

TCP/IP illustrated Vol. 1

The Protocols

Chapter 7 - Ping Program

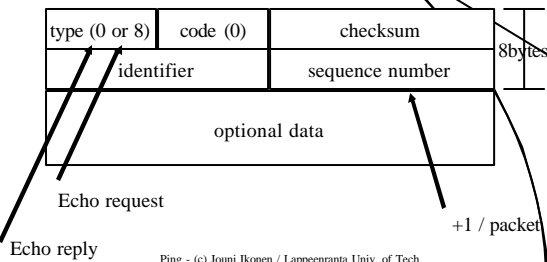
Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Ping

- Tests whether another host is reachable.
- Sends an **ICMP** echo request message to a host, expecting an **ICMP** echo reply to be returned.
- However routers can treat ping messages as less important messages and hosts / organizations can block them due security issues.
- Can also be used to examine the IP record route and timestamp options.

Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Format of the Ping message



Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Ping details

- Most TCP/IP implementations support ping server directly in the kernel.
- Server must echo the *identifier* and *sequence* number fields + optional data.
- *Identifier* field in the ping message is set to process number of the program (at least in Unix), this allows ping program to distinguish between multiple instances.

Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Ping example

```
turgon:-> ping 157.24.23.48
157.24.23.48 is alive

turgon:-> ping -s 157.24.23.48
PING 157.24.23.48: 56 data bytes
64 bytes from dior.it.lut.fi (157.24.23.48): icmp_seq=0. time=3. ms
64 bytes from dior.it.lut.fi (157.24.23.48): icmp_seq=1. time=1. ms
64 bytes from dior.it.lut.fi (157.24.23.48): icmp_seq=2. time=2. ms
64 bytes from dior.it.lut.fi (157.24.23.48): icmp_seq=3. time=1. ms
^C
----157.24.23.48 PING Statistics----
4 packets transmitted, 4 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 1/1/3
```

Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Live capture

```
Frame 1 (74 bytes on wire, 74 bytes captured)
Arrival Time: Sep 22, 2004 11:00:32.420813000
Time delta from previous packet: 0.000000000 seconds
Time since reference or first frame: 0.000000000 seconds
Frame Number: 1
Packet Length: 74 bytes
Capture Length: 74 bytes
Ethernet II, Src: 00:04:76:8e:a9:f2, Dst: 00:0c:ce:0a:40:00
Destination: 00:0c:ce:0a:40:00 (Cisco_Oa:40:00)
Source: 00:04:76:8e:a9:f2 (3Com_8e:a9:f2)
Type: IP (0x0800)
```

Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

...Live capture

```
Internet Protocol Src Addr: 157.24.25.159 (157.24.25.159), Dst Addr: 157.24.24.1 (157.24.24.1)
Version: 4
Header length: 20 bytes
Differentiated Services Field : 0x00 (DSCP 0x00: Default, ECN: 0x00)
0000 00.. = Differentiated Services Codepoint Default (0x00)
... 0.. = ECN - Capable Transport (ECT): 0
... 0.. = ECN - CE: 0
Total Length: 60
Identification: 0x2a0e (10766)
Flags: 0x00
..0.. = Don't fragment Notset
..0.. = More fragments: Notset
Fragment offset: 0
Time to live: 128
Protocol: ICMP (0x01)
Header checksum: 0x4e2 (correct)
Source: 157.24.25.159 (157.24.25.159)
Destination: 157.24.24.1 (157.24.24.1)

Internet Control Message Protocol
Type: 8 (Echo (ping) request)
Code: 0
Checksum: 0xc244 (correct)
Identifier: 0x0200
Sequence number: 0x8917
Data (32 bytes)
```

Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

More ping details

- Round trip time is calculated by storing send time to *optional data* field. When reply is received this value is subtracted from current time.
- Hosts might have different time resolution available and this usually impacts the output format.
- ARP and DNS query can delay sending of the first ping

Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

```
Usage: ping [-t] [-a] [-n count] [-l size] [-f] [-i TTL] [-v TOS]
[-r count] [-s count] [[-j host-list] | [-k host-list]]
[-w timeout] destination-list
```

Options:

- t Ping the specified host until stopped. To see statistics and continue - type Control-Bre To stop - type Control-C.
- a Resolve addresses to hostnames.
- n count Number of echo requests to send.
- l size Send buffer size.
- f Set Don't Fragment flag in packet.
- i TTL Time To Live.
- v TOS Type Of Service.
- r count Record route for count hops.
- s count Timestamp for count hops.
- j host-list Loose source route along host-list.
- k host-list Strict source route along host-list.
- w timeout Timeout in milliseconds to wait for each reply.

Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Windows ping

Windows ping example

```
C:\>ping -n 10 www.funet.fi

Pinging nic.funet.fi[193.166.3.1] with 32 bytes of data:

Reply from 193.166.3.1: bytes=32 time=16ms TTL=248
Reply from 193.166.3.1: bytes=32 time<10ms TTL=248
Reply from 193.166.3.1: bytes=32 time<10ms TTL=248
(some lines removed)
Reply from 193.166.3.1: bytes=32 time<10ms TTL=248

Ping statistics for 193.166.3.1:
    Packets: Sent = 10, Received = 10, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 16ms, Average = 1ms
```


Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

IP route record option

- Most ping programs offer route record option, which sets RR to outgoing ICMP packet.
- This causes “every” router to add its IP address to a list in options list.
- Not all host support all ICMP options
- IP options header has only space for 9 addresses (60-20-3 = 37 bytes / 4 bytes / IP)

Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Format of RR option in IP header



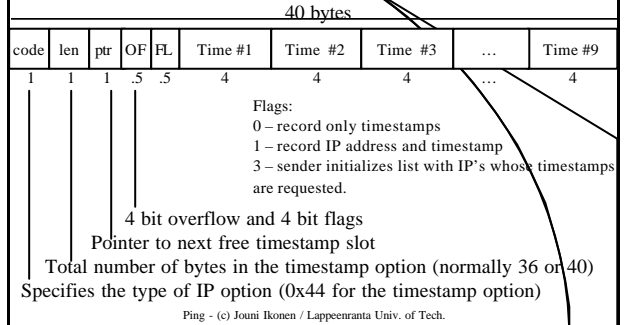
RR ping example

```
jumi:> ping -c 2 -R www.lut.fi
PING www.lut.fi(157.24.8.108): 56 data bytes
64 bytes from 157.24.8.108: icmp_seq=0 ttl=63 time=2.0 ms
RR: jumi.lut.fi (157.24.54.9)
    ty-gw.cc.lut.fi (157.24.10.17)
    www.lut.fi(157.24.8.108)
    www.lut.fi(157.24.8.108)
    ty-gw.cc.lut.fi (157.24.52.1)
    jumi.lut.fi(157.24.54.9)
64 bytes from 157.24.8.108: icmp_seq=1 ttl=63 time=1.9 ms    (same route)

--- www.lut.fi ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 1.9/1.9/2.0 ms
```

Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

IP timestamp option



Timestamp values

- Preferred timestamp values are milliseconds since midnight (UTC).
 - If router has no access to global time information it can use any format it “wishes”, but must set high order bit of the timestamp to indicate the nonstandard time.
 - E.g. devices time can present milliseconds since last reboot.
 - If router can not add timestamp it increments *overflow* field.
- Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Not feasible

- Very limited size to record both IP and timestamp
 - Recording only timestamps is not useful as routes are not usually guaranteed to be fixed.
 - No control of accuracy of timestamps in each router.
 - There are better ways to measure hop times between routers (Traceroute)
- Ping - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

TCP/IP illustrated Vol. 1 The Protocols

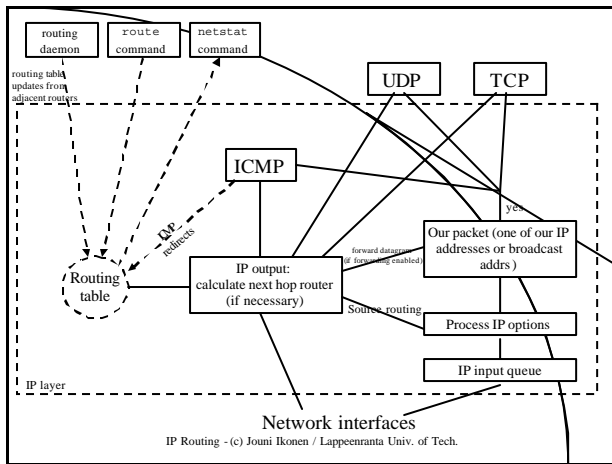
Chapter 9 - IP Routing

IP Routing - (c) Jouni Ikonen / Lappeenranta

IP routing (use of routing table)

- Routing is one of most important functions of IP.
- Datagrams to be routed are generated by:
 - local host
 - some other host (and processing host is configured as a router).
- Routing information between adjacent routers are exchanged by routing protocols, e.g. RIP, BUT now we are interested how a single IP layer makes routing decisions.

IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.



Mechanisms vs. policies

- Routing mechanism
 - Done by the IP-layer.
 - Search through routing table and decide to which interface send packet to.
- Routing policy
 - Set of rules about which rules go into the routing table
 - Normally done by routing daemon

IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Routing table

- Entry has:
 - Flags
 - Destination IP address (host, network, or default)
 - a next-hop router IP address (for an indirect route)
 - a pointer to a local interface to use
- Routing table is searched for every datagram system generates or forwards.

IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Example routing table

```

/sbin/route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Refs Use Iface
157.24.103.12 0.0.0.0 255.255.255.255 UH 0 0 19 eth1
157.24.103.13 0.0.0.0 255.255.255.255 UH 0 0 1 eth2
157.24.103.0 0.0.0.0 255.255.255.0 U 0 0 120 eth0
157.24.123.0 157.24.103.10 255.255.255.0 UG 0 0 0 eth0
127.0.0.0 0.0.0.0 255.0.0.0 U 0 0 3 lo
0.0.0.0 157.24.103.1 0.0.0.0 UG 1 0 299 eth0
    
```

The "distance" to the target

Number of references to a route

Count of lookups for the route

IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Routing table flags

- U The route is UP
- G The route is a gateway (router). If not set the destination is directly connected.
- H The route is to a host, that is, the destination is a complete host address. If not set the route is to a network, and the destination is a network ID
- D The route was created by redirect
- M The route was modified by a redirect

IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Example routing table2

```
C:\>netstat -rn (tai route print)
Active Routes:
Network Destination    Netmask          Gateway          Interface        Metric
0.0.0.0                0.0.0.0          192.168.10.1    192.168.10.2    1
127.0.0.0              255.0.0.0        127.0.0.1      127.0.0.1      1
192.168.10.0           255.255.255.0   192.168.10.2    192.168.10.2    1
192.168.10.2           255.255.255.255 127.0.0.1      127.0.0.1      1
192.168.10.255         255.255.255.255 192.168.10.2    192.168.10.2    1
224.0.0.0              224.0.0.0        192.168.10.2    192.168.10.2    1
255.255.255.255       255.255.255.255 192.168.10.2    192.168.10.2    1
Default Gateway:      192.168.10.1
=====
Persistent Routes:
None
```

IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Routing principles

- Go through routing table
 - Search for matching host address
 - Search for matching network address
 - Search for default entry
- Static routing information can be entered in to the host's configuration files. These can include default gateway.

IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Commands

- Ifconfig, ipconfig
- Route, netstat


IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

What goes to a routing table?

1. Host is not connected to a LAN
 - Only *loopback* interface
2. Host is connected to single LAN
 - *Loopback* interface and *LAN* entry
3. Host is connect to a LAN, which has access to Internet
 - *Loopback*, *LAN* entry and normally *default*
4. Host is connect to a LAN, which has access to Internet and specific routes to other hosts.
 - *Host* entries, *Loopback*, *LAN* entry and normally *default*


IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Example: Routing table for A



IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

Example: Routing table for A



```

/sbin/route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
192.168.1.254 0.0.0.0 255.255.255.255 UH 0 0 19 eth1
192.168.1.0 0.0.0.0 255.255.255.0 U 0 0 1120 eth0
127.0.0.0 0.0.0.0 255.0.0.0 U 0 0 3 lo
0.0.0.0 192.168.1.1 0.0.0.0 UG 1 0 2999 eth0
    
```

IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.


ICMP host & Network Unreachable errors

- How about if there is no match in the Routing table?
 - If datagram is generated by the processing host the error is returned to the application (host / network unreachable)
 - If datagram is sent by some other host (routed), ICMP host unreachable error is returned to the sender.

IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

ICMP redirect

- Host sends a datagram to default router R2
- R2 determines that R1 is correct router in the same network and sends datagram to it and ICMP redirect message to host to correct its routing table



IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

creates a route for each host separately!

ICMP route discovery

- Newer way after statistic routes to use the ICMP router advertisement and solicitation messages. RFC1256
- On boot up a host broadcasts (or multicasts) a solicitation message.
- One or more routers respond with router advertisement messages (with one or more routes).
- Routers send advertisement messages periodically (randomized period).

IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.

ICMP router advertisement

type (9)	code (0)	checksum	8bytes
# of addresses	Address entry size (2)	lifetime	
router address [1]			
preference level [1] (larger value – more preferable)			
router address [2]			
preference level [3]			
...			

IP Routing - (c) Jouni Ikonen / Lappeenranta Univ. of Tech.