

Internet Protocol: ICMP Messages

Linda Wu

(CMPT 471 • 2003-3)

Content

- ICMP protocol
- ICMP messages
- Different types of messages

Reference: chapter 9

ICMP Protocol

- Internet Control Message Protocol
- A required part of IP
- A message mechanism for error reporting and host / management queries
- Reports errors, but does not specify the actions for error correction
- Error is reported to the original source of the datagram, not to intermediate routers


ICMP Messages

- Can be sent by both routers and hosts
- Handled by IP software on the destination machine
- May be lost or discarded
- May cause additional congestion in already congested network
- Contained in the data portion of IP datagram (PROTOCOL = 1)

ICMP Messages (cont.)

ICMP message encapsulation

- ICMP is not a higher level protocol on top of IP
- ICMP message is delivered by IP because it may travel across several physical networks



Notes-7


CMPT 471 • 2003-3

5

ICMP Messages (cont.)

ICMP message general format

- Header: 8 bytes
 - TYPE: the type of the message
 - CODE: the reason for the particular message type
 - CHECKSUM: applies to ICMP message only
 - Rest of the header: specific for each message type
- Data: variable length
 - Error message: include the header and first 64 bits of the datagram causing the problem
 - Query message: extra information based on query type




Notes-7

6

ICMP Messages (cont.)

ICMP error message content



Notes-7

CMPT 471 • 2003-3

7

ICMP Messages (cont.)

ICMP message types

- Error-reporting: report problems that a router or host encounters when processing datagram
- Query: get specific information from a router or host

Category	Type	Message
Error-reporting messages	3	Destination unreachable
	4	Source quench
	5	Redirect (change a route)
	11	Time exceeded for a datagram
	12	Parameter problem on a datagram
Query messages (in pairs)	8 / 0	Echo request / reply
	13 / 14	Timestamp request / reply
	17 / 18	Address mask request / reply
	10 / 9	Router solicitation / advertisement

Notes-7

CMPT 471 • 2003-3

8

ICMP Messages (cont.)

- NO ICMP message will be generated for:
 - The error resulted from the datagram carrying an ICMP error message
 - A fragmented datagram that is not the first fragment
 - A datagram having a multicast address
 - A datagram having a special address such as 127.0.0.0 or 0.0.0.0

Message: Destination Unreachable

- When a router cannot route a datagram or a host cannot deliver a datagram
 - Send *destination unreachable* message back to the original source
 - Discard the datagram
- Possible reasons for unreachable destination
 - Hardware out of service
 - Destination address not exist
 - Router has no route to the destination network
 -
- A router cannot detect all errors that prevent the delivery of datagram

Message: Destination Unreachable (cont.)

◦ Message format

0	8	16	31
TYPE (3)	CODE (0-12)	CHECKSUM	
Unused (must be 0)			
IP header + first 64 bits of datagram			
.....			

- Messages with code 2 or 3: created only by destination host
- Other messages: created only by routers

Message: Destination Unreachable (cont.)

CODE	Meaning
0	Network unreachable
1	Host unreachable
2	Protocol unreachable
3	Port unreachable
4	Fragment needed and DF set
5	Source route failed
6	Destination network unknown
7	Destination host unknown
8	Source host isolated
9	Communication with destination network administratively prohibited
10	Communication with destination host administratively prohibited
11	Network unreachable for type of service
12	Host unreachable for type of service

Message: Source Quench

- For congestion and datagram flow control
- When datagrams arrive too quickly
 - The host or router enqueues the datagram in memory temporarily
 - If memory is exhausted
 - Discard additional datagrams that arrive
 - Send *source quench* message to report congestion to the original source (one message per discarded datagram)
 - Message receiver slows down datagram transmission rate until it stops receiving source quench message
 - Then it gradually increases the rate as long as no further message received (*No mechanism to tell the source the congestion is relieved)

Notes-7

CMPT 471 • 2003-3

13

Message: Source Quench (cont.)

- Source quench message is a request for the source to reduce its current rate of datagram transmission
- Congestion caused by
 - One-to-one communication
 - Source quench message is helpful for this case
 - Many-to-one communication
- Message format

0	8	16	31
TYPE (4)	CODE (0)	CHECKSUM	
Unused (must be 0)			
IP header + first 64 bits of datagram			
.....			

Notes-7

CMPT 471 • 2003-3

14

Message: Redirect

- Routing table of a router
 - Routers exchange routing information periodically to accommodate network changes
 - Routing table of a router is up-to-date
- Routing table of a host
 - Begins with:
 - A limited number of entries
 - one router – default router
 - Hosts do not take part in the dynamic routing update
 - Hosts depend on the routers to update its routing table


Notes-7

CMPT 471 • 2003-3

15

Message: Redirect (cont.)

- Redirect concept
 - Host A sends a datagram destined for B to a wrong router R1 (therefore, a non-optimal route)
 - R1 forwards the datagram to the correct router R2 (* R1 will NOT discard the datagram)
 - R1 sends *redirect* message to A to update A's routing table



Notes-7

CMPT 471 • 2003-3

16

Message: Redirect (cont.)

- Limitation
 - Redirect message is limited between a router and a host on the same network
 - Message format
- | | | | |
|---------------------------------------|--------------|----------|----|
| 0 | 8 | 16 | 31 |
| TYPE (5) | CODE (0 - 3) | CHECKSUM | |
| Correct Router's IP Address | | | |
| IP header + first 64 bits of datagram | | | |
| | | | |
- CODE 0: redirect datagram for the net
1: redirect datagram for the host
2: redirect datagram for the type of service and net
3: redirect datagram for the type of service and host

Notes-7

CMPT 471 • 2003-3

17

Message: Time Exceeded

- Time exceeded message is generated in two cases:
 - A router receives a datagram with $\text{ttl}=0$
 - Datagram is discarded
 - *Time exceeded* message is sent to the original source
 - The destination does not receive all the fragments before the reassembly timer expires
 - All received fragments are discarded
 - *Time exceeded* message is sent to the original source

Notes-7

CMPT 471 • 2003-3

18

Message: Time Exceeded (cont.)

- Message format

0	8	16	31
TYPE (11)	CODE (0 / 1)	CHECKSUM	
Unused (must be 0)			
IP header + first 64 bits of datagram			
.....			

CODE 0: time-to-live count exceeded
CODE 1: fragment reassembly time exceeded

Notes-7

CMPT 471 • 2003-3

19

Message: Parameter Problem

- If there is an ambiguous or missing value in any field of datagram
 - Discard the datagram
 - Send *parameter problem* message to the original source
 - Message format
- Code 0: there is an error or ambiguity in the header fields; POINTER points to the byte with the problem
 - Code 1: a required part of an option is missing; POINTER is not used

0	8	16	31
TYPE (12)	CODE (0 / 1)	CHECKSUM	
POINTER	Unused (must be 0)		
IP header + first 64 bits of datagram			
.....			

Notes-7

CMPT 471 • 2003-3

20

Message: Echo Request / Reply

- To test whether a destination is reachable and responding
 - A host or router sends an ICMP *echo request* message to the destination
 - The destination returns an *echo reply* message to the sender
 - The request contains an optional data area
 - The reply contains a copy of the data sent in the request

Successful receipt of reply → major pieces of the transport system work

Message: Echo Request / Reply (cont.)

- The command to send echo requests – ping
 - Use echo request datagram to elicit an echo reply datagram from another machine
 - Send a series of echo requests, capture responses, and provide statistics about datagram loss
 - Run “ping may” on July

```


PING may (172.16.1.5) from 172.16.1.7 : 56(84) bytes of data.
64 bytes from may (172.16.1.5): icmp_seq=1 ttl=255 time=0.340 ms
64 bytes from may (172.16.1.5): icmp_seq=2 ttl=255 time=0.324 ms
64 bytes from may (172.16.1.5): icmp_seq=3 ttl=255 time=0.325 ms
64 bytes from may (172.16.1.5): icmp_seq=4 ttl=255 time=0.327 ms
64 bytes from may (172.16.1.5): icmp_seq=5 ttl=255 time=0.328 ms
64 bytes from may (172.16.1.5): icmp_seq=6 ttl=255 time=0.330 ms

--- may ping statistics ---
6 packets transmitted, 6 received, 0% loss, time 5003ms
rtt min/avg/max/mdev = 0.324/0.329/0.340/0.005 ms
    
```

Message: Echo Request / Reply (cont.)


- Message format
 - IDENTIFIER & SEQUENCE NUMBER
 - Not formally defined by the protocol
 - Used arbitrarily by the sender to match replies to requests; e.g., identifier is the same as the process id, sequence number keeps track of the echo request messages
 - OPTIONAL DATA
 - Data to be returned to the sender by the destination
 - Variable length

Request: 8
Reply: 0



Message: Timestamp Request / Reply

- Message format
 - Type
 - timestamp request (13)
 - timestamp reply (14)
 - IDENTIFIER & SEQUENCE NUMBER
 - Used by the source to relate replies with requests
 - TIMESTAMP
 - Given in milliseconds since midnight, Universal time



Message: Timestamp Request / Reply (cont.)

- Update the TIMESTAMPS
 - The requesting machine sends a *timestamp request* message to another machine
 - Fill in ORIGINATE TIMESTAMP according to its clock just before the request is sent
 - The receiving machine returns a *timestamp reply* message
 - Copy ORIGINATE TIMESTAMP from the request into the same field of the reply
 - Fill in RECEIVE TIMESTAMP according to its clock immediately upon receipt of the request
 - Fill in TRANSMIT TIMESTAMP according to its clock immediately before the reply is transmitted

Message: Timestamp Request / Reply (cont.)

- One-way and round-trip time calculation
 - Sending time
$$= (\text{RECEIVE TIMESTAMP}) - (\text{ORIGINATE TIMESTAMP})$$
 - Receiving time
$$= (\text{time when the reply arrives at source}) - (\text{TRANSMIT TIMESTAMP})$$
 - Round-trip time
$$= (\text{Sending time}) + (\text{receiving time})$$
 - Sending time and receiving time are accurate

Message: Timestamp Request / Reply (cont.)

- Clock synchronization
 - If the exact one-way time is known
$$\text{Time difference} = (\text{RECEIVE TIMESTAMP}) - (\text{ORIGINATE TIMESTAMP}) - (\text{sending time})$$
- Limitation
 - Exact one-way time is hard to calculate
 - Round-trip time varies time to time
 - Messages may be delayed or dropped, making the calculation inconsistent

Message: Address Mask Request / Reply

- Messages used by a host to obtain its subnet mask
 - The host sends out *address mask request*
 - If the host knows the router's address, sends the message directly
 - Otherwise, broadcasts the message
 - The router returns *address mask reply* to provide the mask

Message: Address Mask Request / Reply (cont.)

o Message format

- Type
 - o Address mask request (17)
 - o Address mask reply (18)
- IDENTIFIER & SEQUENCE NUMBER
 - o To associate replies with requests
- ADDRESS MASK
 - o In request message: all 0s
 - o In reply message: actual subnet mask

0	8	16	31
TYPE (17 / 18)	CODE (0)	CHECKSUM	
IDENTIFIER		SEQUENCE NUMBER	
ADDRESS MASK			

Notes-7

CMPT 471 • 2003-3

29

Message: Router Solicitation / Advertisement

o Router discovery scheme

- After a host boots, it needs to know at least one router on the local network before it can communicate with another network
- The hosts need to know whether a router is alive or functioning

Notes-7

CMPT 471 • 2003-3

30

Message: Router Solicitation / Advertisement (cont.)

o ICMP messages for router discovery:

Router solicitation / advertisement message

- The host broadcasts (multicasts) a *router solicitation* message
- The router(s) receiving the solicitation message broadcast (multicast) their routing information using *router advertisement* message
- Even if no host solicits, a router can periodically send *router advertisement* messages, which announce both its presence and all other routers' presence on the network

Notes-7

CMPT 471 • 2003-3

31

Message: Router Solicitation / Advertisement (cont.)

o Router solicitation message format

0	8	16	31
TYPE (10)	CODE (0)	CHECKSUM	
Reserved			

o Router advertisement message format

0	8	16	31
TYPE (9)	CODE (0)	CHECKSUM	
NUM ADDRS	ADDR SIZE	LIFETIME	
Router address 1			
Preference level 1			
Router address 2			
Preference level 2			
.....			

Notes-7

CMPT 471 • 2003-3

32

Message: Router Solicitation / Advertisement (cont.)

- NUM ADDRESS
 - # of address entries follows
- ADDR SIZE
 - Size of an address in 32-bit units
 - IPv4: ADDR SIZE = 1
- LIFETIME
 - The time (in seconds) a host may use the advertised addresses
- PREFERENCE LEVEL
 - The ranking of the router
 - A host chooses the route with highest preference
 - PREFERENCE LEVEL = 0: default router