Quantum

Alain Renaud
Sustaining engineering

May 2013

Diagnostic with Wireshark.

- Ethernet Frame
- TCP Header
- Using tcpdump.
- Known Stornext ports.

Ethernet Frame.

- Each Ethernet frame size is of MTU + 26 bytes maximum.
- Doing a tcpdump of 96Bytes usually get all the header information.

TCP Header

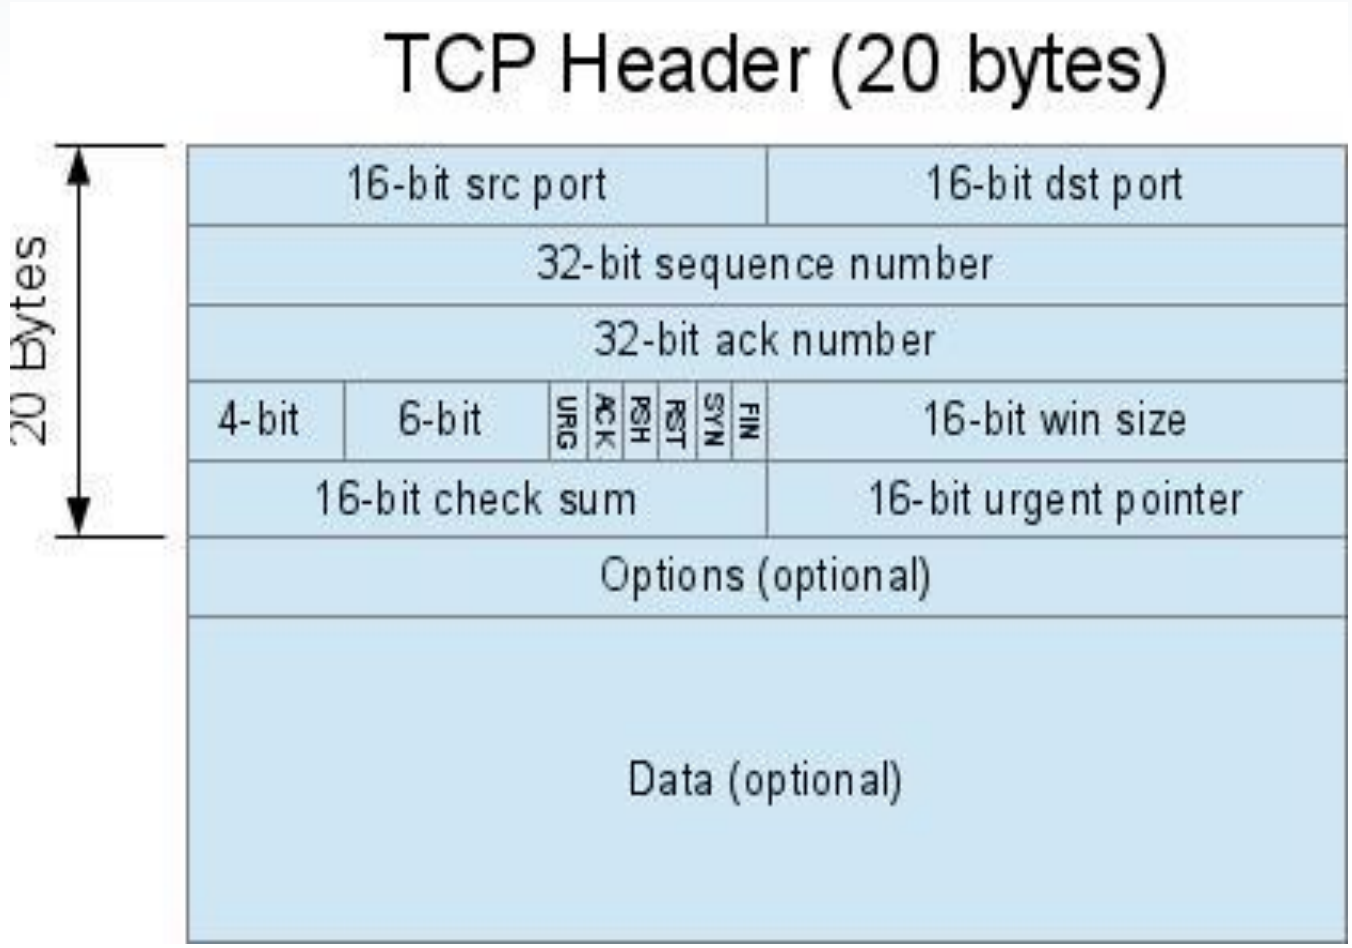
TCP Header (20 bytes)
20 Bytes
16-bit src port
16-bit dst port
32-bit sequence number
32-bit ack number
4-bit
6-bit
URG
ACK
PSH
RST
SYN
FIN
16-bit win size
16-bit check sum
16-bit urgent pointer
Options (optional)
Data (optional)

TCP Header (Cont)

- A TCP connection consist of two data stream.
- The windows size is 16-bit which is limited to 64k
- The ACK return the next expected sequence number.
- Some of the more common TCP options are:
  - window scale: multiplier for windows size.
  - MSS: Maximum segment size. MTU discovery.(default 576)
  - SACK: Selective ACK.

Using tcpdump

- Tcpdump is the preferred tool to grab network packet on Linux and MacOSX. Tcpdump is also available on Windows but people tend to use Wireshark directly to grab packet.
- Exemples:
  - Basic usage ‘-i’ to specify the interface and ‘-nn’ to avoid name resolution for IP and PORT.

```
# tcpdump -i eth0  -nn
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
10:05:13.771021 ARP, Request who-has 10.65.162.31 tell 10.65.162.36, length 46
10:05:13.771723 IP 10.65.178.124.35243 > 10.65.162.11.111: Flags [S], seq 2273551952, win 14600, options [mss 1460,sackOK,TS val 1303682767 ecr 0,nop,wscale 7], length 0
```

Using tcpdump (cont)

- Usually we want to save the dump data to a file. We then use the '-w' flag and also use '-s' flag to avoid grabbing all the data and having a dump file to big.

```
# tcpdump -i eth0 -nn -s 96 -w data.pcap
tcpdump: listening on eth0, link-type EN10MB (Ethernet), capture size 96 bytes
CTRL-C
323 packets captured
323 packets received by filter
0 packets dropped by kernel
```

Using tcpdump (cont)

- Finally we can filter out exactly what we want to grab. This also allow to reduce the size of the dump that we grab and avoid to saturate the disk.
  - By host

```
# tcpdump -i eth0 -nn -s 96 -w data.pcap host 10.65.179.99
```

  - By port

```
# tcpdump -i eth0 -nn -s 96 -w data.pcap port 5164
```

  - Both

```
# tcpdump -i eth0 -nn -s 96 -w data.pcap host 10.65.179.99 and port 5164
```

Known SNFS ports.

- When the ‘fsmpm’ start it will always bind to TCP port 5164 unless told otherwise. That is the only constant port in the protocol. Client at startup will connect to that port on the fsmpm and ask for the information about the cluster.
- The ‘fsmpm’ also maintain a heartbeat over UDP to find this port you can look at the nssdbg.out log file or simply do:

```
# netstat -anp | grep fsmpm
tcp        0      0 :::5164                 :::*                    LISTEN      1798/fsmpm
tcp        0      0 :::43930                :::*                    LISTEN      1798/fsmpm
tcp        0      0 ::ffff:127.0.0.1:43930  ::ffff:127.0.0.1:43726  ESTABLISHED 1798/fsmpm
udp        0      0 :::57645                :::*                                1798/fsmpm
unix  2      [ ]         DGRAM                    11397  1798/fsmpm
unix  2      [ ]         DGRAM                    11385  1797/fsmpm-watcher
```

Known SNFS ports. (cont)

- Finally the ‘fsmpm’ create a ‘*portmap’* which is use for client application to talk to the fsmpm (talk on loopback interface)

- Each ‘fsm’ on the MDC also has a port for each filesystem that it maintain. You can find these port with ‘cvadmin’ command. All communication between clients and MDC for a specific filesystem go ever this port.

Known SNFS ports. (cont)

- Data Lan Server also have a port to transfer data with the the DLC ‘cvdb –x’ will tell you which port. Note the server can listen to multiple port one per defined interface in /usr/cvfs/config/dpserver

```
# cvdb -x | grep 'IP Addr'
  Client   IP Addr  10.0.16.80 (36913)
  Server   IP Addr  10.0.16.245 (36643)
  Client   IP Addr  10.65.176.80 (50220)
  Server   IP Addr  10.65.178.245 (48620)
  Client   IP Addr  10.65.179.192 (52201)
  Server   IP Addr  10.65.178.245 (48620)
```

Wireshark.

References

- http://www.netperf.org/netperf/
- http://code.google.com/p/netperf-win/
- http://sourceforge.net/projects/iperf/

BE CERTAIN