ICMP
Internet Control Message Protocol (1)

- What if a router cannot route or deliver a datagram?
- What if a router experiences a congestion?
- What if the TTL expires?

  …
- Router needs to inform source to take action to avoid or correct the problem
- ICMP allows routers and hosts to send error or control messages to other routers or hosts.
- ICMP is an error reporting mechanism, and can only report condition back to the original source
- ICMP is specified in RFC 792

ICMP (2)

- ICMP messages are encapsulated in IP packets, with the protocol type of 1.
- In the data portion of the packet, the first byte of the ICMP message identifies the message type and the format of the rest of the message.
- Most ICMP messages include the full IP header of the datagram they refer to, plus the first 64 bits (8 bytes) of the data portion of that datagram, to help the sender identify the packet.
- Some ICMP packets have a code that further qualifies the type
- To avoid explosion of ICMP packets:
  - No ICMP packets are generated to report errors on ICMP packets
  - If an ICMP message is generated about a fragmented packet, it is generated only for fragment zero.

Some ICMP Message Types

(see RFC1700)

Type Field	ICMP Message Type
0	Echo Reply
3	Destination Unreachable
4	Source Quench
5	Redirect (change a route)
8	Echo Request
9	Router Advertisement
10	Router Solicitation
11	Time Exceeded for a Datagram
12	Parameter Problem on a Datagram
13	Timestamp Request
14	Timestamp Reply
17	Address Mask Request
18	Address Mask Reply

Echo Request/Echo Reply

- Testing Destination Reachability and status
  - Echo Request Message
  - Echo Reply Message
- Command used to send ICMP echo request is, in most systems, called “ping”.
- Echo request may contain some data, which is returned unchanged in the reply message.
- The ICMP Echo Request/Reply header also contains a sequence number and identifier, to aid the host in matching the request with the reply.

Echo Request/Reply Message Format

- ICMP Echo Request or Reply Message Format

<table>
<tr><td>TYPE (8 OR 0)</td><td>CODE (0)</td><td>CHECKSUM</td></tr>
<tr><td colspan="2">IDENTIFIER</td><td>SEQUENCE NUMBER</td></tr>
<tr><td colspan="3">OPTIONAL DATA</td></tr>
<tr><td colspan="3">.................</td></tr>
</table>

ICMP Destination Unreachable

- Reports of unreachable destinations
  - when a router cannot forward or deliver an IP datagram, it sends a “destination unreachable” message back to the original source.
- Code determines specific condition (see table)

Message Format

TYPE 3	CODE (0-12)	CHECKSUM
UNUSED (MUST BE ZERO)		
INTERNET HEADER + FIRST 64 BITS OF DATAGRAM		
....		

ICMP Destination Unreachable Message format

Destination Unreachable Codes

Code Value	Meaning
0	Network Unreachable
1	Host Unreachable
2	Protocol Unreachable
3	Port Unreachable
4	Fragmentation 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

ICMP Source Quench

- Congestion and Datagram Flow Control
  - ICMP source quench messages report congestion to the original source
  - request for the source to reduce its current rate
- usually sent for each datagram discarded
- can be sent by a host or a router
- some routers may be more sophisticated
  - monitor incoming traffic
  - quench sources that have the highest rate
  - avoid congestion by quenching before datagrams are lost

ICMP Source Quench Format

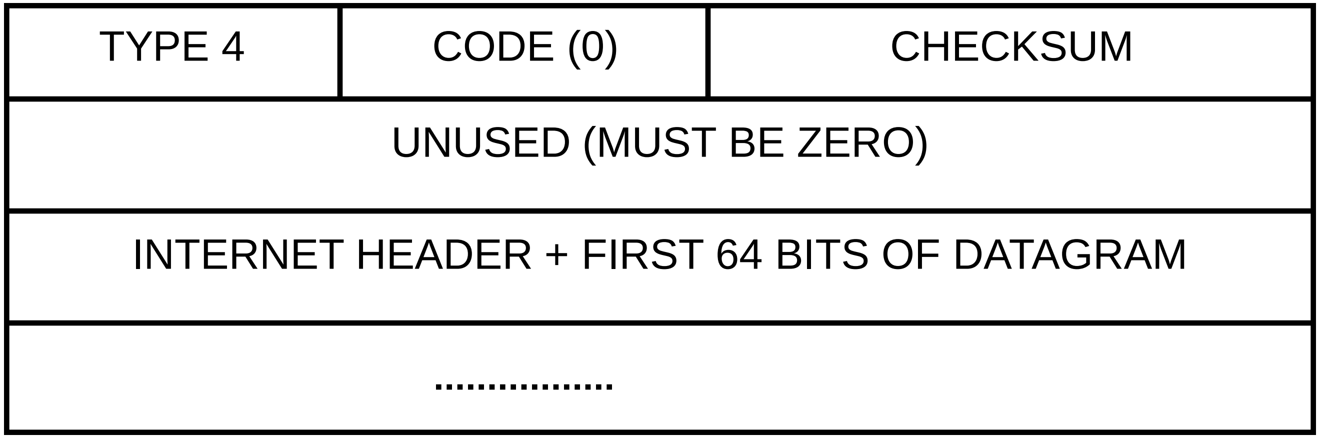

TYPE 4

CODE (0)

CHECKSUM

UNUSED (MUST BE ZERO)

INTERNET HEADER + FIRST 64 BITS OF DATAGRAM

....

ICMP Source Quench Message format

ICMP Redirect Message

- Host sends a datagram to router R1 to be forwarded to a certain destination.
- Router R1 looks at its routing table, and finds the next router in the path, R2.
- If R2 is directly accessible to the sending host, R1 generates an ICMP Redirect Message back to the sending host. R1 also forwards the packet to R2 normally.
- The purpose of the ICMP Redirect message is to inform the host that there is a better route to that destination.

ICMP Redirect Message Format

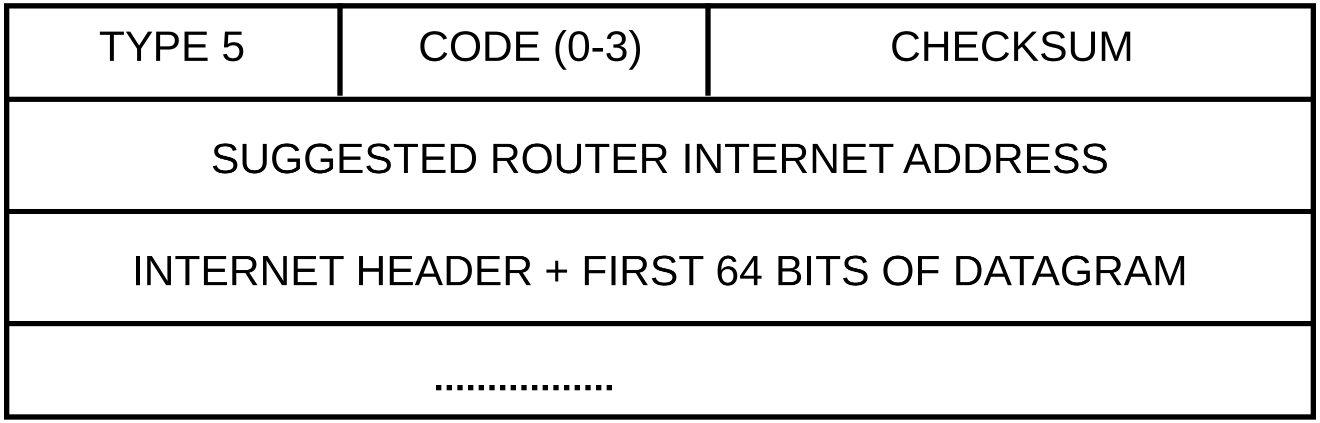

TYPE 5	CODE (0-3)	CHECKSUM
SUGGESTED ROUTER INTERNET ADDRESS		
INTERNET HEADER + FIRST 64 BITS OF DATAGRAM		
....		

ICMP Redirect Message format

Code value	Meaning
0	Redirect datagrams for the Net (now obsolete)
1	Redirect datagrams for the Host
2	Redirect datagrams for the Type of Service and Net
3	Redirect datagrams for the Type of Service and Host

ICMP Time Exceeded Message

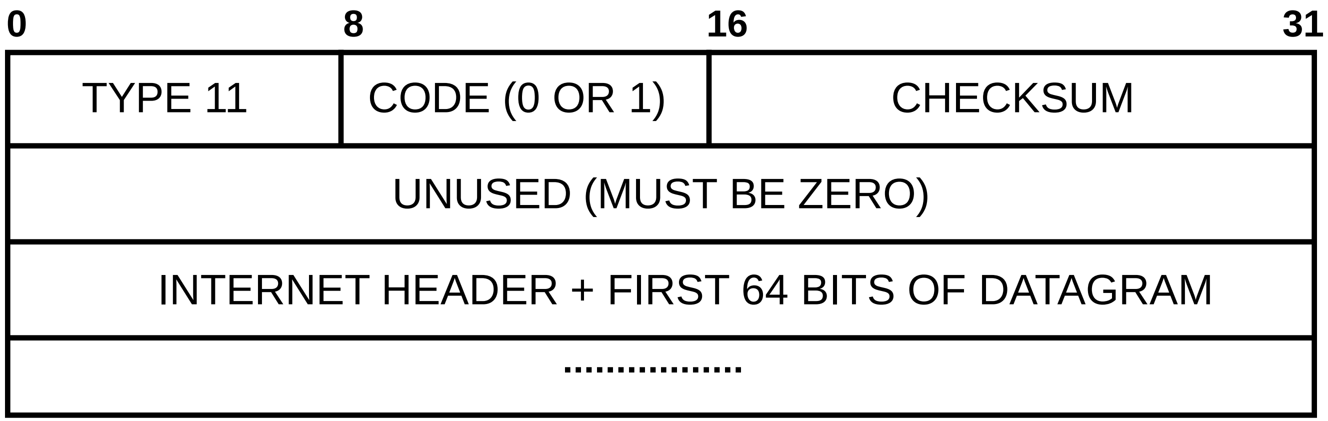

Code value	Meaning
0	Time-to-live count exceeded
1	Fragment re-assembly time exceeded

A router sends this message whenever a datagram is discarded because the time-to-live field in the datagram header has reached zero or because its reassembly timer expired while waiting for fragments.

ICMP Address Mask Request/Reply

- Obtaining a subnet mask
  - ICMP address mask request message
  - ICMP address mask reply message
  - sent to router directly (if known)
  - broadcast (if router unknown)
- Response is unicast if the request contains a valid IP address; otherwise, it is broadcast
- Any host can respond
- Documented in RFC 950

ICMP (16)

TYPE (17 OR 18)	CODE (0)	CHECKSUM
IDENTIFIER		SEQUENCE NUMBER
ADDRESS MASK		

- ICMP address mask request or reply message format. Usually, hosts broadcast a request without knowing which specific router will respond.

ICMP Router Advertisement/Solicitation

- Options for the host to learn the router address(es):
  - Manually enter entries
    - Drawback : not up to date info. & cumbersome
  - Host listens to routing protocol messages
    - Drawback : protocols and their messages differ
  - Use of ICMP messages : “Router Advertisement” and “Router Solicitation”, defined in RFC 1256
- Routers periodically send an ICMP “Router Advertisement”, either broadcast or multicast
- Hosts may solicit Router Advertisement messages with a Router Solicitation message

ICMP Router Advertisement

TYPE = 9	CODE = 0	CHECKSUM
NUM ADDRS	ADDR ENTRY SIZE = 2	LIFETIME (sec)
ROUTER ADDRESS [1]		
PREFERENCE LEVEL [1]		
ROUTER ADDRESS [2]		
PREFERENCE LEVEL [2]		
.		
.		

ICMP Router Solicitation Message

TYPE = 10	CODE = 0	CHECKSUM
RESERVED		

- Default advertisement rate is once every 7 to 10 minutes.
- This router solicitation message causes the router(s) to send their advertisements earlier.
- Advertisement lifetime is typically 30 minutes.

Application: Tracing the Route of a Packet

- Objective: find the path a packet takes between two Internet hosts.
- Originator host sends a series of packets, starting with TTL=1 and increasing the TTL for each packet.
- The first router in the path will drop the TTL=1 packet and send back an ICMP Time Exceeded; host then learns who is the first hop.
- The second router in the path drops the packet with TTL=2 and sends the ICMP Time Exceeded; the third drops the packet with TTL=3 and so on.
- By collecting the ICMP responses, the host can figure out the path taken by the packet.

Traceroute Example

```
elaine1:~> traceroute gatekeeper.dec.com
traceroute to gatekeeper.dec.com (204.123.2.2): 1-30 hops, 38 byte packets
 1  leland-gateway.Stanford.EDU (171.64.15.65)  3.62 ms  3.5 ms  3.49 ms
 2  Core6-gateway.Stanford.EDU (171.64.1.229)  3.74 ms  0.505 ms  0.579 ms
 3  Core4-gateway.Stanford.EDU (171.64.3.114)  1.22 ms  0.725 ms  0.716 ms
 4  i2-gateway.Stanford.EDU (171.64.1.225)  0.792 ms  0.816 ms  0.828 ms
 5  Core-gateway.Stanford.EDU (171.64.1.210)  1.14 ms  1.6 ms  1.9 ms
 6  SUNet-Gateway.Stanford.EDU (171.64.1.34)  1.80 ms  7.77 ms  1.94 ms
 7  f1-0-0.paloalto-cr18.bbnplanet.net (198.31.10.2)  2.39 ms  2.86 ms  2.24
   ms
 8  p3-2.paloalto-nbr2.bbnplanet.net (4.0.3.85)  2.29 ms  2.17 ms  2.76 ms
 9  p1-0.paloalto-nbr1.bbnplanet.net (4.0.5.65)  2.4 ms  3.45 ms  11.9 ms
10  p6-0-0.paix.bbnplanet.net (4.0.1.50)  2.36 ms  1.94 ms  2.45 ms
11  digital-gw1.pa-x.dec.com (4.0.2.114)  3.28 ms 5.84 ms 2.70 ms
12  core-gw1.pa-x.dec.com (204.123.1.1)  2.53 ms 2.76 ms 2.40 ms
13  gatekeeper.dec.com (204.123.2.2)  5.65 ms  4.50 ms  4.5 ms
```

ICMP Traceroute Message

- Current method requires 2*N messages for a N-hop path; will give wrong results if path changes
- ICMP Traceroute (RFC 1393) can do it in N+1 messages.
- Idea: Define a Traceroute IP Option.
- Send an IP packet with this option set.
- Every intermediate system handling this packet will send back an ICMP Traceroute to the source.