

IP network tools & troubleshooting

AFNOG 2010
Kigali, Rwanda
May 2010



Network configuration

- Reminder, configure your network in /etc/rc.conf (x = your IP, from .100 to ...)

```
ifconfig_bge0="196.200.218.x/24"  
defaultrouter="196.200.218.254"  
ipv6_enable="YES"  
ipv6_ifconfig_bge0="2001:4348:0:218:196:200:2  
18:x"  
ipv6_defaultrouter="  
2001:4348:0:218:196:200:218:254"
```

Network configuration

- You can do this from the command line:


```
ifconfig bge0 196.200.218.x/24  
route add default 196.200.218.254
```

```
ifconfig bge0 inet6 2001:4348:0:218:196:200:218:x  
route add -inet6 default  
2001:4348:0:218:196:200:218:254
```

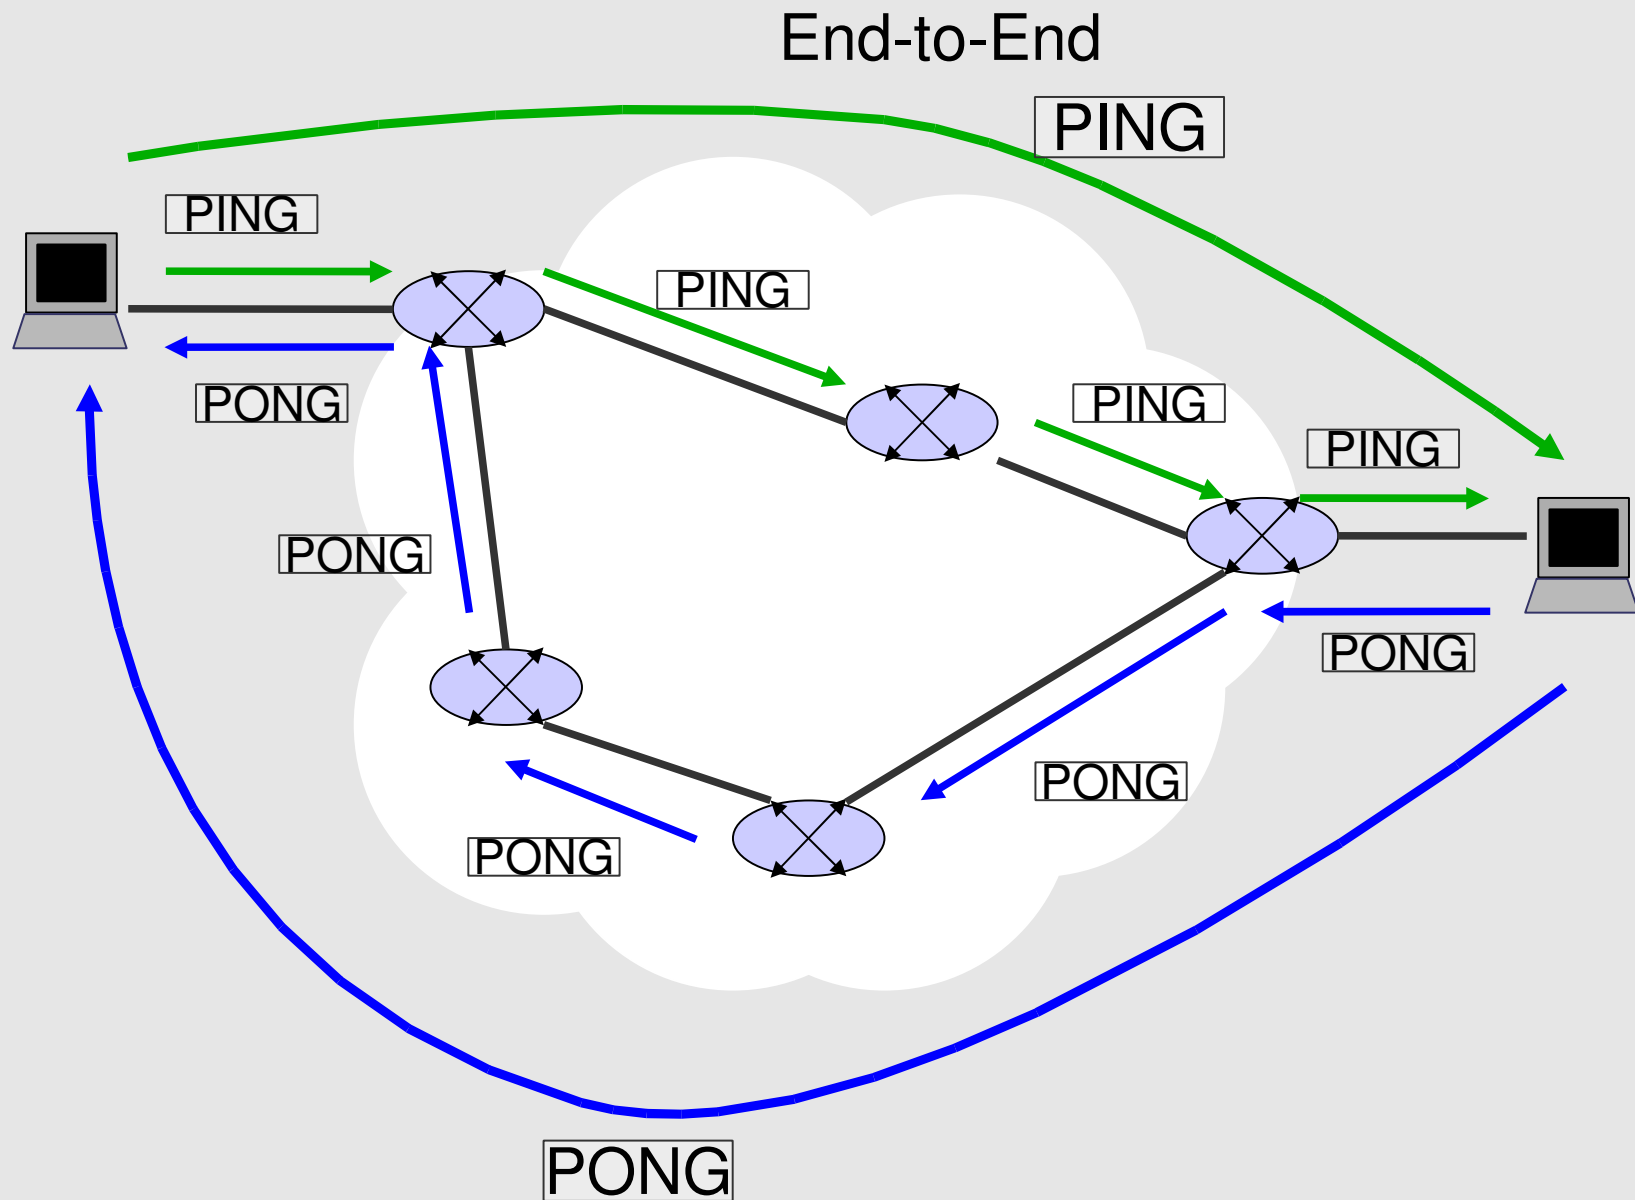
The IP end-to-end principle

- IP is an end-to-end protocol
- The network doesn't keep track of connections
- The host takes a decision on where to send *each packet*
- The network equipment takes a decision on where to forward packets *every time*
- **The path is not necessarily symmetric**
- Cost constraints, reconfiguration of the network, network failures can make the IP packets

IP path




IP path



IP path

End-to-End

TIMEOUT



Network tools

- What network tools can we use to troubleshoot ?
- **ping** – requests echo reply from a computer
- **tracert** – show path taken by IP packets through a network
- **tcpdump** – show network traffic
- **netstat** – show routing entries and listening/active sockets
- **arp** – show/modify the IP <-> MAC address table
- **ndp** – show debug/ndp (Neighb. Disc. Protocol)
- **route** – show/modify the routing table
- **mtr** – combines ping & traceroute

ping

- usage:

ping hostname_or_IP_address

ping6 hostname_or_IPv6_address

- ping sends an ICMP/ICMP6 echo request (type 8), and the responding host sends an ICMP/ICMP6 echo reply (type 0)
- ICMP and ICMP6 sit on top of IP, side by side with TCP and UDP

ping – sample output

```
# ping 196.200.218.254
PING 196.200.218.254 (196.200.218.254): 56 data bytes
64 bytes from 196.200.218.254: icmp_seq=0 ttl=255 time=0.424 ms
64 bytes from 196.200.218.254: icmp_seq=1 ttl=255 time=0.429 ms
64 bytes from 196.200.218.254: icmp_seq=2 ttl=255 time=0.468 ms
...
```

```
# ping6 2001:4348:0:223:196:200:223:254
PING6(56=40+8+8 bytes) 2001:4348:0:218:196:200:218:1 -->
2001:4348:0:223:196:200:223:254
16 bytes from 2001:4348:0:223:196:200:223:254, icmp_seq=0 hlim=64 time=0.426 ms
16 bytes from 2001:4348:0:223:196:200:223:254, icmp_seq=1 hlim=64 time=0.451 ms
16 bytes from 2001:4348:0:223:196:200:223:254, icmp_seq=2 hlim=64 time=0.446 ms
```

Traceroute

- discover path taken by packets on the way to another host
- usage:

```
$ traceroute [-n] hostname_or_IP  
( -n == no DNS lookup )
```

```
# traceroute afnog.org  
traceroute to afnog.org (196.216.2.34), 64 hops max, 40 byte packets  
1 196.200.218.254 (196.200.218.254) 0.435 ms 0.323 ms 0.311 ms  
2 1181-2-205-33-192-81.1181-2.iam.net.ma (81.192.33.205) 1.628 ms 1.330 ms  
1.367 ms  
3 172.20.2.31 (172.20.2.31) 1.485 ms 1.517 ms 1.423 ms  
4 ppp-20-3-217-212.dialup.iam.net.ma (212.217.3.20) 1.360 ms 1.376 ms 1.443  
ms  
5 pal2-almaghrib-2.pal.seabone.net (195.22.197.41) 58.213 ms 58.178 ms  
58.205 ms  
6 POS4-3.BR1.LND9.ALTER.NET.25 (146.188.70.25) 70.771 ms 68.942 ms 70.539 ms
```

Traceroute - IPv6

- usage:

```
$ traceroute6 [-n] hostname_or_IPv6  
( -n == no DNS lookup )
```

```
# traceroute6 -n x1.x0.dk  
traceroute6 to x1.x0.dk (2001:41d0:1:2cc8::1) from  
2001:4348:0:218:196:200:218:1, 64 hops max, 12 byte packets
```

1	2001:4348:0:218:196:200:218:254	0.449 ms	0.363 ms	0.338 ms
2	2001:418:1:101::1	232.759 ms *	232.122 ms	
3	* 2001:418:0:5000::25	252.862 ms	232.198 ms	
4	2001:450:2008:1020::2	235.555 ms	232.634 ms	232.478 ms

Traceroute – how does it work ?

- uses the TTL property of IP packets
 - send the first packet with a TTL of 1, to the destination host.
 - the gateway sees the TTL of 1, decrements it to 0, and returns a “TTL expired” message to the sending host
 - send the second packet, still for the destination host, but this time with a TTL of 2
 - the first gateway lets the packet go through, decrements the TTL from 2 to 1, and passes it on to the next hop
 - the second gateway decrements the TTL from 1 to 0, and returns a TTL expired message to the sending host
 - etc...

netstat

- Allows you to view the status of your network
- The routing table - usage:

```
$ netstat [-n] -r                # v4, v6
$ netstat [-n] -r -f inet        # ipv4
$ netstat [-n] -r -f inet6      # ipv6
```

```
$ netstat -n -r -f inet
Routing tables
```

Internet:

Destination	Gateway	Flags	Refs	Use	Netif
default	196.200.218.254	UGS	1	4610	bge0
127.0.0.1	link#3	UH	0	108	lo0
196.200.218.0/24	link#1	U	0	881	bge0
196.200.218.148	link#1	UHS	0	0	lo0

netstat

- The open connections and listening sockets:

```
$ netstat [-n] -a
```

```
$ netstat -a -n
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address Foreign Address (state)
tcp4 0 0 196.200.218.1.22 196.200.216.49.63843 ESTABLISHED
tcp4 0 0 *.22 *.* LISTEN
tcp6 0 0 *.22 *.* LISTEN
udp4 0 0 *.514 *.*
udp6 0 0 *.514 *.*
Active UNIX domain sockets
Address Type Recv-Q Send-Q Inode Conn Refs Nextref Addr
c55155e8 stream 0 0 c556c770 0 0 0
/tmp/ssh-1To06101I7/agent.983
...
```



ARP

- Used to show IPv4 <-> MAC address lookup tables
- Usually ethernet
- Usage:

```
$ arp -a
```

```
$ arp -a  
? (196.200.218.148) at 00:1e:c9:52:86:be on bge0 permanent  
[ethernet]  
? (196.200.218.254) at 00:25:45:6a:5b:39 on bge0 [ethernet]
```


ARP on v6 ?

- No ARP on v6...
- Use 'ndp'

```
test# ndp -a
Neighbor                               Linklayer Address  Netif  Expire  S  Flags
2001:4348:0:218:196:200:218:1          0:1e:b:b2:f7:e0   em0    20h35m18s  S
2001:4348:0:218:196:200:218:200       0:1e:b:b5:a3:c9   em0    permanent  R
2001:4348:0:218:196:200:218:254       0:1c:58:22:1c:e0  em0    17h15m2s   S R
fe80::21c:58ff:fe22:1ce0%em0          0:1c:58:22:1c:e0  em0    17h14m43s  S R
fe80::21e:bff:feb5:a3c9%em0           0:1e:b:b5:a3:c9   em0    permanent  R
fe80::1%lo0                            (incomplete)      lo0    permanent  R
```

The route command

- The route command is used to modify or query the routing table. Examples for IPv4:

```
route [-n] get default
```

```
route add 196.216.2.34 196.200.218.254
```

```
route add default 196.200.218.253
```

```
route change default 196.200.218.254
```

```
# route get default
  route to: default
destination: default
  mask: default
 gateway: 196.200.218.254
 interface: bge0
  flags: <UP, GATEWAY, DONE, STATIC>
...
```

The route command

- Examples for IPv6

```
route [-n] get -inet6 default
route add 2001:4348:0:223:196:200:223:1
        2001:4348:0:218:196:200:218:254
route add -inet6 default
        2001:4348:0:218:196:200:218:254
route change -inet6 default
        2001:4348:0:218:196:200:218:254
```

```
route -n get -inet6 default
  route to: ::
destination: ::
      mask: default
      gateway: 2001:4348:0:218:196:200:218:254
interface: bge0
      flags: <UP, GATEWAY, DONE, STATIC>
...
```

tcpdump

- tcpdump used to view network traffic on the wire
- basic usage:

```
# tcpdump [-e] [-n] -i if0 [expr.]
```

... where *if0* is your interface (e.g.: bge0)

- To set how much data you want to see, use the 's' option, for example: -s1500
- Expr limits the traffic to certain types (default IPv4)

```
# tcpdump -n -i bge0 icmp
```

```
# tcpdump -n -i bge0 -s1500 tcp and not port
```

```
# tcpdump -n -i bge0 icmp6
```

tcpdump

- Example:

```
# tcpdump -n -i bge0
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on bge0, link-type EN10MB (Ethernet), capture size 96 bytes
20:47:08.700828 IP 196.200.218.148.22 > 196.200.216.43.38378: Flags [P.],
ack 2892018734, win 8326, options [nop,nop,TS val 1949407208 ecr 35128236],
length 192
20:47:08.701938 IP 196.200.216.43.38378 > 196.200.218.148.22: Flags [.] , ac
192, win 999, options [nop,nop,TS val 35128237 ecr 1949407208], length 0
20:47:09.108777 ARP, Request who-has 172.16.0.55 tell 172.16.4.164, length
46
20:47:09.657099 STP 802.1d, Config, Flags [none], bridge-id
80da.00:09:43:a0:79:80.8005, length 43
20:47:09.722271 IP 196.200.218.148.22 > 196.200.216.43.38378: Flags [P.],
ack 1, win 8326, options [nop,nop,TS val 1949408209 ecr 35128237], length
352
```



Tcpdump - IPv6

- Example (with -e to see ethernet addresses)

```
# tcpdump -e -ni bge0 ip6
```

```
19:44:43.434075 00:1e:0b:b2:f7:e0 > 33:33:ff:18:02:00, ethertype IPv6  
(0x86dd), length 86: 2001:4348:0:218:196:200:218:1 > ff02::1:ff18:200:  
ICMP6, neighbor solicitation, who has 2001:4348:0:218:196:200:218:200,  
length 32
```

```
19:44:43.434104 00:1e:0b:b5:a3:c9 > 00:1e:0b:b2:f7:e0, ethertype IPv6  
(0x86dd), length 86: 2001:4348:0:218:196:200:218:200 >  
2001:4348:0:218:196:200:218:1: ICMP6, neighbor advertisement, tgt is  
2001:4348:0:218:196:200:218:200, length 32
```

```
19:44:43.434496 00:1e:0b:b2:f7:e0 > 00:1e:0b:b5:a3:c9, ethertype IPv6  
(0x86dd), length 70: 2001:4348:0:218:196:200:218:1 >  
2001:4348:0:218:196:200:218:200: ICMP6, echo request, seq 0, length 16
```

```
19:44:43.434505 00:1e:0b:b5:a3:c9 > 00:1e:0b:b2:f7:e0, ethertype IPv6  
(0x86dd), length 70: 2001:4348:0:218:196:200:218:200 >  
2001:4348:0:218:196:200:218:1: ICMP6, echo reply, seq 0, length 16
```



mtr

- Can be obtained with `pkg_add -r mtr`
- Combines traceroute & ping – works with v4 & v6

```
# mtr [hostname or IP]
```

```
Keys:  Help      Display mode  Restart statistics  Order of fields  quit
      Packets
      Pings
Host          Loss%  Snt  Last  Avg  Best  Wrst  StDev
1. 2001:4348::216:196:200:217 0.0%   6    2.5   2.7   2.5   2.9   0.2
2. 2001:418:1:101::1          0.0%   6  235.0 235.6 234.0 238.7  1.8
3. fa-4-6.r00.sttlwa01.us.bb 16.7%   6  239.5 239.1 234.8 249.7  6.2
4. 2001:450:2008:1020::2     0.0%   6  235.1 234.9 234.1 235.9  0.6
5. ???
6. ???
7. ???
8. ???
9. ???
10. 2001:41d0:1:2cc8::1       0.0%   5  406.8 405.9 402.0 409.0  2.6
```



Questions ?