

# Exercise V: Analyzing the network traffic

14th October 2003

## Abstract

This work provides basic knowledge of network measurements. It has three parts: use of basic tools, use of data provided by measurement platform and analysing network traces.

## 1 Introduction

This work has two different tasks. In the first one, you will experiment with basic tools studying network performance and its changes. While you are collecting data for the analysis of the first part, you may work for the other analysis task: to analyse data collected by a measurement grid.

## 2 Basic tools

Majority of operating systems provide a set of utility tools to study network status and performance. These provide easy insight to study if network problems are local or if the problem is elsewhere in network.

There are differences between usage of commands in different operating systems and utility versions. Examples here are using program versions listed below on GNU/Linux system (those in classroom). If receive errors or dissimilar function, check usage for version you are using. Examples are designed so that no features special to particular version are used.

**dig** 9.2.2 (prints out in a query)

**ping** netkit-ping 0.10-9 (dpkg -S 'which ping'; dpkg -l netkit-ping)

**traceroute** Version 1.4a12 (version printed out if run without arguments)

**wget** GNU Wget 1.7 (-V)

In examples used here the characters you type are printed in **bold**.

### 2.1 Round-trip latency and connectivity

A very basic test is to check if a host is available. For this purpose, for a part of IP functionality ICMP Echo server was designed. When a host receives ICMP Echo request message, it would respond it by sending ICMP Echo response and copying data part of received message to reply message.<sup>1</sup>

```
prompt> ping -c 3 www.iana.org
PING www.iana.org (192.0.34.162): 56 data bytes
64 bytes from 192.0.34.162: icmp_seq=0 ttl=48 time=190.8 ms
64 bytes from 192.0.34.162: icmp_seq=1 ttl=48 time=191.0 ms
64 bytes from 192.0.34.162: icmp_seq=2 ttl=48 time=190.5 ms

--- www.iana.org ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 190.5/190.7/191.0 ms
```

In example above, `ping` command was instructed to send three ICMP Echo requests to host `www.iana.org`. A line is printed out for each response received including sequence number, time-to-live field value<sup>2</sup>, and time between sent message and reply.

For statistics number of packets transmitted and received with loss percentage and total time for command is printed out. For round-trip time minimum, average, and maximum are calculated.

<sup>1</sup>This helps for detection of data-related errors in transmission.

<sup>2</sup>Can be used to estimate number of routers between as initial values for TTL are 64, 255, 128, or 32 commonly.

### 2.1.1 Nobody home but services still working

For (pseudo-)security reasons many hosts and sites use firewall rules to drop ICMP messages<sup>3</sup> and then one cannot use ICMP echos to test. It is possible to check connectivity for ICMP-blocking servers by contacting the service the server is providing.

An examples of such site one of root servers, `e.root-servers.net`. As shown below, we do not receive any response to ICMP messages but when we query for name servers serving `fi-ccTLD` we got reply in 110 ms. This is quite close to last time provided by traceroute. While one must account for some delays in end systems even with ICMP messages, this is especially important for application tests. However, because root servers are serving only name service info for top-level domains (gTLD like `com`, `net` and ccTLD like `fi`, `fr`, `cn`) but know nothing about particular domains like `example.com`, the response time should be comparable to one of ICMP messages.

```
prompt> ping -q -c 3 e.root-servers.net
PING e.root-servers.net (192.203.230.10): 56 data bytes

--- a.root-servers.net ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss
prompt> traceroute -f 15 e.root-servers.net
traceroute to e.root-servers.net (192.203.230.10), 30 hops max, 40 byte packets
15  so-4-0-0.gar1.SanJose1.level3.net (4.68.112.58) 192.494 ms 192.383 ms 264.260 ms
16  so-10-0.hsa3.SanJose1.Level3.net (209.244.13.226) 194.374 ms 194.452 ms 194.962 ms
17  nasa-level3-fastethernet.SanJose1.Level3.net (209.245.146.6) 194.755 ms 195.517 ms 194.991 ms
18  128.161.3.92 (128.161.3.92) 228.506 ms 260.331 ms 390.873 ms
19  * * *
20  * * *
21  * * *
prompt> dig @e.root-servers.net fi. ns

; <<>> DiG 9.2.1 <<>> @e.root-servers.net fi. ns
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 58797
;; flags: qr rd; QUERY: 1, ANSWER: 6, AUTHORITY: 0, ADDITIONAL: 6

;; QUESTION SECTION:
;fi.                                IN      NS

;; ANSWER SECTION:
fi.                                172800  IN      NS      T.NS.VERIO.NET.
fi.                                172800  IN      NS      NS1-FIN.GLOBAL.SONERA.fi.
fi.                                172800  IN      NS      NS-SE.ELISA.NET.
fi.                                172800  IN      NS      PRIFI.FICORA.fi.
fi.                                172800  IN      NS      HYDRA.HELSINKI.fi.
fi.                                172800  IN      NS      NS.UU.NET.

;; ADDITIONAL SECTION:
T.NS.VERIO.NET.                    172800  IN      A       192.67.14.16
NS1-FIN.GLOBAL.SONERA.fi.          172800  IN      A       193.210.18.31
NS-SE.ELISA.NET.                   172800  IN      A       62.248.254.2
PRIFI.FICORA.fi.                   172800  IN      A       193.229.4.44
HYDRA.HELSINKI.fi.                 172800  IN      A       128.214.4.29
NS.UU.NET.                          172800  IN      A       137.39.1.3

;; Query time: 293 msec
;; SERVER: 192.203.230.10#53(e.root-servers.net)
;; WHEN: Mon Oct 13 21:06:50 2003
;; MSG SIZE rcvd: 282
```

For other services, such as `http`, there is more processing at end system needed so response time may be longer and more variable.

You can use `traceroute` to check how long is the common path to two destination. For example running two traceroutes below on `t.ns.verio.net` and `ns.uu.net` one can notice that up to 16th router the routes are same.

---

<sup>3</sup>Either all message types that results severe hard-to diagnose connectivity problems or just “useless” types (including `echo`).

```

prompt> traceroute T.NS.VERIO.NET
removed lines..
14  sl-bb21-msq-10-0.sprintlink.net (144.232.19.29) 100.766 ms
15  sl-bb20-msq-15-0.sprintlink.net (144.232.9.109) 100.439 ms
16  sl-bb20-nyc-11-3.sprintlink.net (144.232.9.102) 101.779 ms
17  sl-gw39-nyc-1-0.sprintlink.net (144.232.13.62) 101.394 ms
removed lines..
prompt> traceroute ns.uu.net
removed lines..
14  sl-bb21-msq-10-0.sprintlink.net (144.232.19.29) 100.850 ms
15  sl-bb20-msq-15-0.sprintlink.net (144.232.9.109) 101.626 ms
16  sl-bb20-nyc-11-3.sprintlink.net (144.232.9.102) 101.845 ms
17  204.255.174.225 (204.255.174.225) 105.584 ms 105.624 ms
removed lines..

```

### 3 Passive measurements

Passive measurements can be often very complicated, especially if you want to make in-depth flow-level measurements real-time. However, we do only basic measurements using simple text-based tools.

#### 3.1 Tcpcmdump

The very basic tool for network traces is `tcpdump` that is available for multiple operating system. For simple analysis, you can work with text-based `tcpdump` output.

Short summary of important options:

- n do not resolve addresses (should be used here everytime as packet trace has addresses changed and thus names are totally useless and non-informative).
- q provides more condensed output. It is useful here as it prints out tcp and udp user data lengths.
- v provides more information about packets, no useful here as most of information is removed.
- t removes timestamp from lines
- tt prints timestamp as seconds.microsecond. Usefull if you want to do some time-based calculations
- r file read packets from file instead of network device

`Tcpdump` has an expression language to select packets from network trace. For example expression `udp port 53` selects those packets that have UDP protocol and either source or destination port is 53 (name service). Expression `tcp src port 80 or tcp src port 8080 or tcp src port 8888` selects those tcp packets that have one of typical web server port numbers as source port.

Text output is easy to analyse using `awk` program. With following line (combine to one line) one can find out number of bytes in tcp traffic, number of packets and average packet size.

```

/usr/sbin/tcpdump -r /p/edu/s-38.180/trace/10-minute.pcap -n -q tcp
| awk '{i+=$6+40;c++}END{print i, c, i/c}'

```

#### 3.2 Timing commands

If you do not want to sit beside computer to run commands at specified times, you can utilise `at` command. Easiest way is to create command script that should be run at specified times:

```

date >> log
ping -c 4 m.root-servers.net >> log
dig @m.root-servers.net fi. ns | grep time: >> log

```

If you save framgment above file as `ping-dig`, you can then give following commands to run script after one, two, three, and four hours:

```

prompt> at -f ping-dig now + 1 hour
prompt> at -f ping-dig now + 2 hour
prompt> at -f ping-dig now + 3 hour
prompt> at -f ping-dig now + 4 hour

```

Output is recoded to file `log`. If there are errors, they will be mailed to your account.

## 4 Tasks

Answers should be returned by Friday 2003-11-03.

### 4.1 Check latency for some ccTLD name servers

To randomly distribute ccTLDs, there is a program at `/p/edu/s-38.180/trace/cc` you need to run to find out ccTLD assigned to you.<sup>4</sup> Compare rtt times measured using ICMP messages and DNS queries. If some of those does not respond to ICMP messages, use traceroute to locate last responding router and compare ping rtt to DNS queries (select some domain you know and ask its name servers).

Select three of those name servers and check round-trip time with dns query (as above). Repeat test at least three times with more than one hour intervals. Make some conclusions about stability of network delay. Was some of those servers different from the others?

### 4.2 Measurement platform studies

Study AMP measurements `http://amp.nlanr.net/active/` for RTT and loss measurements between amp-hutf and amp-surf at 2003-10-06<sup>5</sup>. What anomalies you can identify from that day and how it affects on network performance? Is the reverse path similar?

### 4.3 Traffic analysis

There is an anonymised network trace (about 10 minute long, 20000 packets) in `/p/edu/s-38.180/trace/10-minute.pcap`. Based on that data, answer following:

1. What is proportion of each transport-layer protocol (tcp, udp, icmp, igmp) by packets and by bytes.
2. Select 5 different TCP application (well-known port numbers) and report proportion for each of those<sup>6</sup>
3. Report top-10 hosts by packet count (senders) and by bytes
4. How much there are ACK-only TCP segments.

Remember to use `-n` switch for `tcpdump` to avoid unnecessary (and false, as addresses have been modified) DNS lookups!

---

<sup>4</sup>The program checks that there is at least 5 name servers for given ccTLD and tries another one until it finds one with enough name servers.

<sup>5</sup>`http://watt.nlanr.net/active/cgi-bin/weekly.cgi?amp-hutf/international/data/amp-surf`

<sup>6</sup>You can check from `/etc/services` for applications corresponding to port numbers.