Internet Control Message Protocol

ICMP

1

The Context

- IP-Related Protocol: ICMP

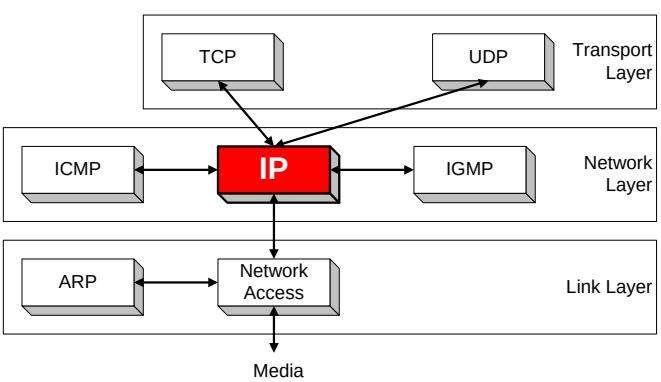

TCP/IP Protocol Stack

2

Internet Control Message Protocol (ICMP)

- The Internet Control Message Protocol (ICMP) is used by routers and hosts to send network control information to each other
- From a layering point of view, ICMP is a separate protocol that sits above IP and uses IP to transport messages
- In practice, ICMP is an integral part of IP and all IP modules must support the ICMP protocol

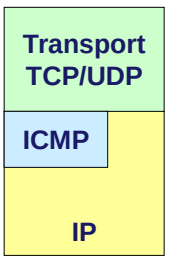

3

Internet Control Message Protocol

- ICMP messages are encapsulated in IP datagrams

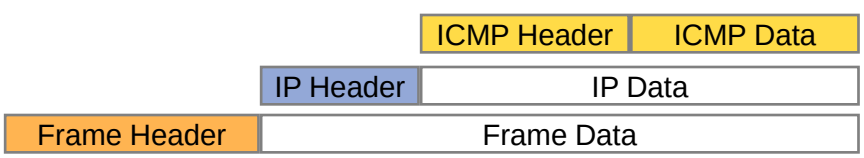

- ICMP frames are identified by IP Protocol field value 1
- Used by IP to send error and control messages
- Uses IP to send its messages

4

ICMP Message Format

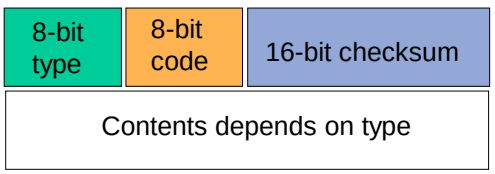

- ICMP message: content contains the first 8 bytes of IP datagram causing error, plus other things

5

Types of ICMP Messages

- Information messages
  - Sender sends a query to another machine (either host or router) and expects an answer. For example, a host might want to know if a router is alive
- Error indication messages
  - The IP software on a host or router has encountered a problem processing an IP datagram. For example, it may be unable to route the datagram to its destination.

6

ICMP Types of Messages

Category	*Type*	*Message*
Error-reporting messages	3	Destination unreachable
	4	Source quench
	11	Time exceeded
	12	Parameter problem
	5	Redirection
Query messages	8 or 0	Echo request or reply
	13 or 14	Timestamp request or reply
	17 or 18	Address mask request or reply
	10 or 9	Router solicitation or advertisement

7

ICMP Error Reporting

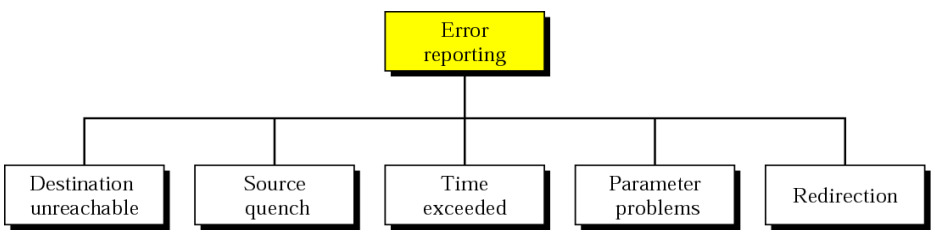

10

ICMP Error Messages

ICMP always reports error messages to the original source.

- Important Points
  - No ICMP error message will be generated in response to a datagram carrying an ICMP error message
  - No ICMP error message will be generated for a fragmented datagram that is not the first fragment
  - No ICMP error message will be generated for a datagram having a multicast address
  - No ICMP error message will be generated for a datagram having a special address such as 127.0.0.0 or 0.0.0.0

11

ICMP Data for Error Messages

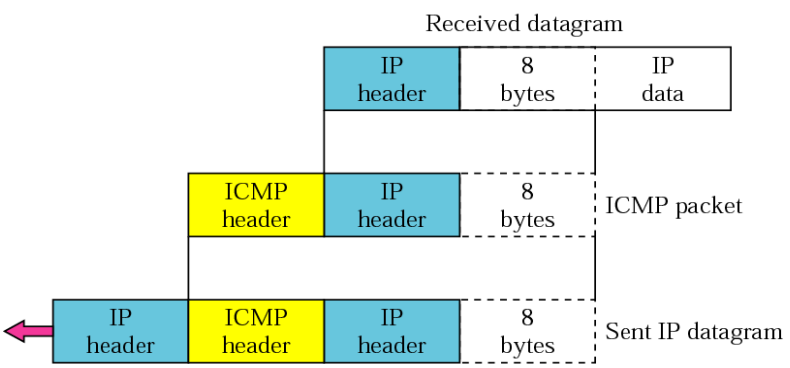

12

ICMP Error Messages (1)

- Destination Unreachable (type 3)
  - When a gateway (router) cannot route a datagram (e.g., it doesn't have an appropriate route in its local table, or it needs to fragment and the DF bit is set), it discards the message and returns an ICMP "destination unreachable" message to the sending host.

- Source Quench (type 4)
  - When a gateway becomes congested and runs out of buffer space, it may discard a datagram and return a source quench message. Source quench messages are used to request that the sender reduce the rate at which it is sending datagrams

13

ICMP Error Messages (2)

- Time Exceeded (type 11)
  - As a datagram is processed, routers decrement its time-to-live (TTL) field. If the TTL value reaches 0, the gateway discards the datagram and sends a time exceeded message (code 0) to the sender.
  - Code 1 is used by a destination host to show that not all fragments have arrived within a set time.

14

Recall: Traceroute

- Traceroute records the route that packets take
- A clever use of ICMP and the TTL field
- When a router receives a packet, it decrements TTL
- If TTL=0, send ICMP "Time exceeded" back to sender
- To determine a route
  - Send a packet with TTL = 1 (hop)
  - The first router discards the packet and sends ICMP "Time Exceeded"; when ICMP "Time Exceeded" is received, record the sender's (router's) address
  - Increment TTL
  - Repeat until the destination host is received or an error occurs

15

Traceroute (contd.)

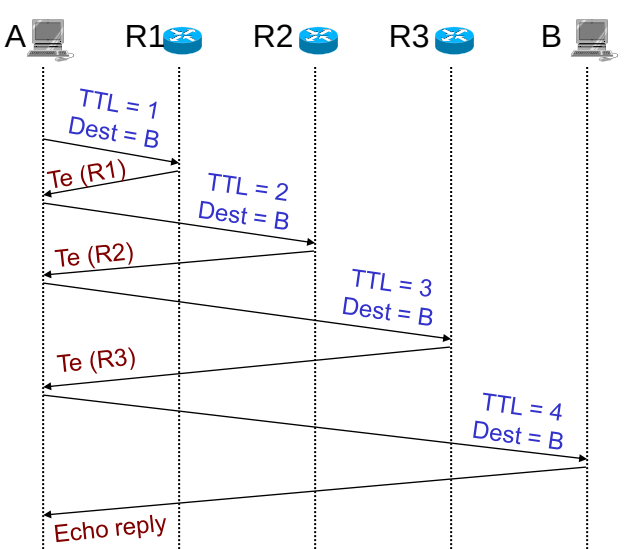

16

ICMP Error Messages (3)

- Parameter Problem (type 12)
  - When a host or gateway encounters a problem parsing an IP datagram, it returns a parameter problem message to the datagram's sender

- Redirection (type 5)
  - Sent from a router to a local host on the same network
  - Informs the source of a better route to the destination
  - A host usually starts with a small routing table that is gradually augmented and updated. Redirection helps it.

17

ICMP Redirect Example (1)

Host A has one default router, which is R2

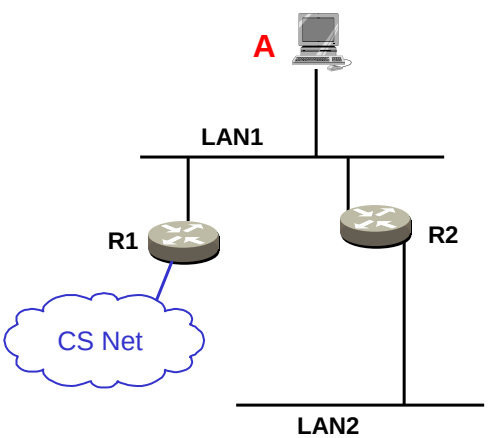

18

ICMP Redirect Example (2)

When A wants to send a message to the Campus Net, it sends it to the default router (R2)

R2 forwards the message to R1

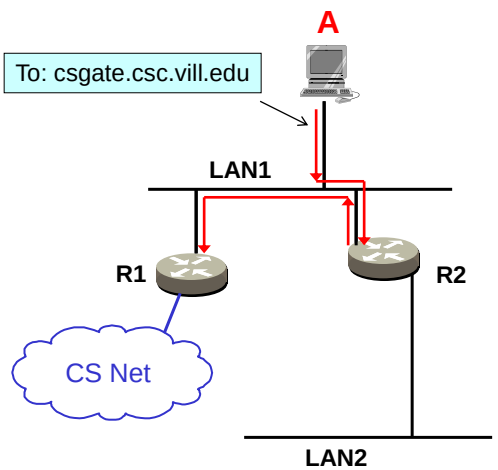

19

ICMP Redirect Example (3)

R2 also sends an ICMP redirect message to A, telling it to use R1 for connections to csgate.vill.edu

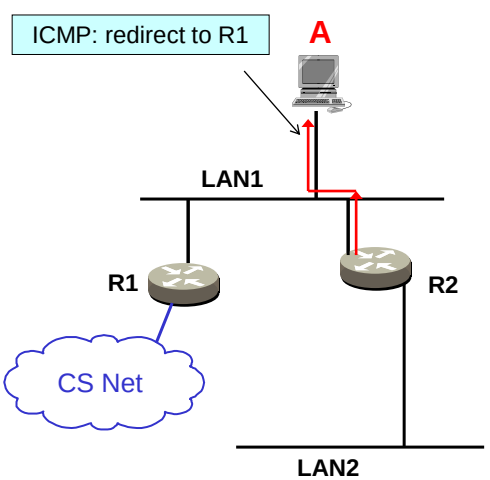

20

ICMP Redirect

Type: 5	Code: 0 to 3	Checksum
IP address of the target router		
Part of the received IP datagram including IP header plus the first 8 bytes of datagram data		

Redirection Message Format

ICMP Redirect Example (4)

A sends subsequent packets directly to R1

Note: some hosts deliberately ignore ICMP Redirect messages as a precaution against network attacks.

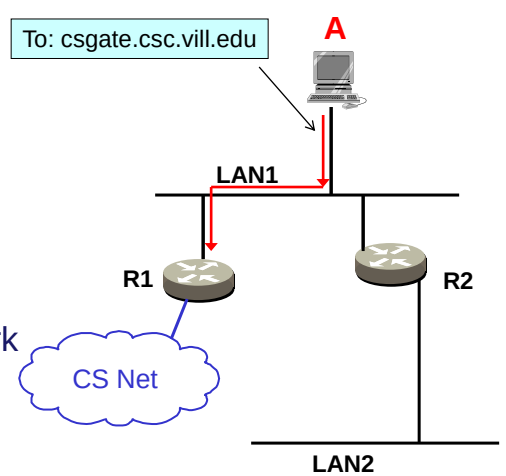

ICMP Queries

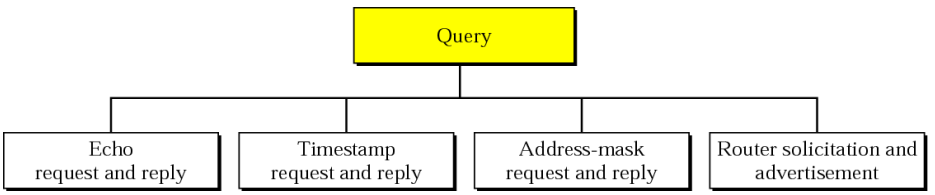

ICMP Query Messages

- Used to diagnose network problems
- In this type of ICMP message, a node sends a message that is answered in a specific format by the destination node.

- Echo Request / Reply (types 8 / 0)
  - If machine A sends an ICMP echo request message to machine B, machine B is required to respond with an ICMP echo reply
  - In UNIX, the program ping allows a user to check whether a machine is reachable and functioning

Ping

- Uses ICMP Echo request/reply to
  - test destination reachability
  - compute round trip time
  - count the # of hops to destination
- Source sends ICMP echo request message to the destination address
  - echo request packet contains timestamp also
- Destination replies with an ICMP echo reply message containing the data in the original request message
- Source can calculate RTT of packets
- If no echo reply comes back, destination unreachable

25

Ping (contd.)

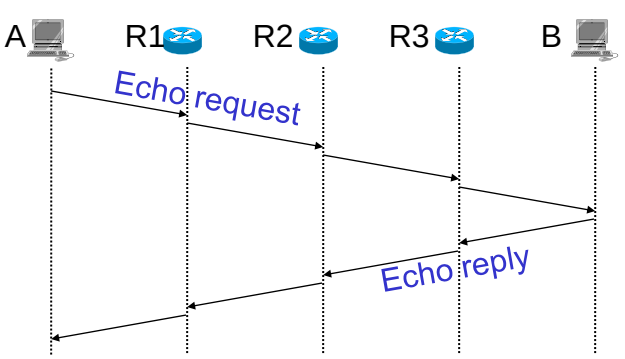

- Sample output:

  Reply from 164.107.144.3: 48 bytes in 47 msec. TTL: 253

26

ICMP Echo Request / Reply Message

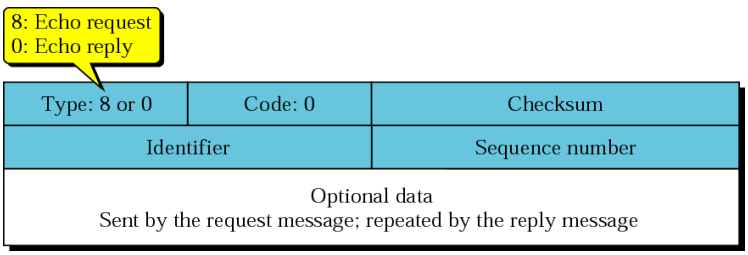

- The identifier and sequence number may be used by the echo sender to aid in matching the replies with the echo requests.

27

ICMP Timestamp Request / Reply Message

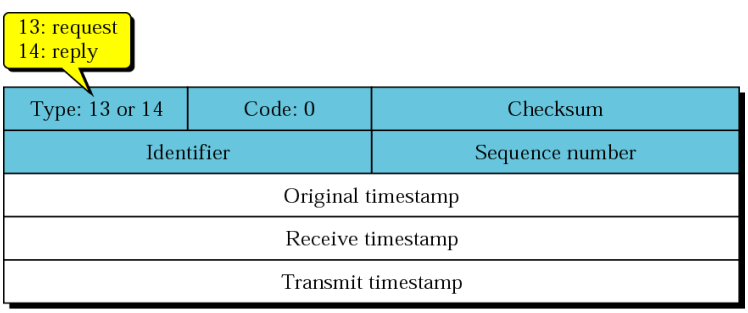

- Timestamp-request and timestamp-reply messages can be used to calculate the round-trip time between a source and a destination machine.

28

ICMP Mask Request / Reply Message

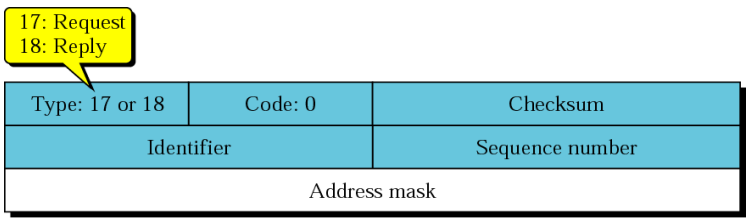

- Timestamp-request and timestamp-reply messages can be used to calculate the round-trip time between a source and a destination machine.

29

ICMP Router Solicitation / Advertisement

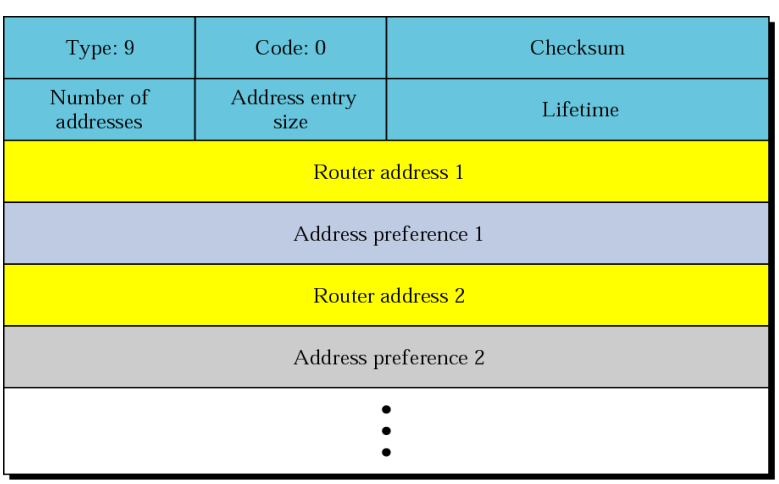

30

New ICMP Version:

ICMPv6

31

Context for ICMPv6

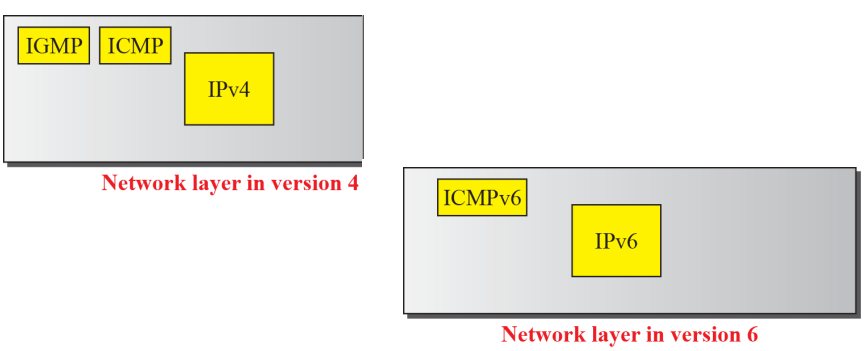

ICMPv6 follows the same strategy and purposes of version 4. It is only slightly more complex than ICMPv4.

32

Internet Control Message Protocol ICMPv6

- ICMPv6 is more complex than ICMPv4:
  - some protocols that were independent in version 4 are now part of ICMPv6
  - new messages have been added to ICMPv6 to make it more useful
- Introduces some simplifications by eliminating obsolete types of messages no longer in use

ICMPv6 Header

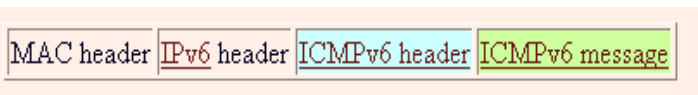

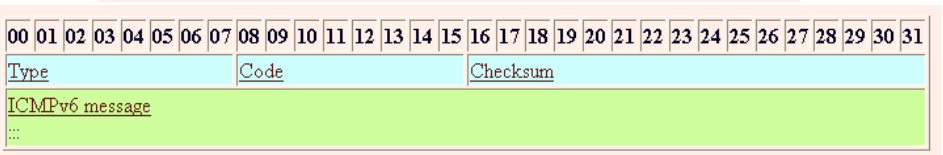

- Type (8 bits)
  - High order bit = 0 (0 – 127) indicates error message
  - High-order bit = 1 (128 – 255) indicates information message.
- Code (8 bits)
  - depends on the message type
- Checksum (16 bits)
  - Used to detect errors in ICMP and part of IPv6

ICMPv6 Messages

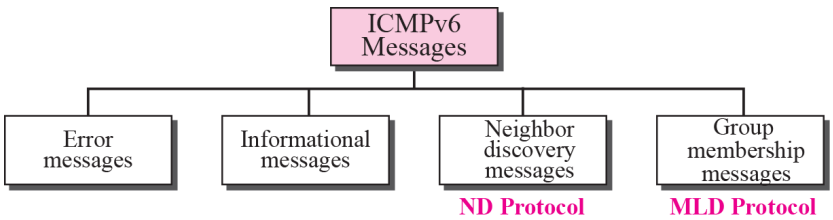

- Transported within an IPv6 packet in which extension headers can also be present.
- Identified by a value of **58 in the Next Header** field of the IPv6 header or of the preceding Header.

35

ICMPv6 Error-Reporting Messages

- Destination-Unreachable
- Packet-Too-Big
- Time-Exceeded
- Parameter-Problem

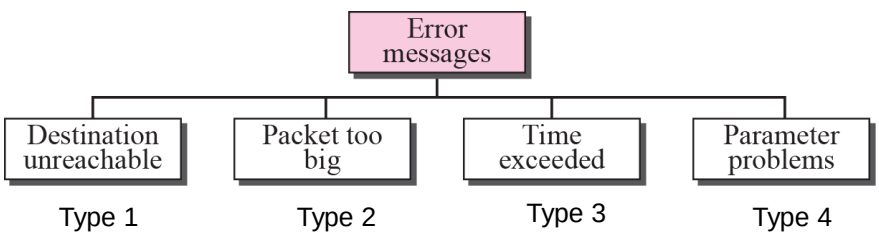

36

ICMPv6 Error Reporting Messages

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

37

Error Reporting in ICMP (v4 vs. v6)

Type of Message	*Version 4*	*Version 6*
Destination unreachable	Yes	Yes
Source quench	Yes	No
Packet too big	No	Yes
Time exceeded	Yes	Yes
Parameter problem	Yes	Yes
Redirection	Yes	Yes

38

ICMPv6 Informational Messages

- Echo-Request
- Echo-Reply
- Router-Solicitation
- Router-Advertisement
- Neighbor-Solicitation
- Neighbor-Advertisement
- Redirect

39

ICMPv6 Echo Request / Reply

- Echo-Request: Sent to test connectivity to another device
- Echo-Reply: Sent in response to Echo request

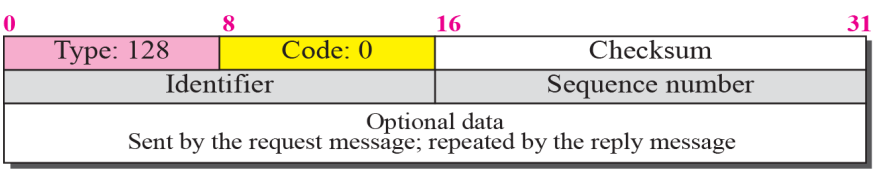

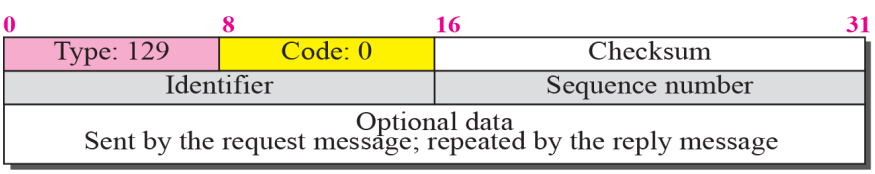

40

Neighbor Discovery Messages

- Used by nodes (hosts or routers) on the same link
  - Router-Solicitation Message
  - Router-Advertisement Message
  - Neighbor-Solicitation Message
  - Neighbor-Advertisement Message

41

ICMPv6 Neighbor Discovery Messages

Type	Message Name	Summary Description of Message Type
133	Router Solicitation	Prompts a router to send a Router Advertisement
134	Router Advertisement	Sent by routers to tell hosts on the local network that the router exists and describe its capabilities
135	Neighbor Solicitation	Sent by a device to request the MAC address of another local device and provide its own
136	Neighbor Advertisement	Provides information about a host to the local network

42

Information Messages in ICMP (v4 vs. v6)

Type of Message	*Version 4*	*Version 6*
Echo request and reply	Yes	Yes
Timestamp request and reply	Yes	No
Address mask request and reply	Yes	No
Router solicitation and advertisement	Yes	Yes
Neighbor solicitation and advertisement	ARP	Yes
Group membership	IGMP	Yes

43

Path MTU Discovery for ICMPv6 (PMTUD)

PMTUDv6 Overview

- 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

45

IPv4 vs. IPV6 MTU

- Increased Default MTU
  - In IPv4 minimum required MTU = **576 bytes**.
  - In IPv6 minimum required MTU = **1280 bytes**.
  - Improves efficiency by increasing the ratio of maximum payload to header length, and reduces the frequency of fragmentation

- Elimination of En Route Fragmentation
  - In IPv4, datagrams may be fragmented by either the source device, or by routers during delivery.
  - In IPv6, only the source node can fragment; **routers do not**.
  - The source must therefore fragment to the size of the smallest MTU on the route before transmission.

46

How Do Hosts Know What Size to Use?

Two choices:

1. Use Default MTU

- Use the default MTU of 1280, which all physical networks must be able to handle.
- Good choice especially for short communications or for sending small amounts of data.

2. Use Path MTU Discovery feature

- A node sends messages over a route to determine what the overall minimum MTU for the path is

47

Path MTU Discovery

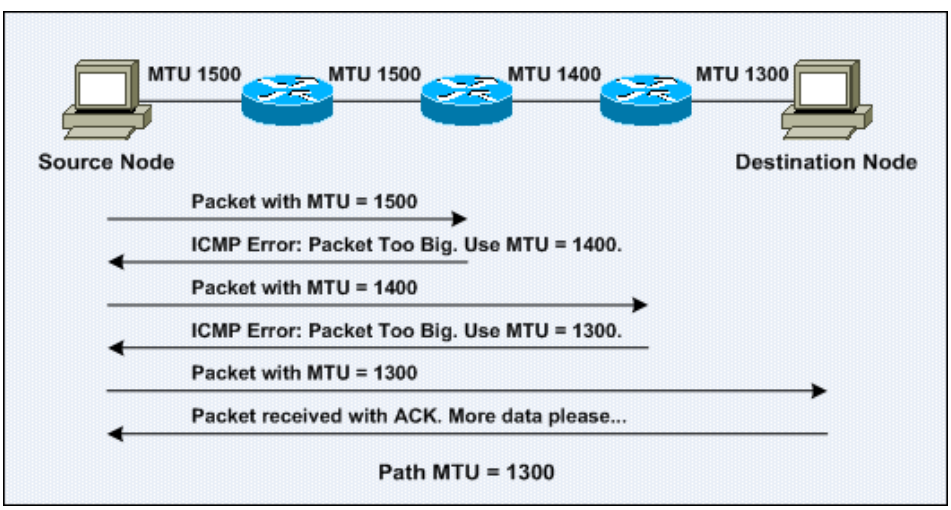

48