and notes that the address of the originating host is a neighbor. Additionally, it notes that both the originating host's address and the next-hop address are on the same link.
3. The router sends the originating host a Redirect message. In the Target Address field of the Redirect message is the next-hop address of the node to which the originating host should send subsequent packets addressed to the destination.
4. The router forwards the packet to the appropriate next-hop address, using address resolution if needed to obtain the link-layer address of the next hop.
   For packets redirected to a router, the Target Address field is set to the link-local address of the router. For packets redirected to a host, the Target Address field is set to the destination address of the packet originally sent.
   The Redirect message includes the Redirected Header option. It might also include the Target Link-Layer Address option.
5. Upon receipt of the Redirect message, the originating host updates the destination address entry in the destination cache with the address in the Target Address field. If the Target Link-Layer Address option is included in the Redirect message, its contents are used to create or update the corresponding neighbor cache entry.

Redirect messages are sent only by the first router in the path between the originating host and the destination. Hosts never send Redirect messages and routers never update routing tables based on the receipt of a Redirect message. Like ICMPv6 error messages, Redirect messages are rate limited.

Redirect Example

Host A has the Ethernet MAC address of 00-AA-00-11-11-11 and a corresponding link-local address of FE80::2AA:FF:FE11:1111. Host A also has the site-local address of FEC0::1:2AA:FF:FE11:1111. Router 2 has the Ethernet MAC address of 00-AA-00-22-22-22 and a corresponding link-local address of FE80::2AA:FF:FE22:2222. Router 2 also has the site-local address of FEC0::1:2AA:FF:FE22:2222. Router 3 has the Ethernet MAC address of 00-AA-00-33-33-33 and a corresponding link-local address of FE80::2AA:FF:FE33:3333. Router 3 also has the site-local address of FEC0::1:2AA:FF:FE33:3333. Host A sends a packet to an off-link host at FEC0::2:2AA:FF:FE99:9999 (not shown) and uses Router 2 as its current default router. However, Router 3 is the better router to use to reach this destination.

2.2.5 Path MTU (PMTU) Discovery

Path MTU Discovery (**PMTUD**) is a standardized technique in computer networking for determining the maximum transmission unit (MTU) size on the network path between two Internet Protocol (IP) hosts, usually with the goal of avoiding IP fragmentation. PMTUD was originally intended for routers in Internet Protocol Version 4 (IPv4). However, all modern operating systems use it on endpoints. In IPv6, this function has been explicitly delegated to the end points of a communications session.

Sending the largest possible packets maximizes efficient use of the network capacity when bulk data transfer is performed. Because IPv6 routers no longer support fragmentation, the sending host must either fragment its payload or discover the maximum sized packet that can be send to the destination and send unfragmented packets at that size.

The PMTU is the smallest link MTU supported by any link in the path between a source and a destination. The link MTU is the maximum-sized link-layer payload that can be sent on the link. This corresponds to the maximum-sized packet that can be sent on the link but it differs from the maximum-sized frame that can be sent on the link. The maximum-sized frame includes the link-layer header and trailer.
For example:
For Ethernet link using Ethernet II encapsulation, the link MTU is 1500 bytes and the maximum-sized frame is 1526 bytes
PMTU is discovered through the following process: -

1. The sending node assumes that the destination PMTU is the MTU of the interface on which the traffic is being forwarded.
2. The sending node sends IPv6 packets at the assumed PMTU size
3. If a router on the path is unable to forward the packet because the forwarding interface ha a link MTU that is smaller than the size of that packet, it sends an ICMPv6 packet Too Big message back to the sending node and discard the IPv6 packet.
4. The sending node sets the new assumed PMTU for packets being sent to the destination to the value of the MTU field in the IVMPv6 Packet Too Big message.
5. The sending node starts again at step 2 and repeats steps 2 through 4.

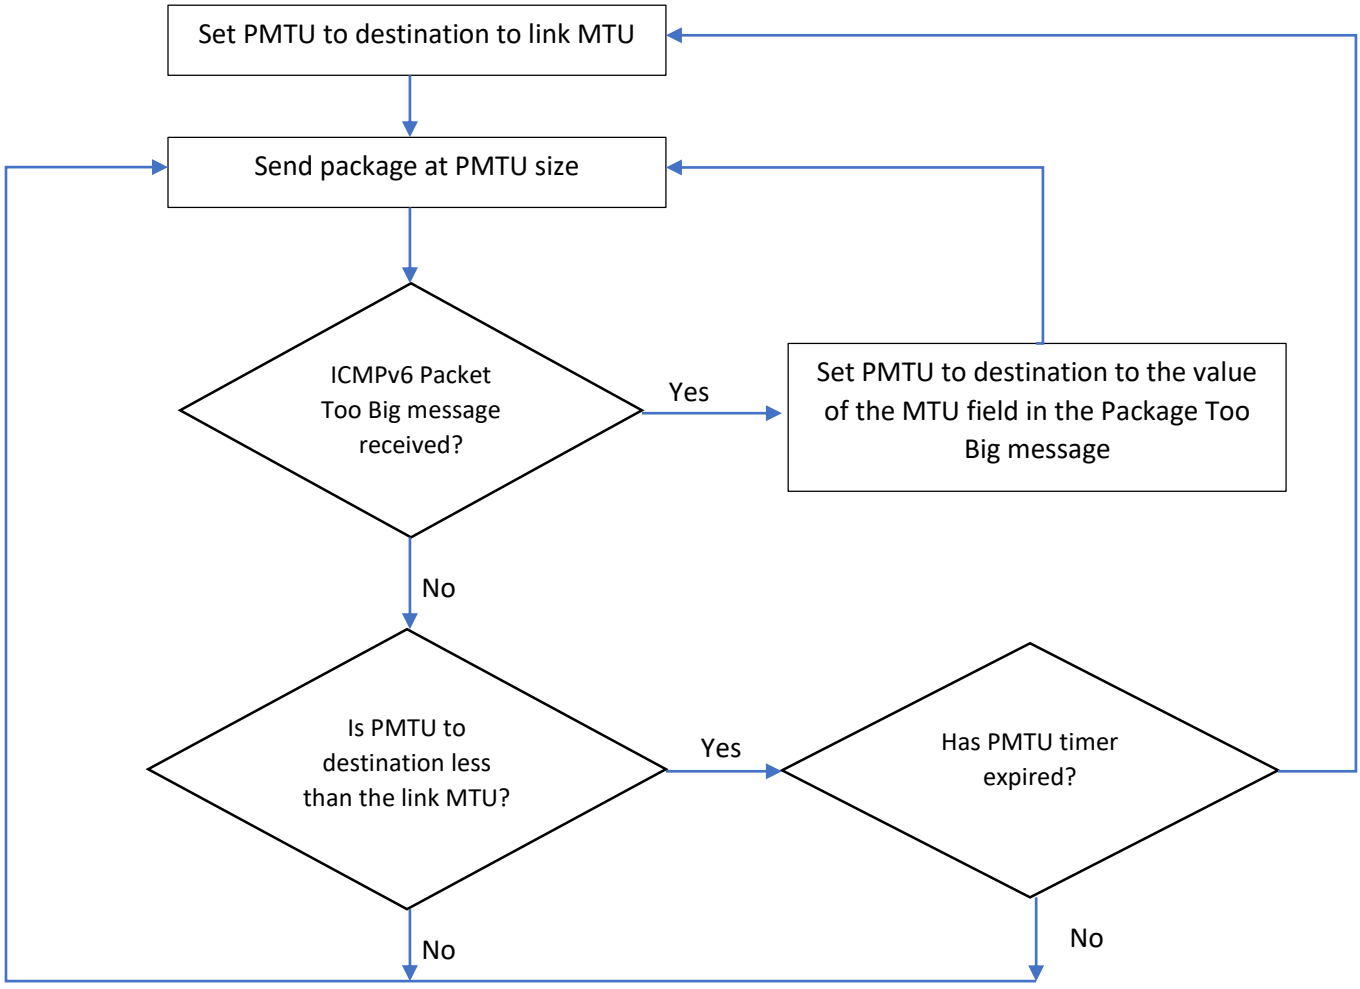

Fig. The PMTU Discovery Process