

IPv6: Deep Dive SHARE Session 13232



Laura Knapp
WW Business Consultant
Laurak@aesclever.com

What is IPv6?

Addressing

128-bit addresses hierarchically assigned

Routing

Strongly hierarchical (route aggregation)

Performance

Simple datagram

Extensibility

New flexible option header format

Improved support for extensions and options

Multimedia

Better support for QoS

Multicast

Compulsory-better scope control

Security

Built in security (IPSEC)

Auto-configuration

Stateless and state-full address configuration

Mobility

Better efficiency and security



IPv6 Header

IPv4 Header

Vers:HD	TOS	Payload length
Fragment ID		Fragment Info.
TTL	Protocol	Header Checksum
Source Address		
Destination Address		

IPv6 Header

Vers:Class	Flow Label	
Payload Length	Next hdr	Hop limit
Source Address		
Destination Address		

- IPv4 header is 20 bytes: IPv6 header is 40 bytes
- Address increased from 32 to 128 bits
- Fragmentation fields moved out of base header
- Header checksum
- Time to Live replaced with 'Hop Limit'
- Protocol replaced with 'Next Header'
- TOS replaced with 'Flow Label'
- Alignment changed from 32 to 64 bits

Items to Be Discussed


- **IP Addressing**
- **ICMPv6**
 - Error Messages
 - Informational Messages
 - Neighbor Discovery Protocol
 - Multicast Listener Discovery Protocol
 - Packet MTU Size
 - Fragmentation
 - Other ICMPv6 functions



Addressing Format

1080:0002:4544:0000:8532:9A14:0648:417A

IPv6



Format Prefix are the high order bits with fixed values

- **Defined in RFC 3513:**
 - 40,282,366,920,938,463,374,607,431,768,211,456 addresses
 - 40 trillion trillion trillion addresses
- **Addresses are assigned to interfaces**
- **Multiple address can be defined to a single interface**
- **Address structure**
 - Ipv6 address = Prefix + Interface id
- **Separation of 'who you are' from 'where you are connected'**
- **Assignments by ARIN, APNIC, RIPE**

IPv6 Address Types


unicast:

for one-to-one communication




multicast:

for one-to-many communication




anycast:

for one-to-nearest communication




Link Local Address



- FE80 prefix
- Similar to IPv4 APIPA (169.254.0.0/16)
- Only for on-link communication, not routable
- Used for
 - Auto configured addresses
 - Neighbor discovery process

Multicast Address



Flags

0: well known address, 1: transient address

Scope

1: Node Local (FF01::1), 2: Link Local (FF02::1)


All routers group: FF02::2)

Group ID

1: All nodes, 2: All routers, 101: all NTP servers

- Multicast replaces Broadcast
- All IPv6 nodes must support multicast
- You must enable IGMP snooping

Global Unicast Address




Address Type	Binary Prefix	Prefix
Unspecified	000...0	::/128
Loopback	0000...01	::1/128
ULA	1111 110	FC00::/7
Assigned to RIRs	001	2003:/3
Global Unicast	Everything else!!	

Korea: 2001:0200 – 099F

ATT: 2001:0408/32

Verizon: 2001:0506:0000/48

Unique Local Address (ULA)



- L=1
- FC00::/7 prefix
- Local or site local communications
- Most likely will be unique and not expected to be routable
- Well known, somewhat like the RFC1918

Windows and IPv6

IPv6 is preferred

Nameserver query

Try to reach IPv6

Try to reach IPv4

Timeout

```
Wireless LAN adapter Wireless Network Connection 2:

    Media State . . . . . : Media disconnected
    Connection-specific DNS Suffix  . :

Wireless LAN adapter Wireless Network Connection:

    Media State . . . . . : Media disconnected
    Connection-specific DNS Suffix  . :

Ethernet adapter Local Area Connection:

    Connection-specific DNS Suffix  . : hawaii.rr.com
    Link-local IPv6 Address . . . . . : fe80::6947:83b1:88e2:73d4%13
    IPv4 Address. . . . . : 192.168.1.146
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 192.168.1.1

Ethernet adapter Bluetooth Network Connection:

    Media State . . . . . : Media disconnected
    Connection-specific DNS Suffix  . :

Tunnel adapter Teredo Tunneling Pseudo-Interface:

    Connection-specific DNS Suffix  . :
    IPv6 Address. . . . . : 2001:0:4137:9e76:10f7:1e4b:9d69:43f8
    Link-local IPv6 Address . . . . . : fe80::10f7:1e4b:9d69:43f8%15
    Default Gateway . . . . . : ::

Tunnel adapter isatap.hawaii.rr.com:

    Media State . . . . . : Media disconnected
    Connection-specific DNS Suffix  . : hawaii.rr.com
```

Address Type Prefixes

- Unspecified
 - ▶ used when there is no address

0000 0000 (::/128)

- Loopback

0000 0001 (::1/128)

- Link Local Unicast

1111 1110 1000 0000 (fe80::/16)

- Multicast

1111 1111 (ffxx::/8)

- Unicast + Anycast

- hierarchical
- /13 - /32 to LIRs (ISPs)
- /48 or /56 to end-users / sites

The rest, 2000::/3, which is 1/8th of total IPv6 space

2001::/16 = RIRs

2001::/32 = Teredo

2002::/16 = 6to4

3ffe::/16 = 6bone*

fd00::/8 = ULA

* = 6bone shut down on 6/6/6

- ▶ “Site Local” used to exist (fec0::/10) but this has been deprecated in favor of ULA

<http://www.iana.org/assignments/ipv6-address-space>

Items to Be Discussed

- IP Addressing
- **ICMPv6**
 - Error Messages
 - Informational Messages
 - Neighbor Discovery Protocol
 - Multicast Listener Discovery Protocol
 - Packet MTU Size
 - Fragmentation
 - Other ICMPv6 functions



IPv6: Autoconfiguration

Combination

ARP : ICMP router discovery : ICMP redirect

Neighbor discovery

Multicast and unicast datagrams

Establishes MAC address on same network

ICMPv6 router solicitation

ICMPv6 router advertisement

ICMPv6 neighbor solicitation

ICMPv6 redirect

ICMPv6 includes IGMP protocol for Multicast IP

Reduces impact of finding hosts

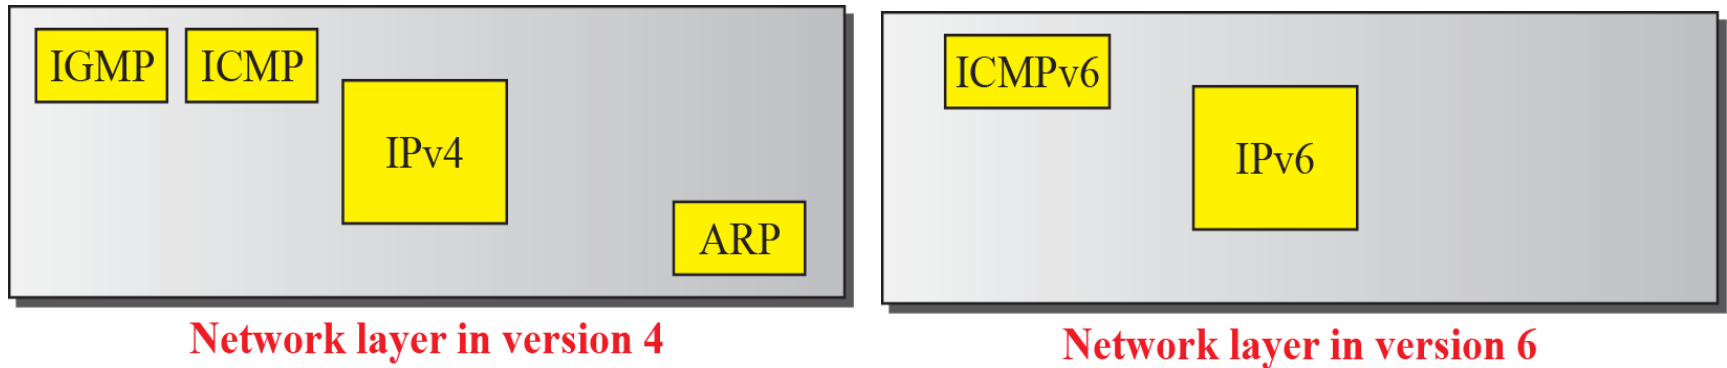
Stateless: router configures a host with IPv6 address

Stateful: DHCP for IPv6

Link Local Address: IPv6 connectivity on isolated LANs



ICMPv4 and ICMPv6 Quick View



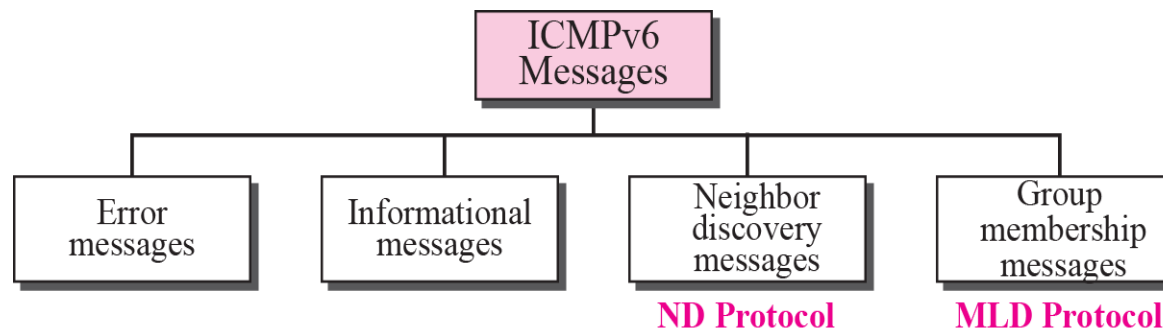
ICMPv6 is more complicated than ICMPv4.

Protocol consolidation occurred in IPv6.

Additional messages have been added.

ICMPv6

- **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 is an integral part of IPv6 and MUST be fully implemented by every IPv6 node**
- **ICMPv6 messages are grouped into two classes:**
 - **error messages - Types 0-127**
 - **informational messages - Types 128-255**
- **IPv6 next 'header' value for ICMPv6 is 58**




ICMPv6 Functions

Reports:

- packet processing errors
- intranetwork communications path diagnosis
- multicast membership

New functions:

- Neighbor Discovery
 - allows nodes on the same link to discover each other
 - allows nodes to discover each other's addresses
 - finds routers for paths to other networks
 - determines the fully qualified name of a node
 - path MTU discovery determines the maximum path size along a path



ICMPv6 Header

Three Fields

Type (8 bits)

- Indicates the type of the message.
- If the high order bit = 0 (0- 127) → error message
- if the high-order bit = 1 (128 – 255) → information message.

Code (8 bits)

- content depends on the message type, and it is used to create an additional level of message granularity.

Checksum (16 bits)

- Used to detect errors in the ICMP message and in part of the IPv6 message.

MAC header	IPv6 header	ICMPv6 header	ICMPv6 message
------------	-------------	---------------	----------------

00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<u>Type</u>								<u>Code</u>								<u>Checksum</u>															
<u>ICMPv6 message</u>																															
...																															

ICMPv6 Messages

ICMPv6 messages are grouped into two classes:

- **Error messages**
 - To provide feedback to a source device about an error that has occurred.
 - Generated specifically in response to some sort of action, usually the transmission of a datagram
 - Identified as such by having a zero in the high-order bit of their message
 - Type field values 0 to 127.
- **Informational messages**
 - Used to let devices exchange information, implement certain IP-related features, and perform testing.
 - Message Types from 128 to 255.

Error messages

Type	Description	References
1	<u>Destination unreachable.</u>	<u>RFC 2463</u>
2	<u>Packet too big.</u>	<u>RFC 2463</u>
3	<u>Time exceeded.</u>	<u>RFC 2463</u>
4	<u>Parameter problem.</u>	<u>RFC 2463</u>

Informational messages

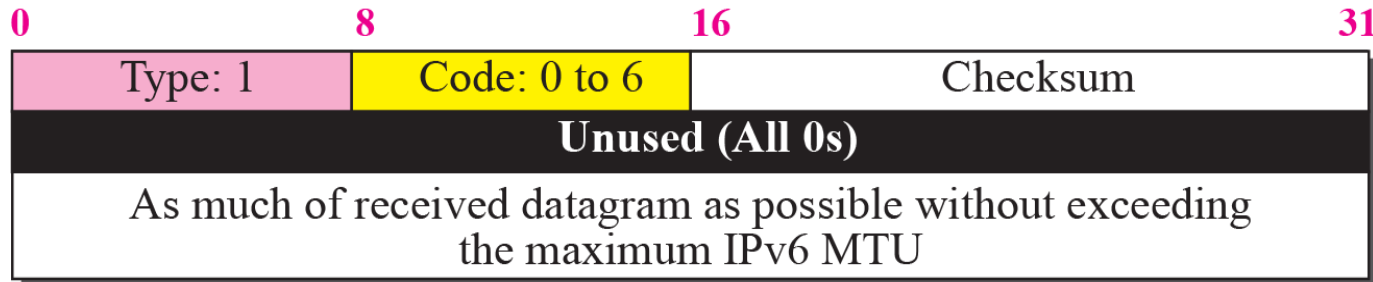
Type	Description	References
128	<u>Echo request.</u>	<u>RFC 2463</u>
129	<u>Echo reply.</u>	<u>RFC 2463</u>

Many of these ICMP types have a "code" field.

ICMPv6 Error Messages

Type Value	Message Name	Summary Description of Message Type
1	Destination Unreachable	Indicates that a datagram could not be delivered to its destination. <i>Code</i> value provides more information on the nature of the error.
2	Packet Too Big	Sent when a datagram cannot be forwarded because it's too big for the MTU of the next hop in the route. This message is only needed in IPv6 because routers cannot fragment oversized messages in IPv6, but they can in IPv4.
3	Time Exceeded	Sent when a datagram has been discarded prior to delivery because the <i>Hop Limit</i> field was reduced to zero.
4	Parameter Problem	Indicates a miscellaneous problem (specified by the <i>Code</i> value) in delivering a datagram.

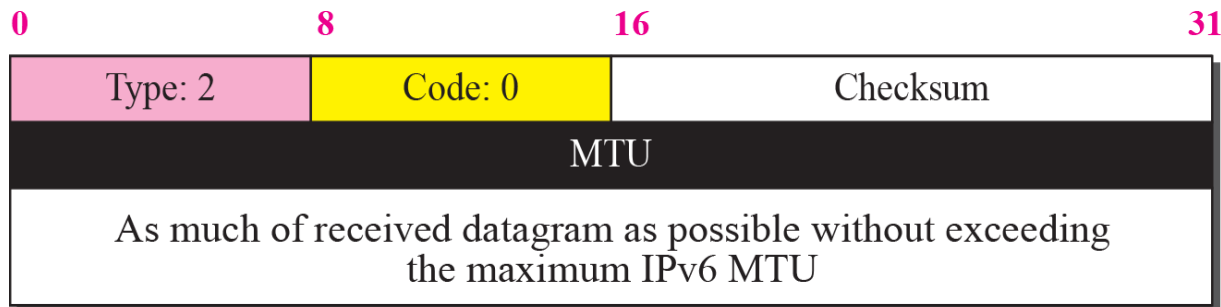
ICMPv6 Error Messages



ICMPv6 error messages:


- 1 Destination unreachable
 - code=0 no route to destination
 - code=1 communication with destination prohibited
 - code=2 (not assigned)
 - code=3 address unreachable
 - code=4 port unreachable
 - code=5 source address failed
 - code=6 reject route to destination

ICMPv6 Error Messages

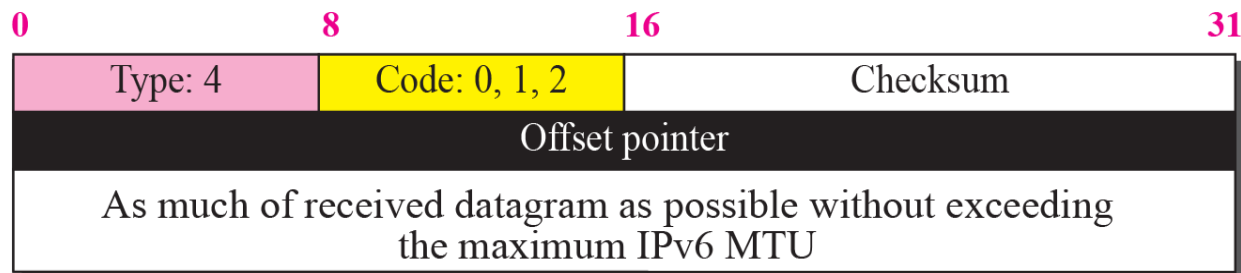


2 Packet too big

code=0 next byte contains the maximum transmission MTU of the next hop



3 Time exceeded



4 Parameter problem

code=0 erroneous header field encountered

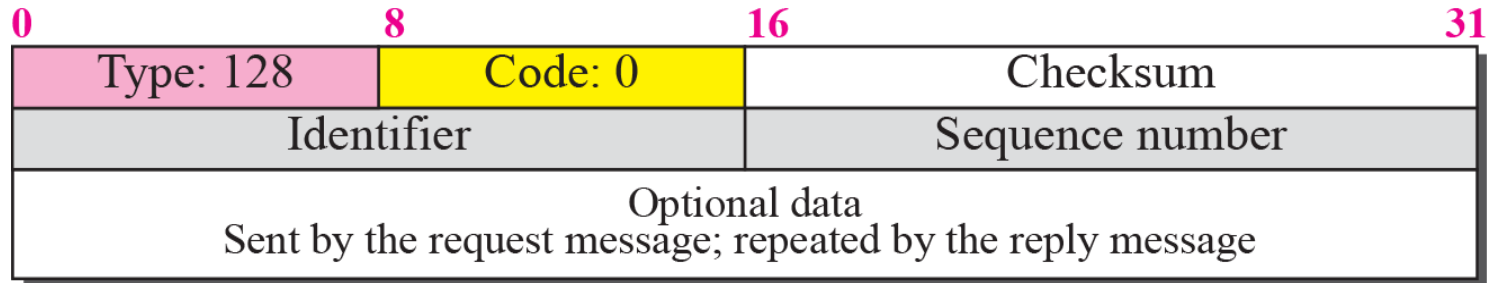
code=1 unrecognized next header type encountered

code=2 unrecognized IPv6 option encountered

ICMPv6 Informational Messages

ICMPv6 Informational Messages	128	<i>Echo Request</i>	Sent by a device to test connectivity to another device on the internetwork.	2463
	129	<i>Echo Reply</i>	Sent in reply to an <i>Echo (Request)</i> message; used for testing connectivity.	2463
	133	<i>Router Solicitation</i>	Prompts a router to send a <i>Router Advertisement</i> .	2461
	134	<i>Router Advertisement</i>	Sent by routers to tell hosts on the local network the router exists and describe its capabilities.	2461
	135	<i>Neighbor Solicitation</i>	Sent by a device to request the layer two address of another device while providing its own as well.	2461
	136	<i>Neighbor Advertisement</i>	Provides information about a host to other devices on the network .	2461
	137	<i>Redirect</i>	Redirects transmissions from a host to either an immediate neighbor on the network or a router.	2461
	138	<i>Router Renumbering</i>	Conveys renumbering information for router renumbering.	2894

ICMPv6 Informational Messages

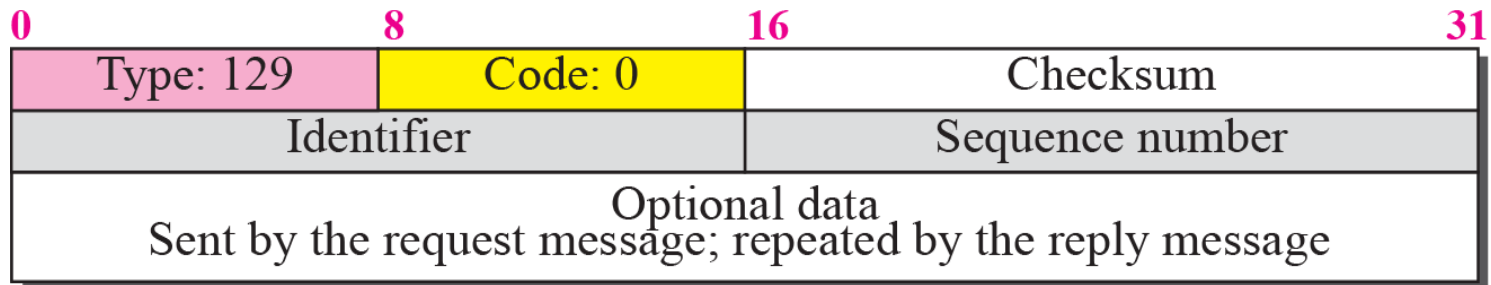


128 Echo request

code=0 and Identifier and sequence number carried

129 Echo reply

code=0 and identifier and sequence number carried



ICMPv6 Neighbor Discovery Protocol (NDP)

Defined in RFC 2461

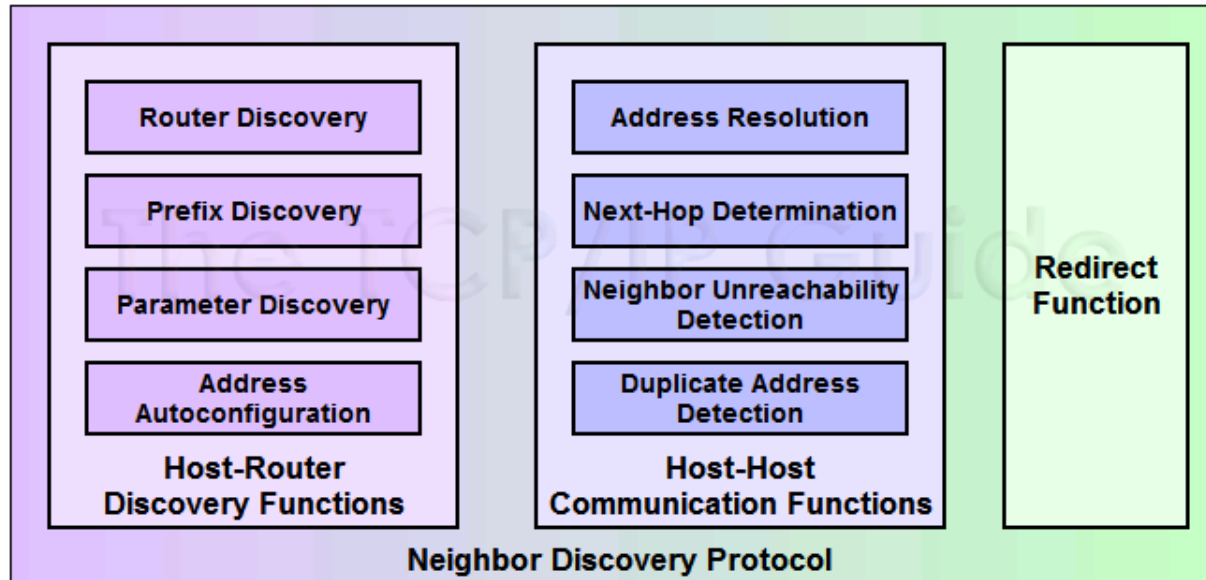
- Combines prior IPV4 functions
 - ARP (RFC 826)
 - Router Discovery (RFC 1256)
 - Redirect Message (RFC 792)

Mechanisms to:

- Discover routers
- Prefix discovery for on-link
- Parameter discovery (i.e link MTU)
- Address autoconfiguration
- Address resolution
- Next hop determination
- Neighbor unreachable
- Duplicate address
- Redirect



NDP Groups



Main three functions:

1. Host-Router Functions
2. Host-Host Communication Functions
3. Redirect Function

NDP Functional Groups

Host-Router Discovery Functions

- **Router Discovery**
 - Core function of this group: the method by which hosts locate routers on their local network.
- **Prefix Discovery**
 - Closely related to the process of router discovery is prefix discovery.
 - Determines what network they are on, which tells them how to differentiate between local and distant destinations and whether to attempt direct or indirect delivery of datagrams.
- **Parameter Discovery**
 - A host learns important parameters about the local network and/or routers, such as the MTU of the local link.
- **Address Autoconfiguration**
 - Hosts in IPv6 are designed to be able to **automatically configure themselves**, but this requires information that is normally provided by a router.

Host-Host communications

- **Address Resolution**
 - The process by which a device determines the layer two address of another device on the local network from that device's layer three (IP) address.
 - Performed by ARP in IP version 4.
- **Next-Hop Determination**
 - Looking at an IP datagram's destination address and determining where it should next be sent.
- **Neighbor Unreachability Detection**
 - Determining whether or not a neighbor device can be directly contacted.
- **Duplicate Address Detection (DAD)**
 - Determining if an address that a device wishes to use already exists on the network.

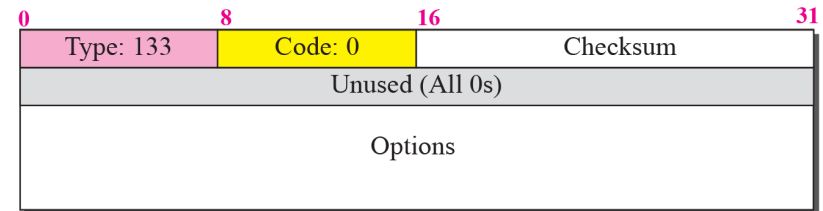
Redirect Function

- The technique whereby a router informs a host of a better next-hop node to use for a particular destination.

ICMPv6 Router Solicitation/Advertisement

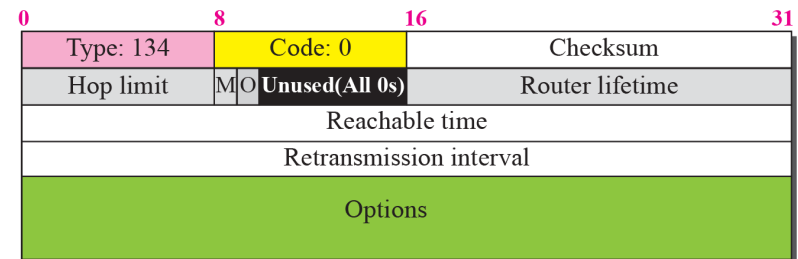
Router Solicitation (ICMPv6 Type 133)

Sent by hosts to request that any local routers send a *Router Advertisement* message so they don't have to wait for the next regular advertisement message.



Router Advertisement (ICMPv6 Type 134)

Sent regularly by routers to tell hosts that they exist and to provide them with important prefix and parameter Information.



Sent on periodic basis from router to the 'all nodes address'

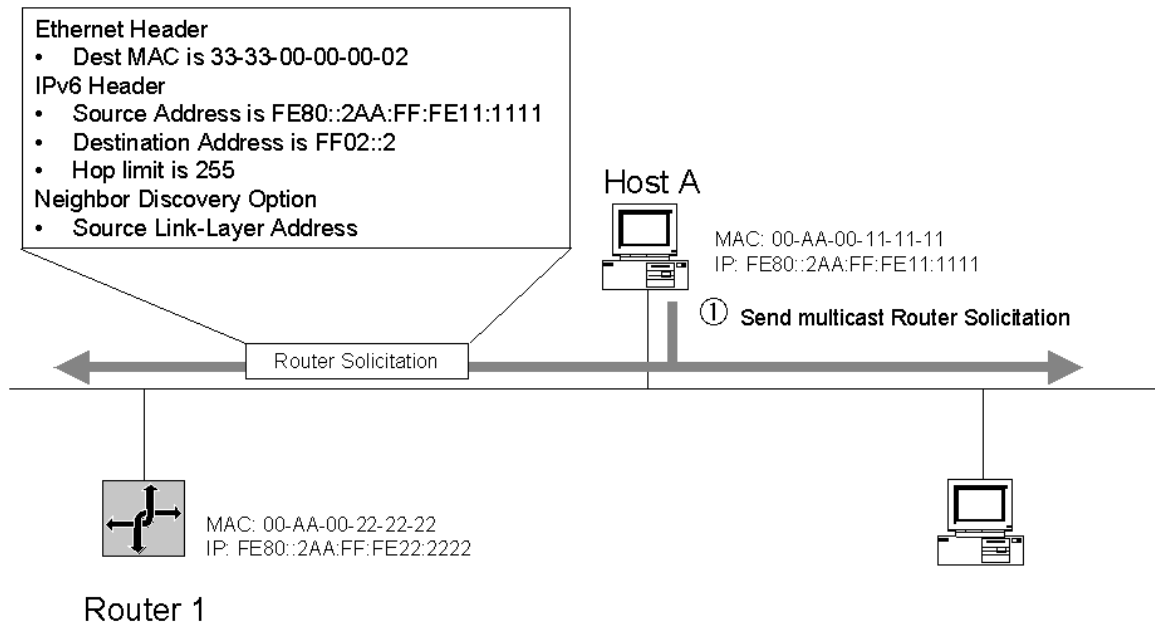
Hop limit should be 255

Could include security header

M=1 use DHCP for address configuration

O=1 use stateful protocol for address configuration

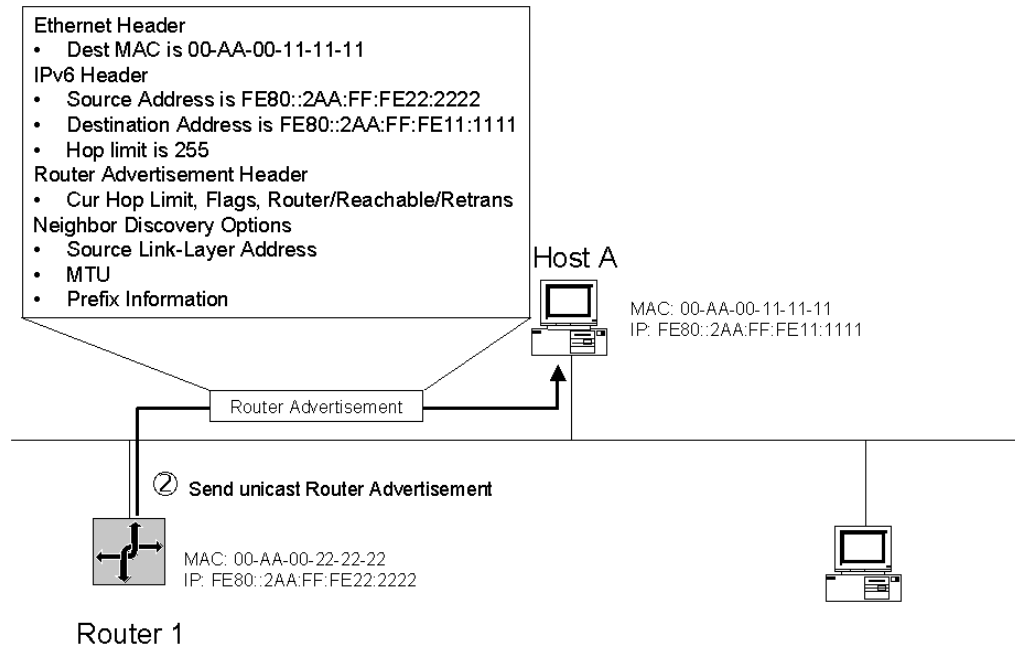
IPv6 Router Discovery



To forward packets to off-link destinations,
Host A must discover the presence of Router 1.

Host A sends a multicast Router Solicitation to the address FF02::2

Router Discovery Response

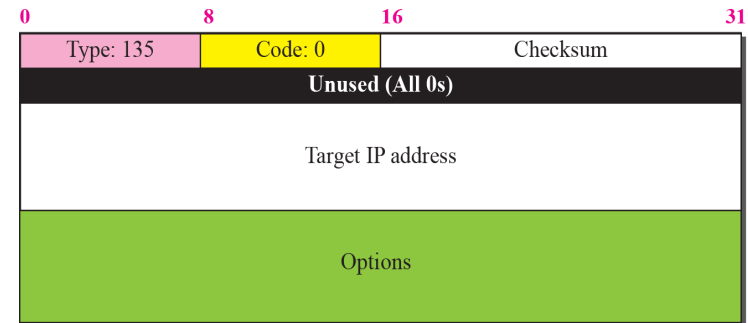


Router 1, having registered the multicast address of 33-33-00-00-00-02 with its Ethernet adapter, receives and processes the Router Solicitation. Router 1 responds with a unicast Router Advertisement message containing configuration parameters and local link prefixes

ICMPv6 Neighbor Messages

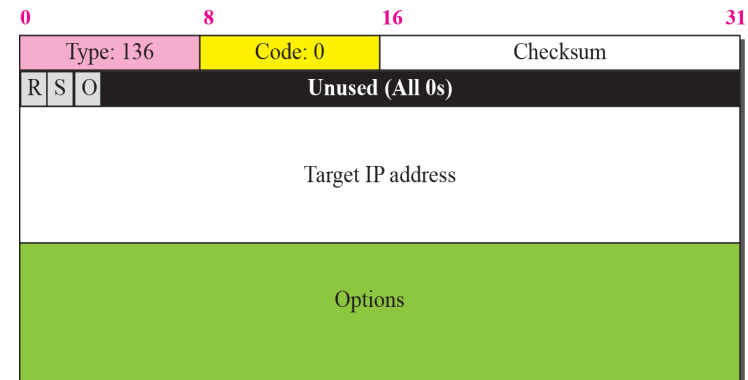
Neighbor Solicitation (ICMPv6 Type 135)

- Nodes ask for link layer address of a target while providing their own link layer address to the target.
- Multicast to resolve an address in the range FF02:::001:FF00:000 to FF02:::001:FFF:FFF
- Take low order 32 bits of address and append to the following prefix: FF02:::001.
- Unicast to verify the reachability of a neighbor.




Neighbor Advertisement (ICMPv6 Type 136)

- Sent by nodes in response to Neighbor solicitation message.
- Can be sent unsolicited to quickly ask for information
- Identify sender as router, destination address, or over-ride existing cache

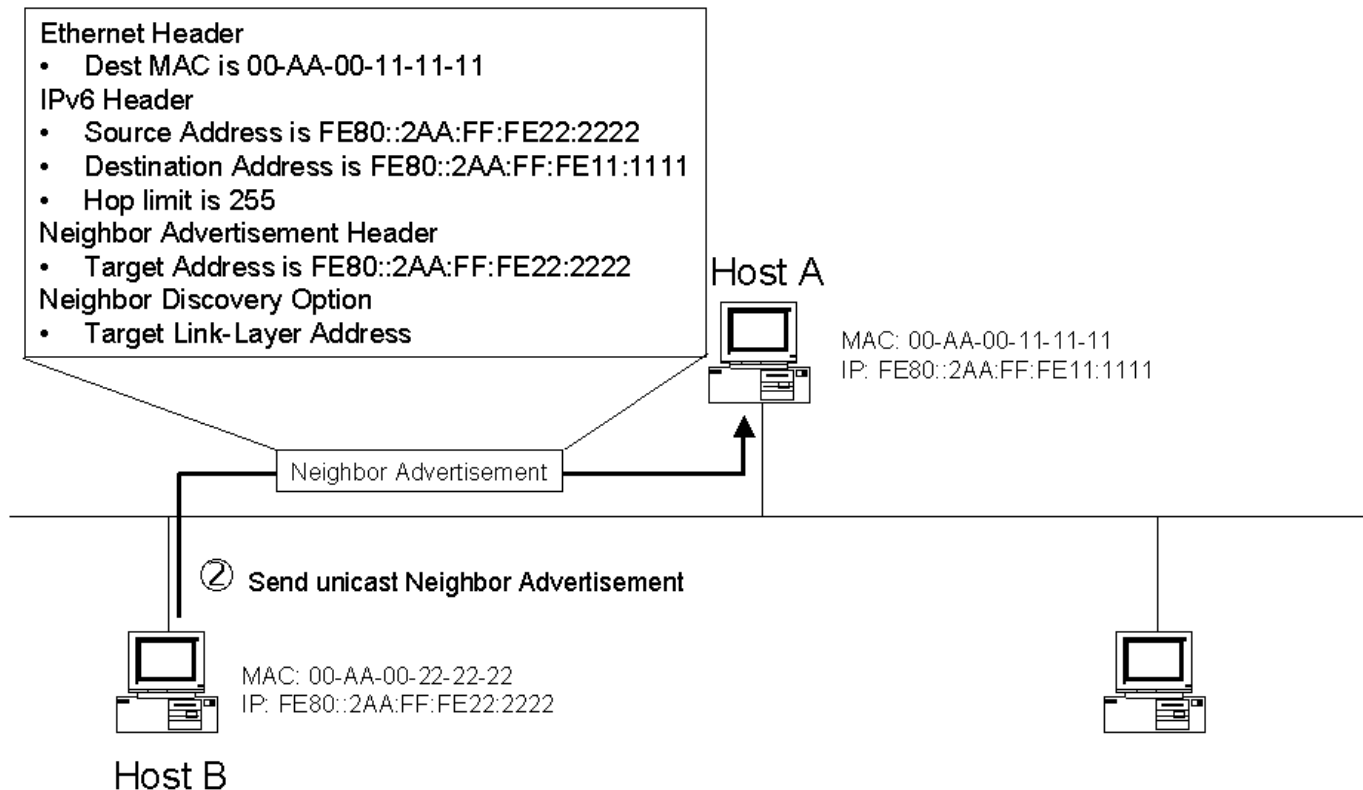


Address Resolution: Multicast Neighbor Solicitation



To send a packet to Host B, Host A must use address resolution to resolve Host B's link-layer address.

Address Resolution: Unicast Neighbor Notification



Host B, having registered the solicited-node multicast address of 33-33-FF-22-22-22 with its Ethernet adapter, receives and processes the Neighbor Solicitation. Host B responds with a unicast Neighbor Advertisement message

Neighbor Solicitation and Advertisement

Traces Query Builder Packet Summary Session Summary						
Packet Summary						
ID	Timestamp	Datagram Size	Local IP	Rmt. IP	Protocol	Messages
320	06:14:34:0405	72	2001:428:3804:0	FF02::1:FF00:1	ICMPv6	
321	06:14:34:0460	161	10.2.0.236	239.255.255.250	UDP	
322	06:14:34:0596	72	FE80::1	2001:428:3804:0	ICMPv6	

Traces | Query Builder | Packet Summary | Session Summary | **Packet Details**

Packet Details [Hex](#)

Packet ID : 320
Time : 4/10/2012 06:14:34:0405 HAT

IP Version 6
Source : 2001:428:3804:0:D78:D8B8:F88D:8A5A
Destination : FF02::1:FF00:1
Traffic Class : 0x000
Flow Label : 0x000
Payload Length : 32
Next Header(Protocol) : ICMPv6
Hop Limit : 255

ICMPv6 Informational Message:
Type: Neighbor Solicitation (135)
Code: 0
Checksum: 0xEE6B
Target Address: FE80::1

ICMPv6 Option
Type: Source Link_layer Address(1)
Length: 8 bytes
Link-layer address: EC:55:F9:C1:E1:51

Packet Details [Hex](#)

Packet ID : 322
Time : 4/10/2012 06:14:34:0596 HAT

IP Version 6
Source : FE80::1
Destination : 2001:428:3804:0:D78:D8B8:F88D:8A5A
Traffic Class : 0x000
Flow Label : 0x000
Payload Length : 32
Next Header(Protocol) : ICMPv6
Hop Limit : 255

ICMPv6 Informational Message:
Type: Neighbor Advertisement (136)
Code: 0
Checksum: 0xD8D5
Flags:
1... = Router: Set
.1.. = Solicited: Set
..1. = Override: Set
Target Address: FE80::1

ICMPv6 Option
Type: Target Link_layer Address(2)
Length: 8 bytes
Link-layer address: 00:08:E2:60:18:1A

Neighbor Discovery Table

```
RouterA#show ipv6 neighbors
```

IPv6 Address	Age	Link-layer Addr	State	Interface
FEC0::1:200:86FF:FE4B:F9CE	0	0000.864b.f9ce	REACH	FastEthernet0/0

```
<waiting of 10 minutes>
```

```
RouterA#show ipv6 neighbors
```

IPv6 Address	Age	Link-layer Addr	State	Interface
FEC0::1:200:86FF:FE4B:F9CE	2	0000.864b.f9ce	STALE	FastEthernet0/0
FE80::200:86FF:FE4B:F9CE	10	0000.864b.f9ce	STALE	FastEthernet0/0

Adding a Static Entry in the Neighbor Discovery Table (Cisco Feature)

```
RouterA(config)#ipv6 unicast-routing
```


```
RouterA(config)#ipv6 neighbor fec0::1:0:0:1:b fastEthernet 0/0 0080.12ff.6633
```

```
RouterA(config)#exit
```

```
RouterA#show ipv6 neighbors
```

IPv6 Address	Age	Link-layer Addr	State	Interface
FEC0::1:200:86FF:FE4B:F9CE	15	0000.864b.f9ce	STALE	FastEthernet0/0
FEC0::1:0:0:1:B	-	0080.12ff.6633	REACH	FastEthernet0/0
FE80::200:86FF:FE4B:F9CE	15	0000.864b.f9ce	STALE	FastEthernet0/0

IPv6 Auto-configuration



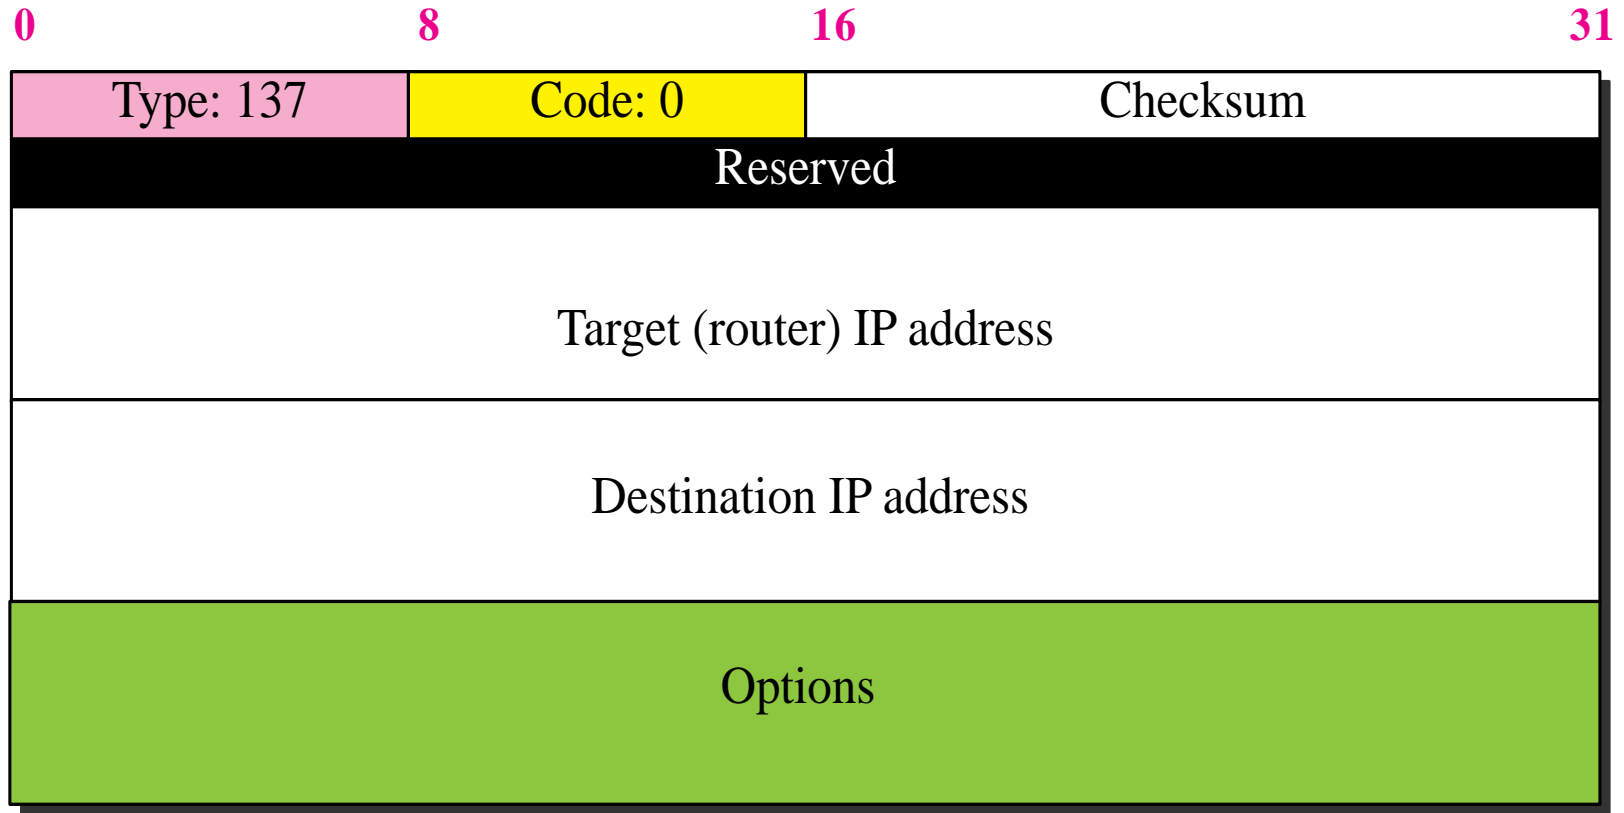
- Host 1 comes on line and generates a link local address.
- Host 1 sends out a query called neighbor discovery to the same address to verify uniqueness. If there is a positive response, a random number generator is used to generate a new address.
- Host 1 multicasts a router solicitation message to all routers.
- Routers respond with a router advertisement that contains the IPv6 Address prefix and other information.
- Host 1 automatically configures its global address by appending its interface ID to the AGA
- Host 1 can now communicate

Prefix Advertisement

Packet Summary											
ID	Timestamp	Datagram Size	Local IP	Rmt. IP	Protocol	Messages	Local Port	Rmt. Port	Seq. Number	Ack. Number	Window Size
132	06:13:39:2874	104	FE80::1	FF02::1	ICMPv6						

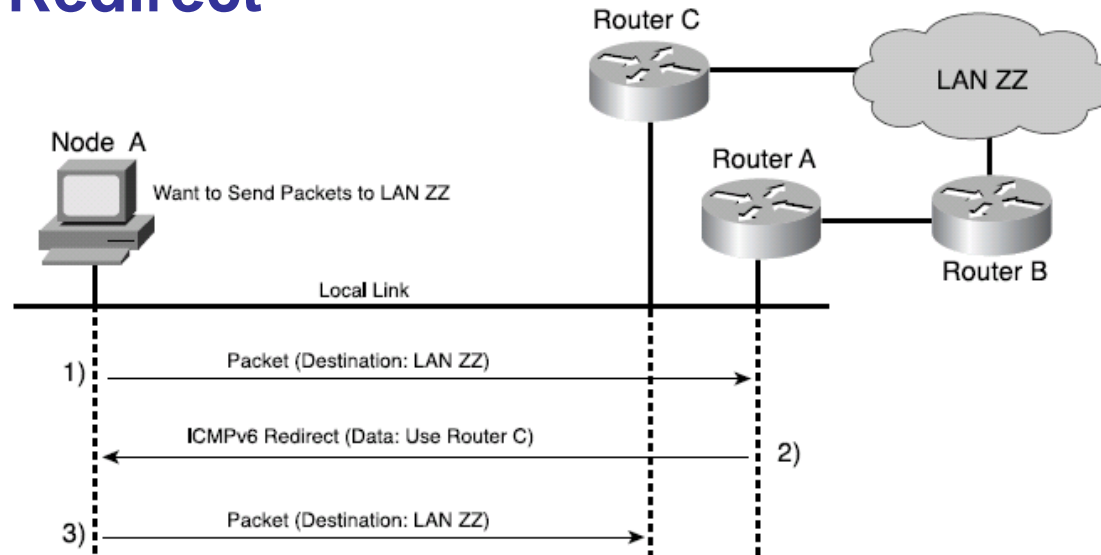
Packet Details Packet Details Hex Packet Details	<pre> Packet ID : 132 Time : 4/10/2012 06:13:39:2874 HAT IP Version 6 Source : FE80::1 Destination : FF02::1 Traffic Class : 0x000 Flow Label : 0x000 Payload Length : 64 Next Header(Protocol) : ICMPv6 Hop Limit : 255 ICMPv6 Informational Message: Type: Router Advertisement (134) Code: 0 Checksum: 0xC673 Cur hop limit: 64 Flags: 1... = Managed address configuration: Set .0.. = Other configuration: Not Set ..0. = Home Agent: Not Set ...0 0... = Default Router Preference: Medium 0.. = Proxy: Not Set Router lifetime (s): 1800 Reachable time (ms): 0 Retrans timer (ms): 0 </pre>	<pre> ICMPv6 Option Type: Source Link_layer Address(1) Length: 8 bytes Link-layer address: 00:08:E2:60:18:1A ICMPv6 Option Type: MTU(5) Length: 8 bytes MTU: 1500 ICMPv6 Option Type: Prefix Information(3) Length: 32 bytes Prefix Length: 64 Flags: 1... = On-link flag(L): Set .1.. = Autonomous address-configuration flag(A): Set Valid Lifetime: 2592000 Preferred Lifetime: 604800 Prefix(IPv6 address): 2001:428:3804:: </pre>
---	---	--

ICMPv6 Redirect



An option is added to let the host know the target router's physical address.

Router Redirect



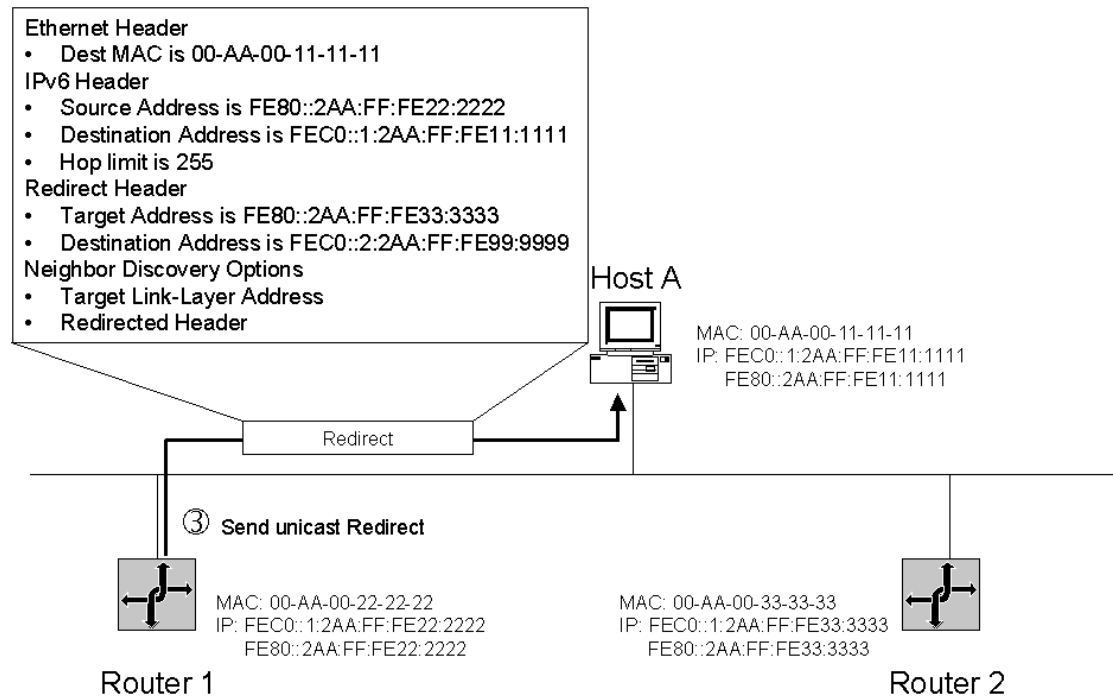
1. **A router informs an originating host of the IP address of a router available on the local link that is “closer” to the destination.**

“Closer” is 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 best 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

Router Redirect Process



To inform Host A that subsequent packets to the destination of FEC0::2:2AA:EE:FE99:9999 should be sent to Router 2, Router 1 sends a Redirect message to Host A.

ICMPv6 Multicast Listener (MLD)

Took pieces from IGMP (Internet Group Management Protocol) (RFC 1112 and RFC 2236) and merged into new protocol.

Defined in RFC 2710.

MLD is a sub-protocol of ICMPv6.

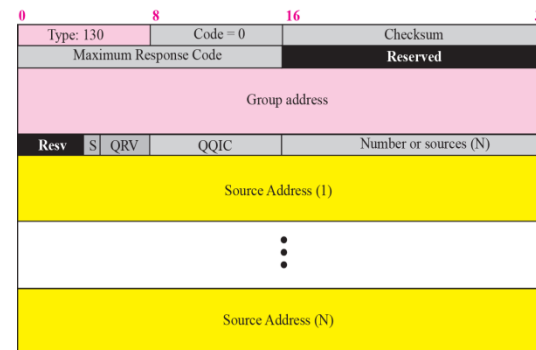
Allows routers to discover nodes that wish to receive multicast packets on all the routers links.

Query can be general or specific:

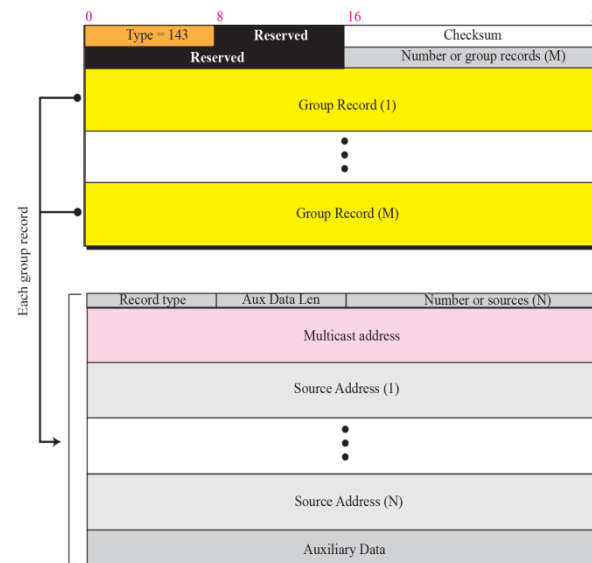
- Tell me all nodes with multicast address x
- Tell me all nodes and their multicast addresses

Maximum response delay only is used with the Query message.

Membership Query



Membership Report



Trace Multicast Listener Query

Packet Summary			
ID	Timestamp	Datagram Size	Local Address
380	06:14:42.4013	72	FE80::...
381	06:14:42.5287	78	10.2.2.1
382	06:14:42.7010	78	FE80::...

Traces	Query Builder	Packet Summary	Session Summary	Packet Details
Packet Details Packet Details Hex EBCDIC ASCII Hex Decode				
<pre> Packet ID : 380 CTRACE Header L 0 E-ID Time 1 CI Ld LINK/JOB Time 2 08 01 0000 C654D068 050000 04 DDDCCCD4 CC444444 0000DDF 0E 00 0004 96B545E7 040050 08 36721320 66000000 000020D LOPEACK FF IPv6 Header V Flow PL N H Source Dest 6 000 02 0 0 F8000000DDFF222 F000000000000000 0 000 00 0 1 E000000020DFE0C0 F2000000000000001 ICMPv6 Header T C CS 3 0 00 A 0 52 RU Data 00008050210000000000000000000000 001020E6700000000000000000000000^'..... </pre>				

Type – 3A (ICMPv6)
 Code – 00
 Checksum -0502

82=130decimal=MLQ
 Maximum Response Delay=
 27 10 hex= 10000ms

Multicast Listener Report

Packet Summary

ID	Timestamp	Datagram Size	Local IP	Rmt. IP
380	06:14:42:4013	72	FE80::D2D0:FDF	FF02::1
381	06:14:42:5287	78	10.2.0.123	10.2.255.255
382	06:14:42:7249	72	FE80::E488:BE1	FF02::1:3

83=131decimal=MLR
 Maximum Response Delay=
 00 00hex= 0ms
 Multicast Address FF02::1:3

Traces | Query Builder | **Packet Summary** | Session Summary | Packer | Window Size

Packet Details

[Packet Details](#) | [Hex](#) | ☐ EBCDIC | ☒ ASCII

[Hex Decode](#)

Packet ID : 382

CTTRACE Header

L	0	E-ID	Time	1	CI	Ld	LINK/JOB	Ti
08	01	0000	C6542058	000010	04	CCECDF44	CC444444	00
0E	00	0004	96B632E7	0A0050	08	39236100	66000000	00
						CISCO1	FF	

IPv6 Header

V	Flow	PL	N	H	Source	Dest
6	000	02	0	0	F8000000E8B1112C	F000000000000000
0	000	00	0	1	E000000048E60F02	F2000000000000103

ICMPv6 Header

T	C	CS
3	0	00
A	0	52

RU Data

```

000080A90000F000000000000000
001030CD0000F2000000000000103
.....
  
```

ICMPv6 Path MTU Discovery

RFC 1981

To enable hosts to discover the min. MTU on a path to a particular destination.

Fragmentation in IPv6 is not performed by intermediary routers.

The source node may fragment packets by itself only when the path MTU is smaller than the packets to deliver.

PMTUD for IPv6 uses ICMPv6 error message

- Type 2 Packet Too Big

MTU Size Error Feedback

- If a router is forced to try sending a datagram that is too large over a physical link, it must drop the datagrams, since it cannot fragment them.
- A feedback process has been defined using ICMPv6 that lets routers tell source devices when the datagrams they are using are too large for the route.

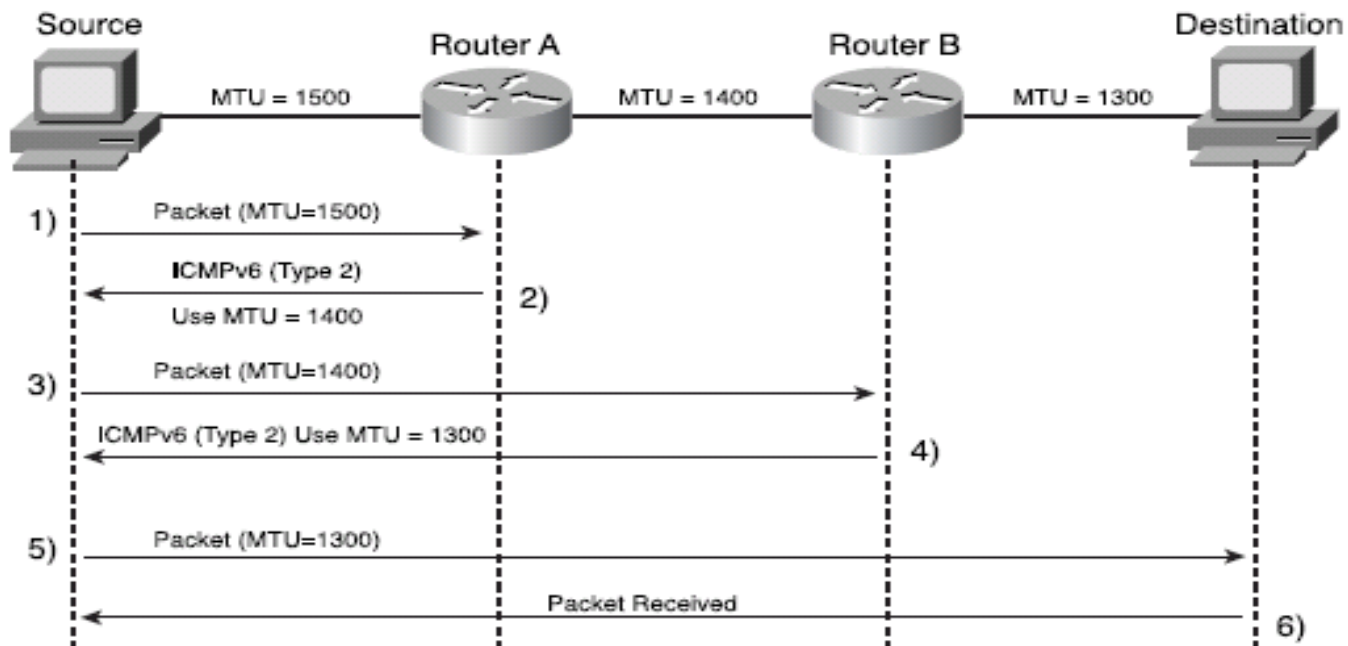
How Does a Node know what MTU size to Use?

1. Use Default MTU

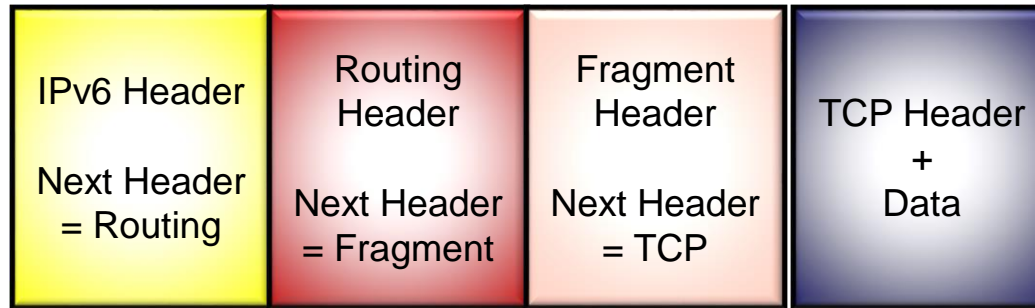
Use the default MTU of **1280**, which all physical networks must be able to handle.

2. Use Path MTU Discovery feature

A node sends messages over a route to determine the overall minimum MTU.



Fragmentation



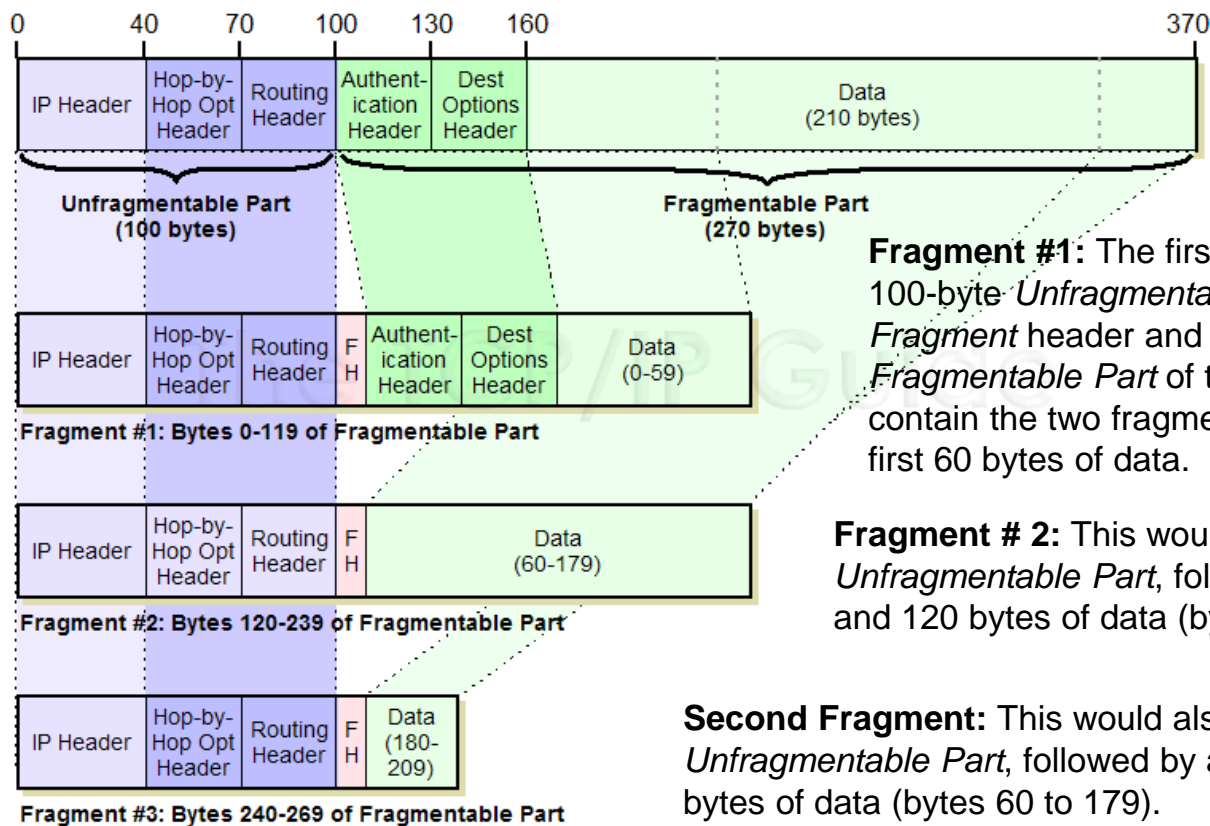
For purposes of fragmentation, IPv6 datagrams are broken into two pieces:

- **Unfragmentable Part**
Includes the main header of the original datagram + any extension headers that need to be present in each fragment - ***Hop-By-Hop Options***, ***Destination Options*** (for those options to be processed by devices along a route) and ***Routing***.
- **Fragmentable Part**
Data portion of the datagram + other extension headers if present - *authentication Header*, *Encapsulating Security Payload* and/or *Destination Options* (for options to be processed only by the final destination).

The ***Unfragmentable Part*** must be present in each fragment, while the ***Fragmentable Part*** is split up amongst the fragments.

Fragmentation Example

Suppose we need to send this over a link with an MTU of only 230 bytes. Three fragments are created. This is due to the need to put the two 30-byte unfragmentable extension headers in each fragment and the requirement that each fragment be a length that is a multiple of 8.




Fragment #1: The first fragment would consist of the 100-byte *Unfragmentable Part*, followed by an 8-byte *Fragment* header and the first 120 bytes of the *Fragmentable Part* of the original datagram. This would contain the two fragmentable extension headers and the first 60 bytes of data.

Fragment # 2: This would also contain the 100-byte *Unfragmentable Part*, followed by a *Fragment* header and 120 bytes of data (bytes 60 to 179).²

Second Fragment: This would also contain the 100-byte *Unfragmentable Part*, followed by a *Fragment* header and 120 bytes of data (bytes 60 to 179).

ICMPv6 Model Host



Each host is to maintain the following:

- Neighbor Cache
- Destination Cache
- Prefix List
- Default Router List
- LinkMTU
- CurHopLimit
- BaseReachable Time
- Reachable Time
- Retransmit Timer

Changes Needed to Implement IPv6

Hosts

- ✓ Implement IPv6 code in operating system
- ✓ TCP/UDP aware of IPv6
- ✓ Sockets/Winsock library updates for IPv6
- ✓ Domain Name Server updates for IPv6

Domain Name Server (DNS)


- ✓ Many products already support 128 bit addresses
- ✓ Uses 'AAAA' records for IPv6
- ✓ IP6.INT (in_addr_arpa in IPv4)

Routers

- ✓ IPv6 forwarding protocols
- ✓ Routing protocols updated to support IPv6
- ✓ Management needs to support ICMPv6
- ✓ Implement transition mechanisms

IPv6 Protocol Status

- ✓ RIPv6 - Same as RIPv2
- ✓ OSPFv6 - Updated for IPv6
- ✓ EIGRP - Extensions implemented
- ✓ IDRP - Recommended for exterior protocol over BGP4
- ✓ BGP4+ - Preferred implementation in IPv6 today



Vielen
Dank

ありがとうございました

Köszönettel

Obi Спасибо

ขอบคุณ

شكراً

Bedankt

Gracias

شكراً

धन्यवाद

Ευχαριστώ

THANK YOU

Merci

Díky

Grazie

Danke

Hvala

Merci

ขอบคุณ

Teşekkürler

תודה

धन्यवाद
Hindi

Gracias

laurak@aesclever.com

www.aesclever.com

650-617-2400

감사합니다

நன்றி
Tamil

Obrigado

IPv6 References

IPv6 Home Page

<http://www.ietf.org/>

<http://playground.sun.com/pub/ipng/html/ipng-main.html>

http://www.getipv6.info/index.php/IPv6_Presentations_and_Documents<http://www.6ren.net>

<http://www.ipv6forum.com>

<http://arin.net>

<http://www.internet2.edu>

<http://www.ipv6.org>

<http://ipv6.or.kr/english/natpt.overview>

<http://www.research.microsoft.com/msripv6>

<http://www.ipv6.org.uk>

Books

New Internet Protocol - Prentice Hall - ISBN 0-13-241936-x

IPNG and the TCP/IP Protocols - John Wiley and Sons - ISBN-0-471-13088-5

IPv6 The New Internet Protocol - ISBN-0-13-24-241936

IPNG Internet Protocol Next Generation - ISBN-0-201-63395-7

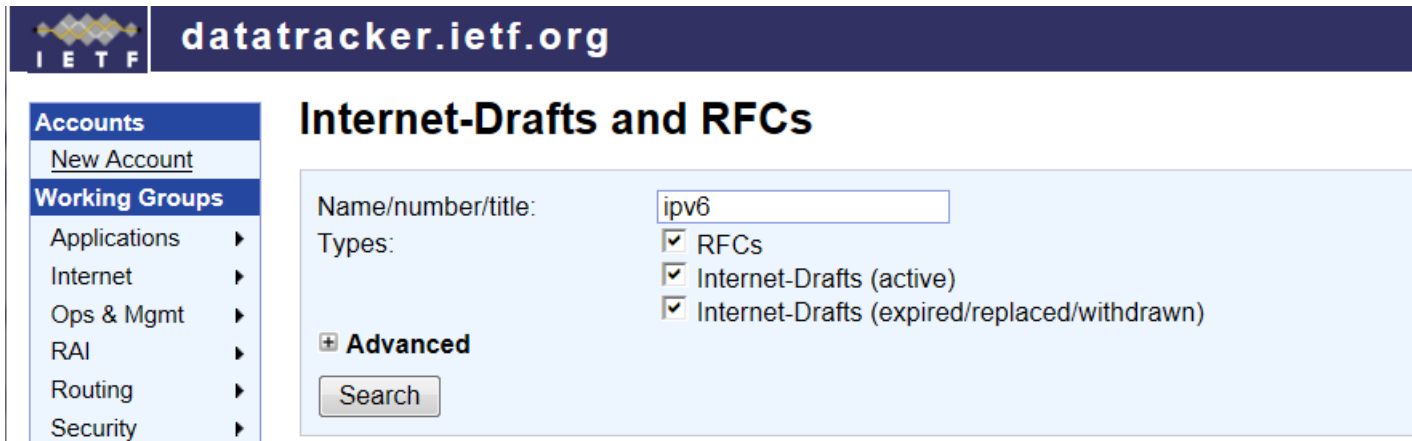
Internetworking IPv6 with Cisco Routers - ISBN 0-07-022831-1



IPv6 RFCs

View any IPv6 RFC

<http://datatracker.ietf.org/doc/search/>



The screenshot shows the datatracker.ietf.org website. On the left is a navigation menu with 'Accounts' (containing 'New Account') and 'Working Groups' (listing Applications, Internet, Ops & Mgmt, RAI, Routing, and Security). The main content area is titled 'Internet-Drafts and RFCs'. It features a search form with a text input field containing 'ipv6', a 'Types:' section with three checked checkboxes for 'RFCs', 'Internet-Drafts (active)', and 'Internet-Drafts (expired/replaced/withdrawn)', an 'Advanced' link, and a 'Search' button.

datatracker.ietf.org

Internet-Drafts and RFCs

Name/number/title:

Types:

- ☒ RFCs
- ☒ Internet-Drafts (active)
- ☒ Internet-Drafts (expired/replaced/withdrawn)

[Advanced](#)