Internet Control Message Protocol (ICMP)

Overview

- The IP (Internet Protocol) relies on several other protocols to perform necessary control and routing functions:
  - Control functions (ICMP)
  - Multicast signaling (IGMP)
  - Setting up routing tables (RIP, OSPF, BGP, PIM, …)

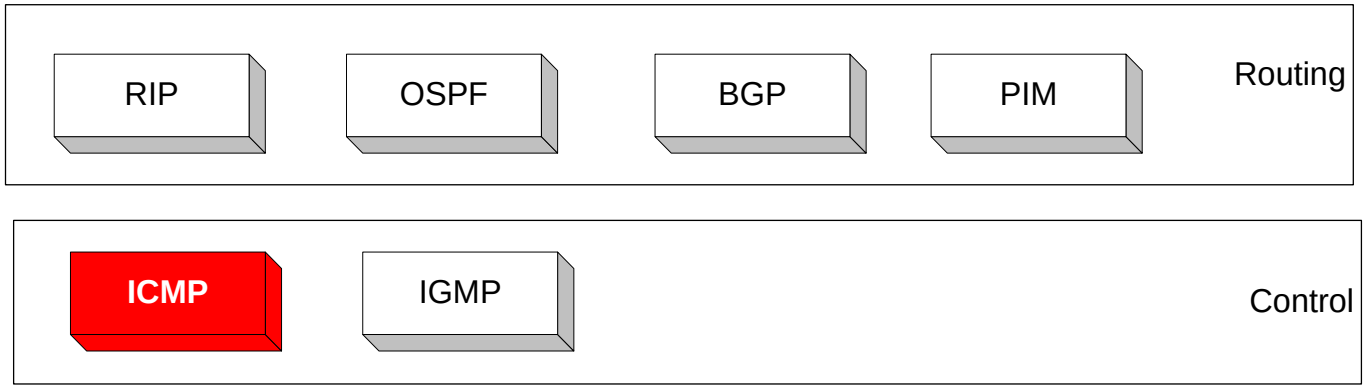

Overview

- The **Internet Control Message Protocol (ICMP)** is a helper protocol that supports IP with facility for
  - Error reporting
  - Simple queries

- ICMP messages are encapsulated as IP datagrams:

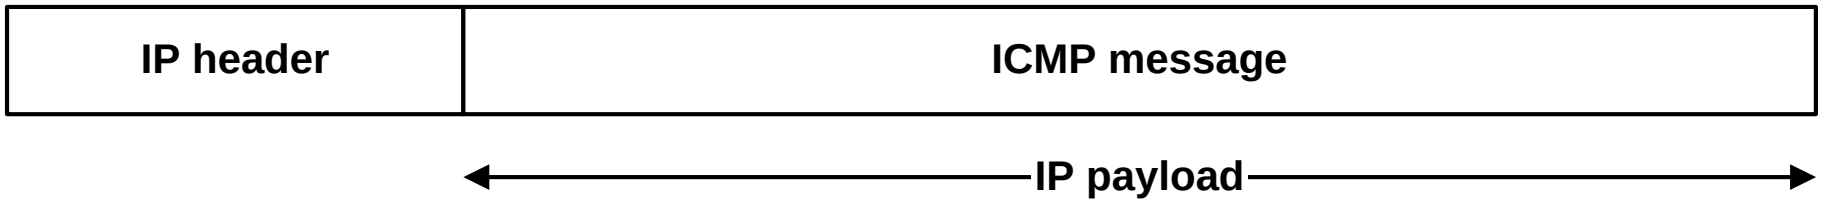

ICMP: Internet Control Message Protocol

(1) ICMP Introduction

- Motivation
  - IP may fail to deliver datagrams because
    - the destination is not available
    - the time-to-live counter expires
    - routers become congested
  - We need to let the sender know what has happened
  - ICMP is a required part of IP

- Purpose
  - ICMP allows routers (and hosts) to send error or control messages to other routers or hosts
  - ICMP provides communication between the Internet Protocol software on one machine and the Internet Protocol software on another

- Restrictions
  - ICMP messages are not generated for errors that result from datagrams carrying ICMP error messages. Why?
  - ICMP is only sent to the original source. Why?

- ICMP Encapsulation
  - ICMP is encapsulated in an IP packet, but is considered part of the IP or Internet layer.

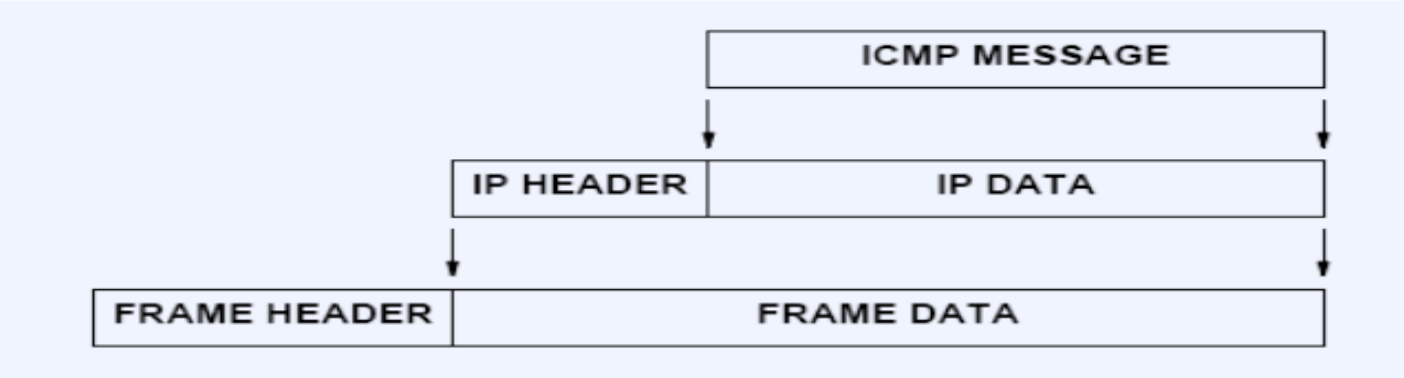

ICMPv4 and ICMPv6

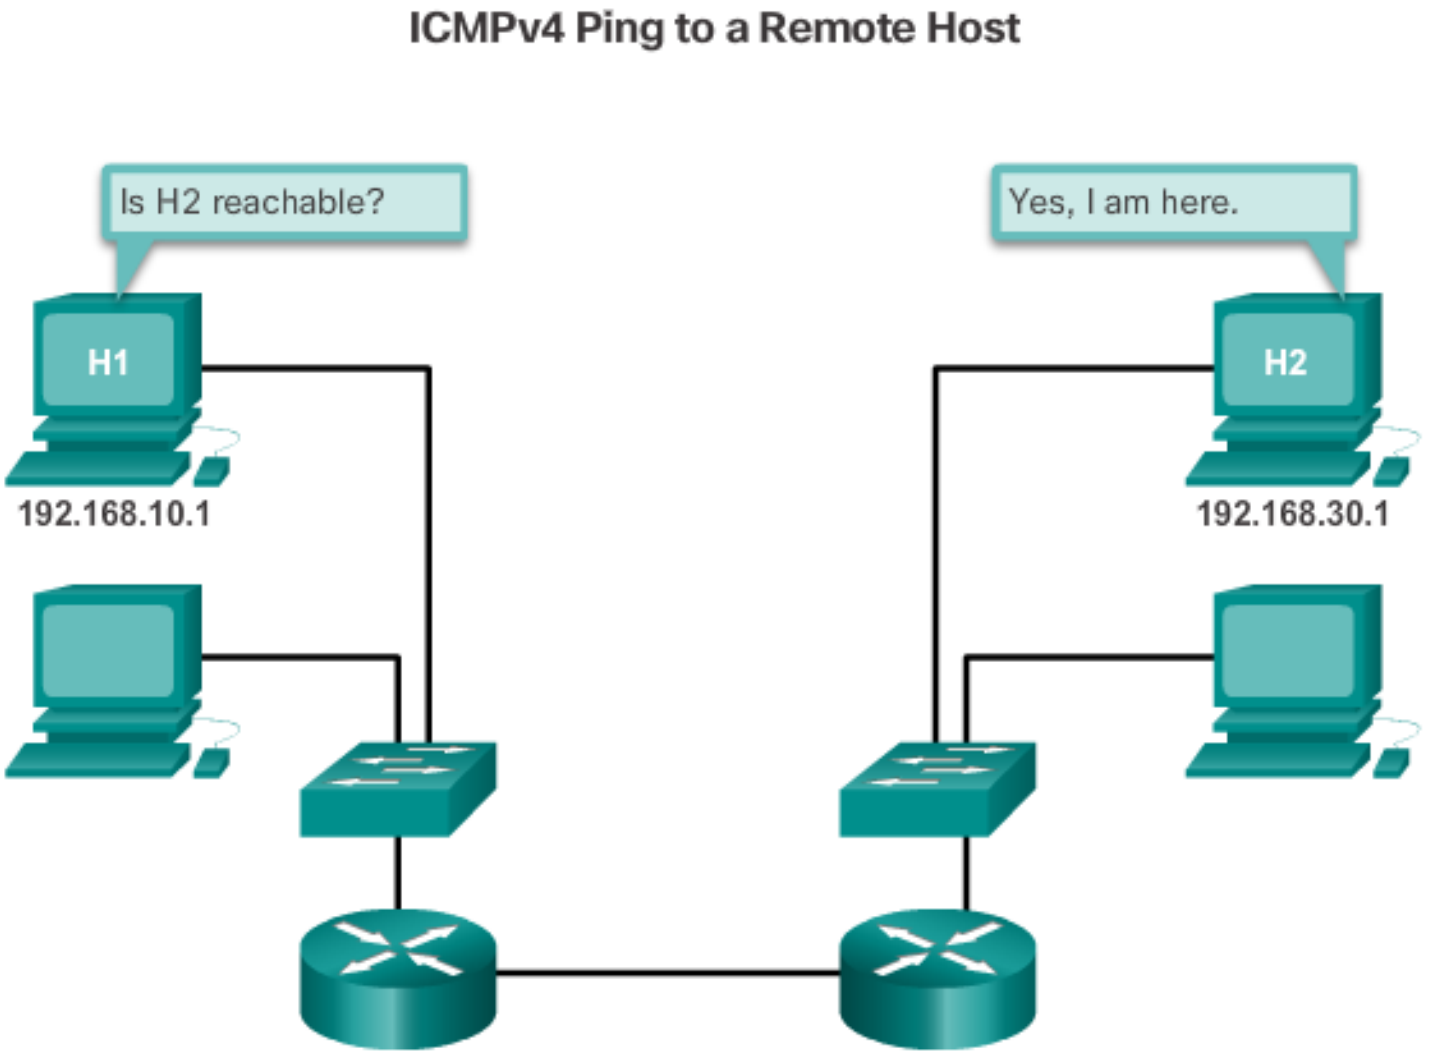

ICMPv4 and ICMPv6

- ICMP messages common to both ICMPv4 and ICMPv6 include:
  - Host confirmation
  - Destination or service unreachable
  - Time exceeded
  - Route redirection
- Although IP is not a reliable protocol, the TCP/IP suite provides for messages to be sent in the event of certain errors. They are sent using the services of ICMP.

ICMP Host confirmation Query message

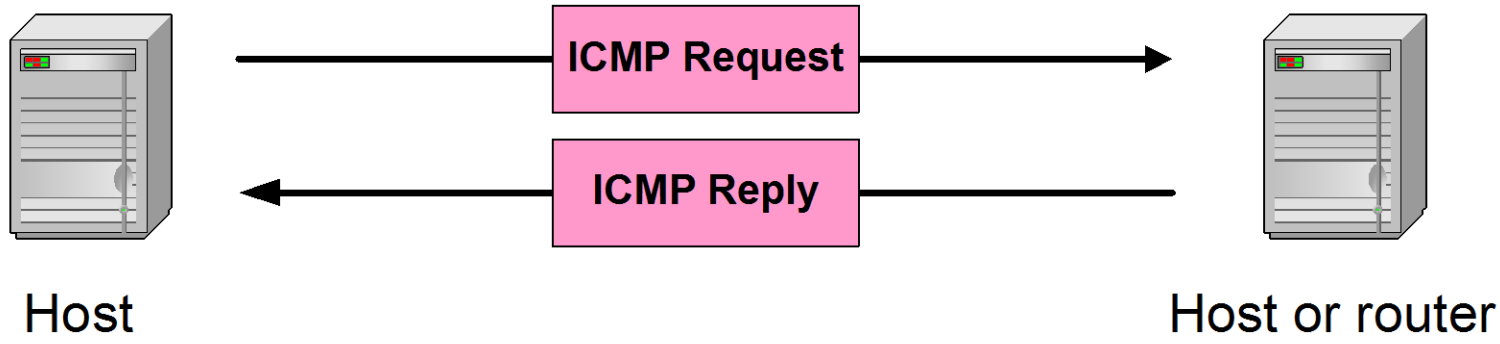

ICMP query:

- Request sent by host to a router or host
- Reply sent back to querying host

ICMP Error message

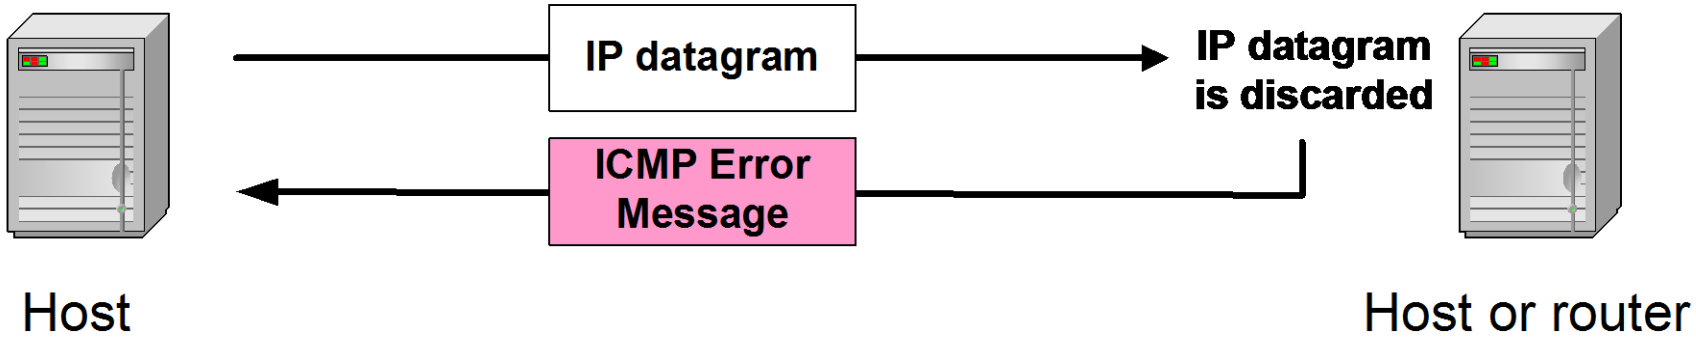

- ICMP error messages report error conditions
- Typically sent when a datagram is discarded
- Error message is often passed from ICMP to the application program

ICMP Error message

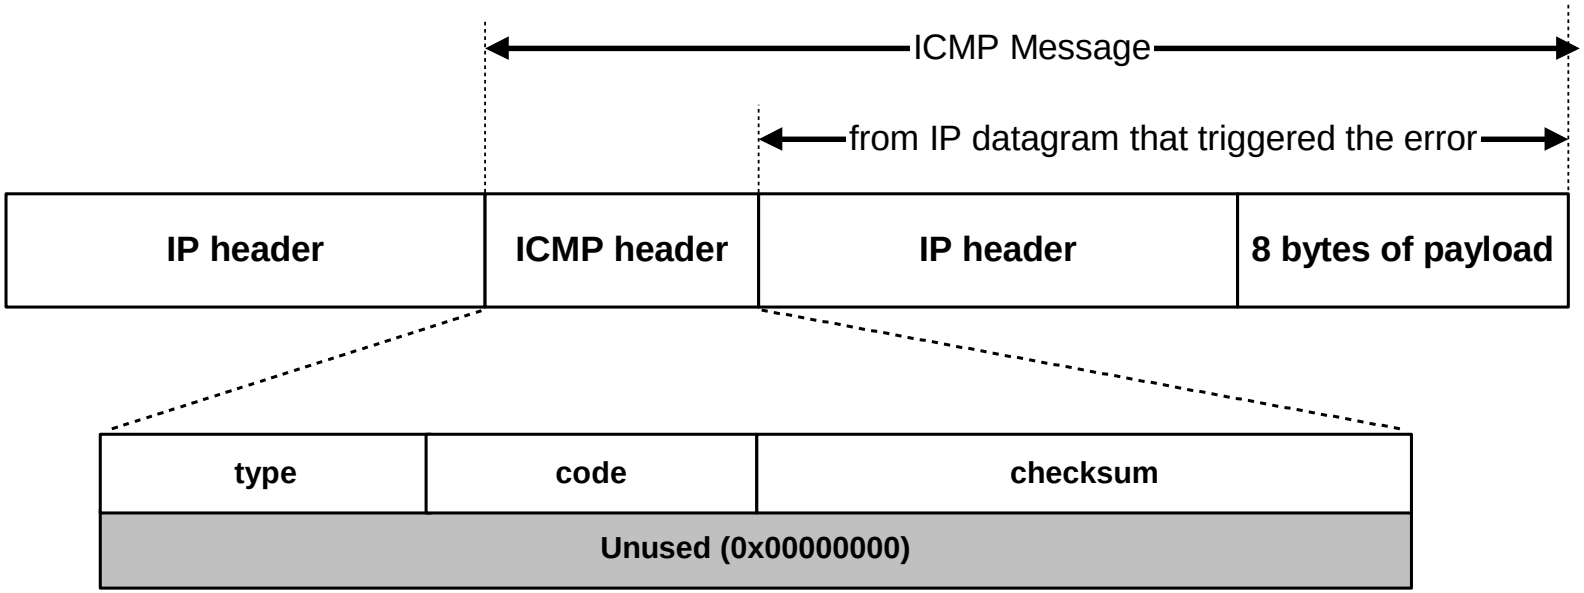

- ICMP error messages include the complete IP header and the first 8 bytes of the payload (typically: UDP, TCP)

Frequent ICMP Error message

Type	Code	Description	
3	0–15	Destination unreachable	Notification that an IP datagram could not be forwarded and was dropped. The code field contains an explanation.
5	0–3	Redirect	Informs about an alternative route for the datagram and should result in a routing table update. The code field explains the reason for the route change.
11	0, 1	Time exceeded	Sent when the TTL field has reached zero (Code 0) or when there is a timeout for the reassembly of segments (Code 1)
12	0, 1	Parameter problem	Sent when the IP header is invalid (Code 0) or when an IP header option is missing (Code 1)

Some subtypes of the “Destination Unreachable”

Code	Description	Reason for Sending
0	Network Unreachable	No routing table entry is available for the destination network.
1	Host Unreachable	Destination host should be directly reachable, but does not respond to ARP Requests.
2	Protocol Unreachable	The protocol in the protocol field of the IP header is not supported at the destination.
3	Port Unreachable	The transport protocol at the destination host cannot pass the datagram to an application.
4	Fragmentation Needed and DF Bit Set	IP datagram must be fragmented, but the DF bit in the IP header is set.

Example: ICMP Port Unreachable

- RFC 792: If, in the destination host, the IP module cannot deliver the datagram because the indicated protocol module or process port is not active, the destination host may send a destination unreachable message to the source host.
- Scenario:

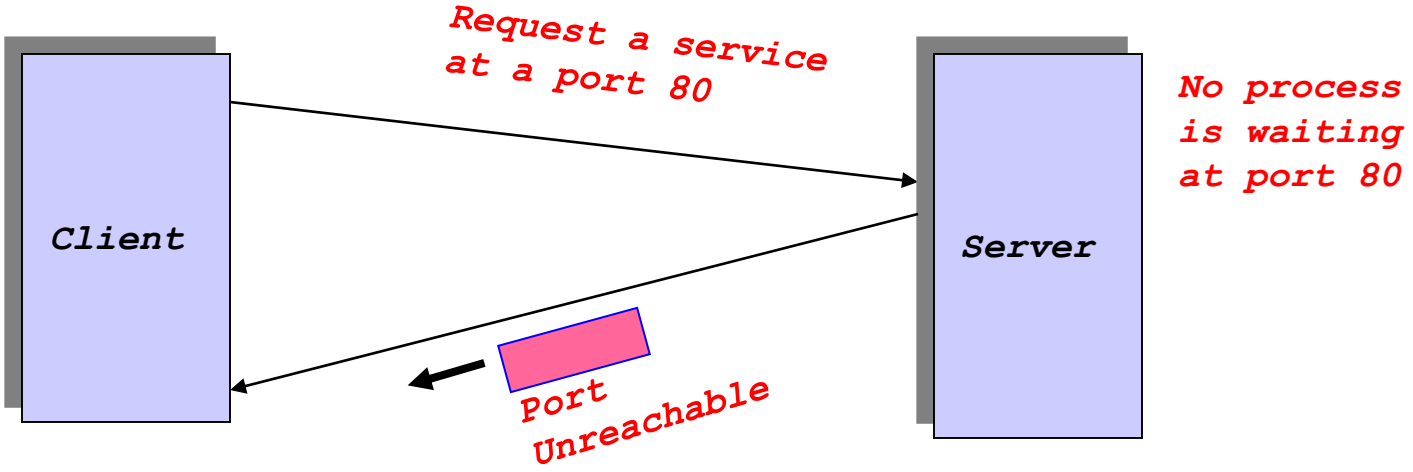

(3) Attacks Using ICMP Messages

- Mapping Network Topology
  - Mapping a target network is a very strategic part of most intelligently planned attacks. This initial step in reconnaissance attempts to discover the live hosts in a target network. An attacker then can direct a more focused scan or exploit toward live hosts only.
  - Sending individual ICMP echo: this is what the `ping` command does.
  - Sending ICMP echo requests to the broadcast addresses of a network.
  - Sending ICMP echo requests to network and broadcast address of subdivided networks
  - Sending an ICMP address mask request to a host on the network to determine the subnet mask to better understand how to map efficiently.

- Smurf Attack
  - Ping a broadcast address, with the (spoofed) IP of a victim as source address
  - All hosts on the network respond to the victim
  - The victim is overwhelmed
  - Keys: Amplification and IP spoofing
  - Protocol vulnerability; implementation can be "patched" by violating the protocol specification, to ignore pings to broadcast addresses
  - ICMP echo just used for convenience
  - All ICMP messages can be abused this way
  - "Fraggle" is the equivalent with UDP

- Ping of Death
  - ICMP echo with fragmented packets
  - Maximum legal size of an ICMP echo packet:
    65535 - 20 - 8 = 65507
  - Fragmentation allows bypassing the maximum size:
    (offset + size) > 65535
  - Reassembled packet would be larger than 65535 bytes
  - OS crashes
  - Same attack with different IP protocols

ICMP Redirect

ICMP Redirect message is sent by a router (R1) to the sender of an IP datagram (host) when the datagram should have been sent to a different router (R2).

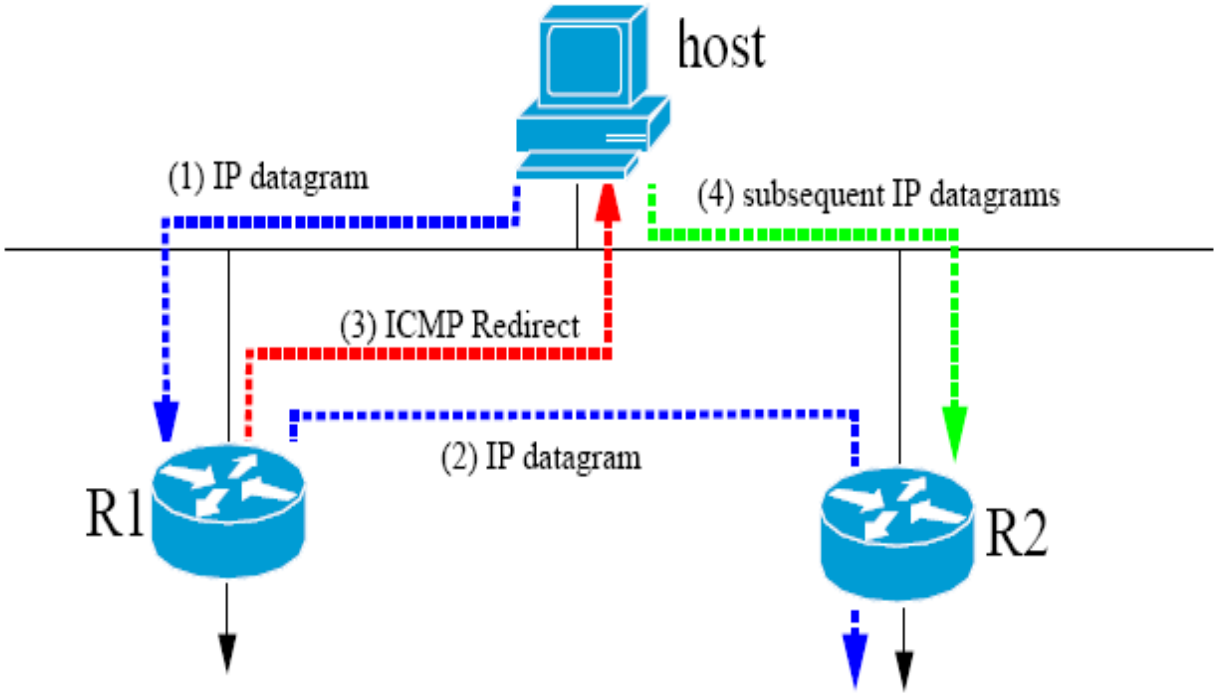

Traceroute

- traceroute works by increasing the "time-to-live" value of each successive batch of packets sent. The first three packets have a time-to-live (TTL) value of one (implying that they make a single hop). The next three packets have a TTL value of 2, and so on. When a packet passes through a host, normally the host decrements the TTL value by one, and forwards the packet to the next host. When a packet with a TTL of one reaches a host, the host discards the packet and sends an ICMP time exceeded (type 11) packet to the sender. The traceroute utility uses these returning packets to produce a list of hosts that the packets have traversed en route to the destination. The three timestamp values returned for each host along the path are the delay (aka latency) values typically in milliseconds (ms) for each packet in the batch. If a packet does not return within the expected timeout window, a star is traditionally printed. traceroute may not list the real hosts, it indicates that the first host is at one hop, the second host at two hops. IP does not guarantee that all the packets take the same route.

Traceroute

- Traceroute capitalizes on the “Time To Live” field in the IP packet header to measure the forwarding path from one host to another. Why might a hop in traceroute might show a “*” (i.e., no IP address for the router at that hop in the path).

  - *The packet may have been lost on the forward path.*
  - *The TIME_EXCEEDED packet may have been lost on the reverse path.*
  - *The router may have been configured not to send TIME_EXCEEDED messages.*
  - *The packet may have encountered a firewall or NAT that drops the packet.*

Traceroute

- Traceroute may report a path that does not exist. For example, traceroute could return a path A-B-C (where A, B, and C are IP addresses) when there is no link between routers B and C.

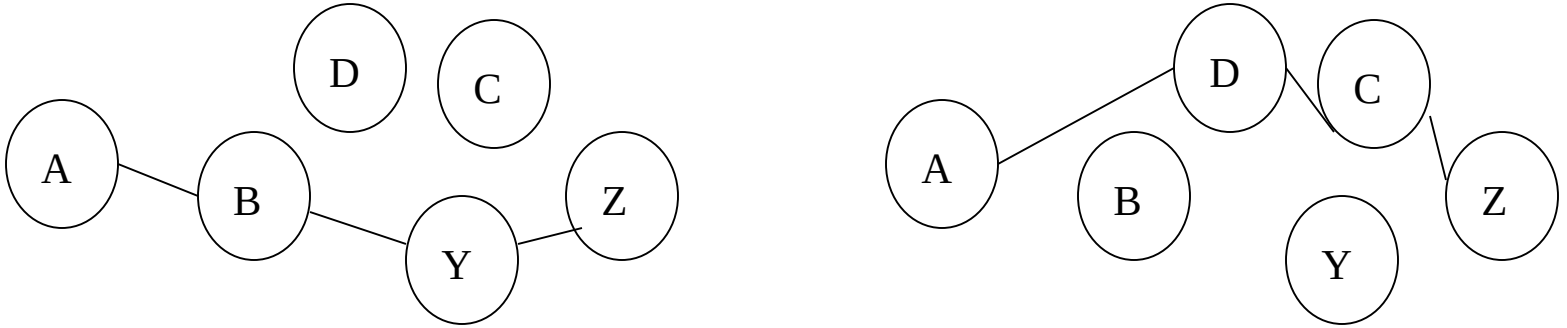

- *The route to the remote destination may have changed during the measurement process. For example, suppose the route from A to Z changed from A-B-Y-Z to A-D-C-Z. Then, the first traceroute probe would return "B" and the second would return "C", even though nodes B and C are not directly connected.*

ICMPv6 Router Solicitation and Router Advertisement Messages

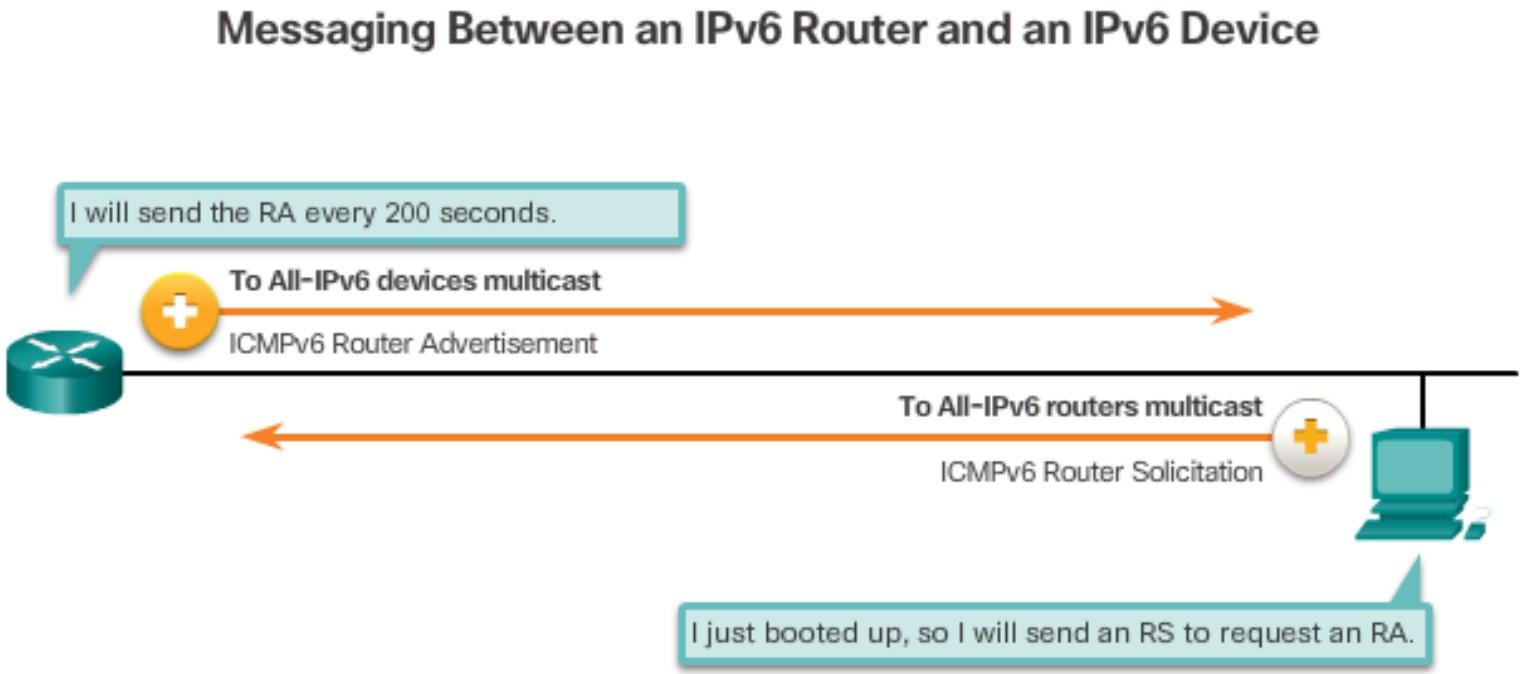

RA messages are sent by routers to provide addressing information to hosts using SLAAC. The RA message can include addressing information for the host such as the prefix, prefix length, DNS address and domain name. A router will send an RA message periodically or in response to an RS message. A host using SLAAC will set its default gateway to the link-local address of the router that sent the RA.

ICMPv6 Router Solicitation and Router Advertisement Messages

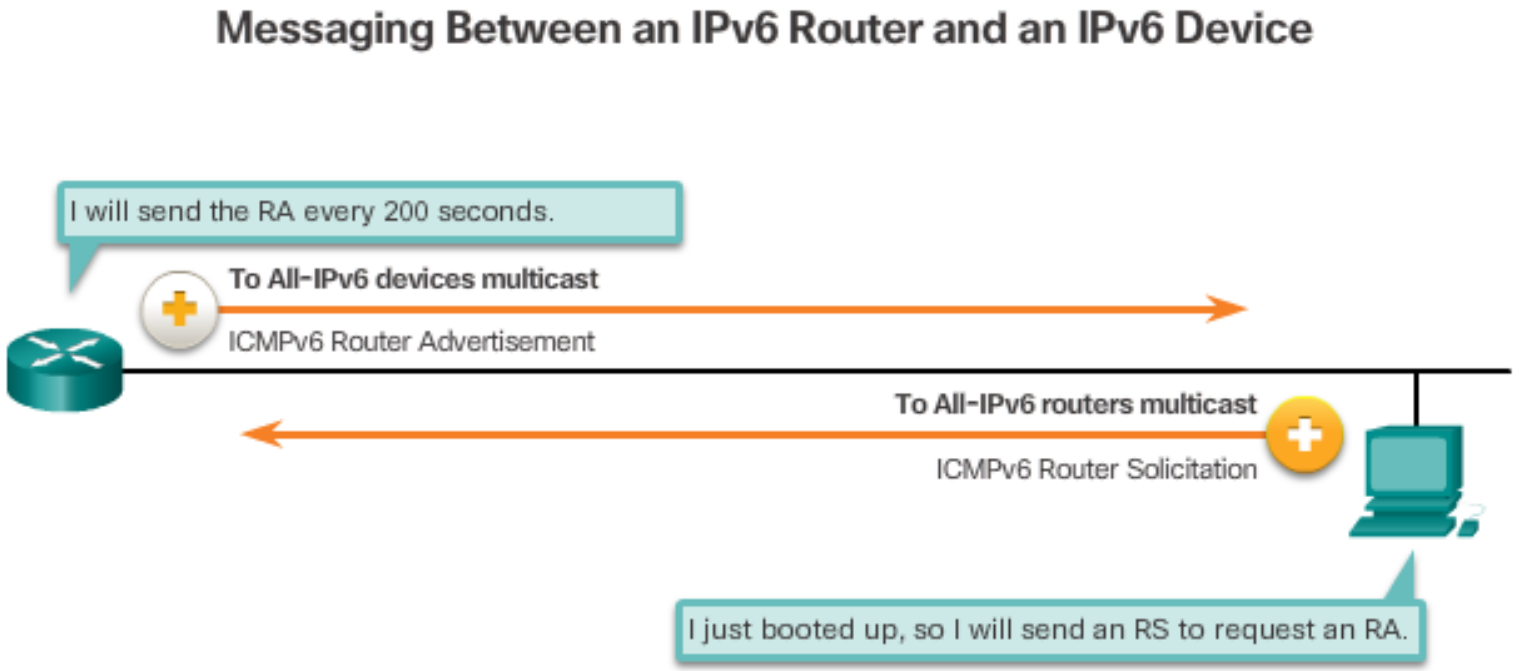

When a host is configured to obtain its addressing information automatically using Stateless Address Autoconfiguration (SLAAC), the host will send an RS message to the router requesting an RA message.

ICMPv6 Router Solicitation and Router Advertisement Messages (cont.)

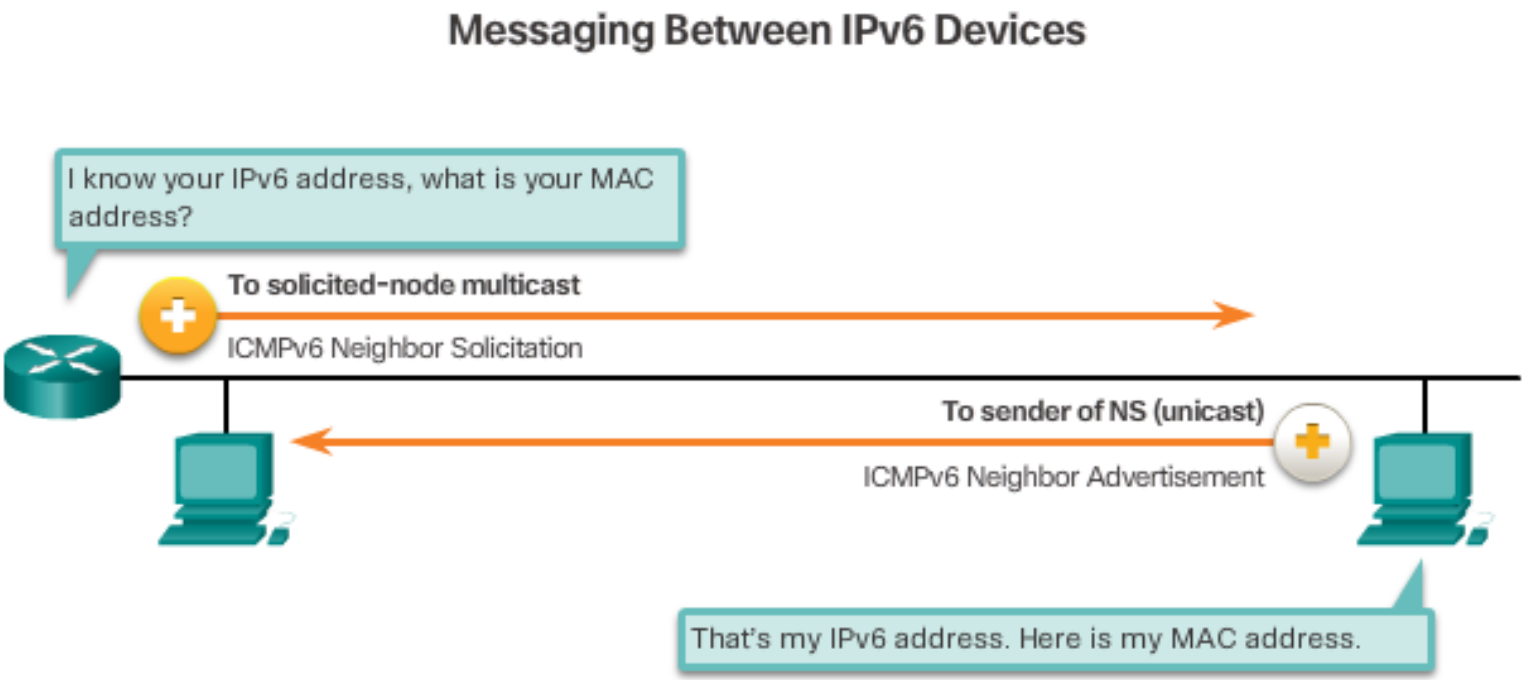

NS messages are sent when a device knows the IPv6 address of a device but does not its MAC address. This is equivalent to an ARP Request for IPv4.

ICMPv6 Router Solicitation and Router Advertisement Messages (cont.)

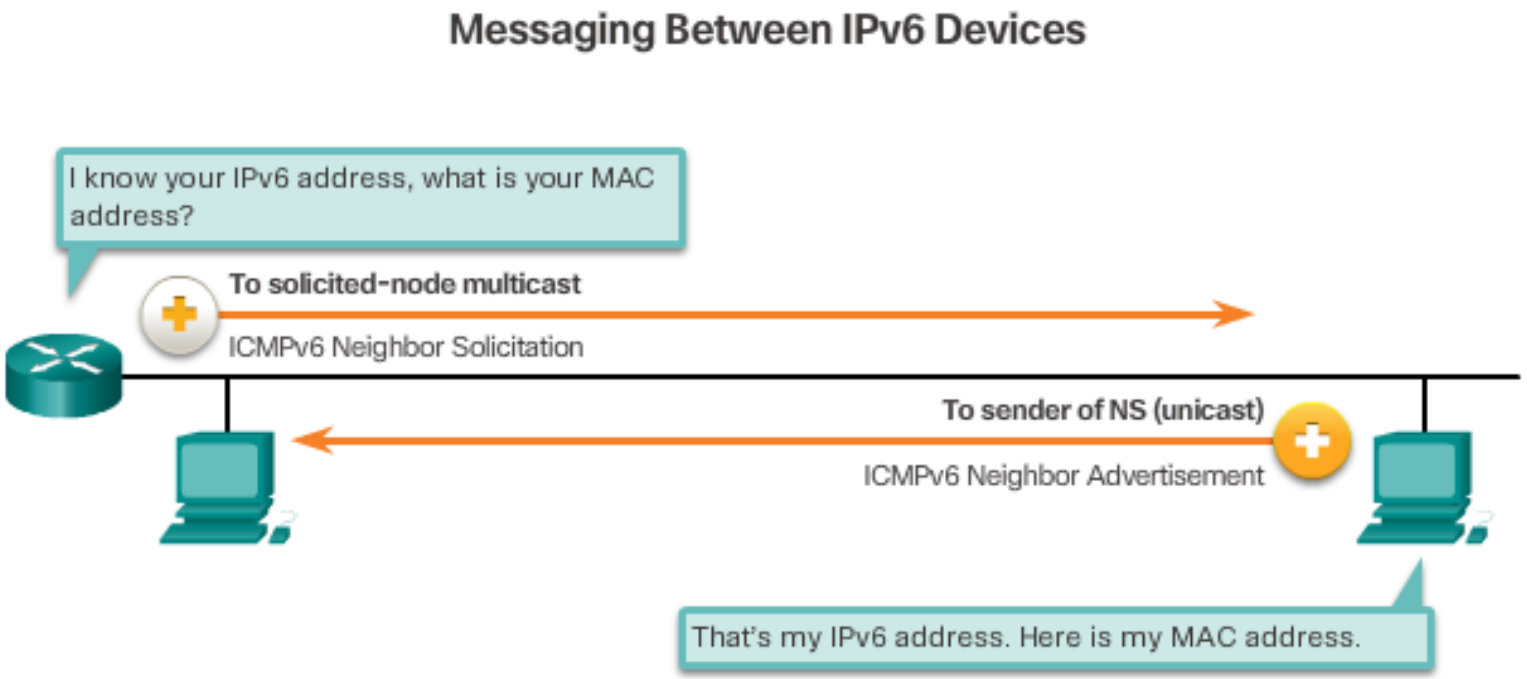

NA mesages are sent in response to an NS message and matches the target IPv6 address in the NS. The NA message includes the device's Ethernet MAC address. This is equivalent to an ARP Reply for IPv4.

ICMPv6 Router Solicitation and Router Advertisement Messages (cont.)

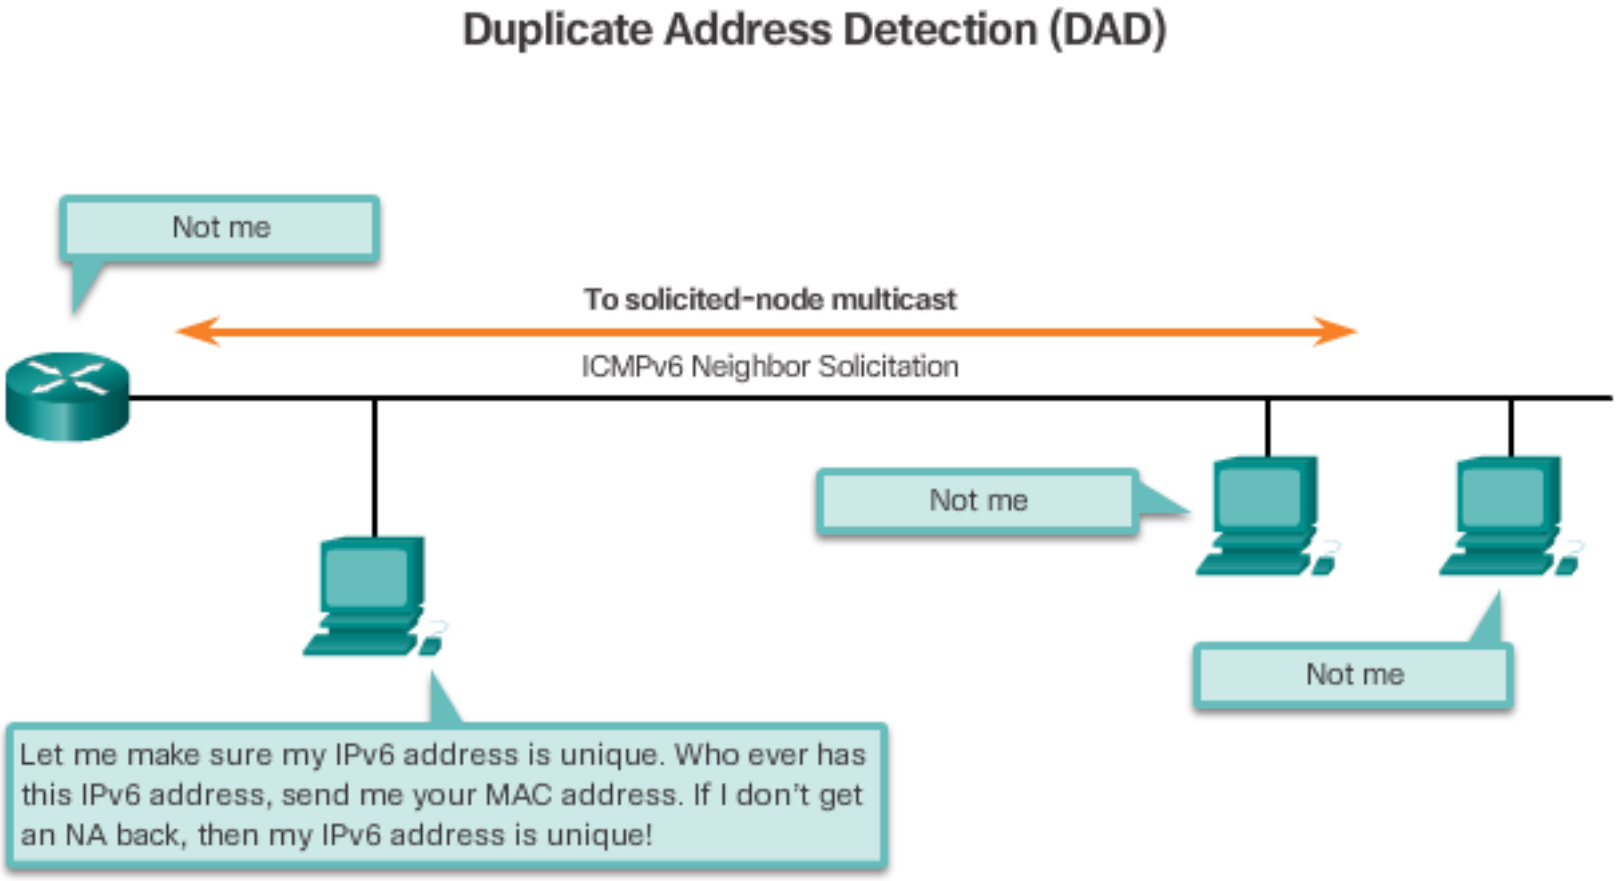

References

- http://www.redbooks.ibm.com/abstracts/gg243376.html
- http://en.wikipedia.org/wiki/Internet_Control_Message_Protocol