Wireshark报文分析

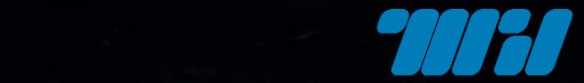

北京网瑞达科技有限公司
Beijing WRD Technology Co., Ltd.

Linux Performance Observability Tools

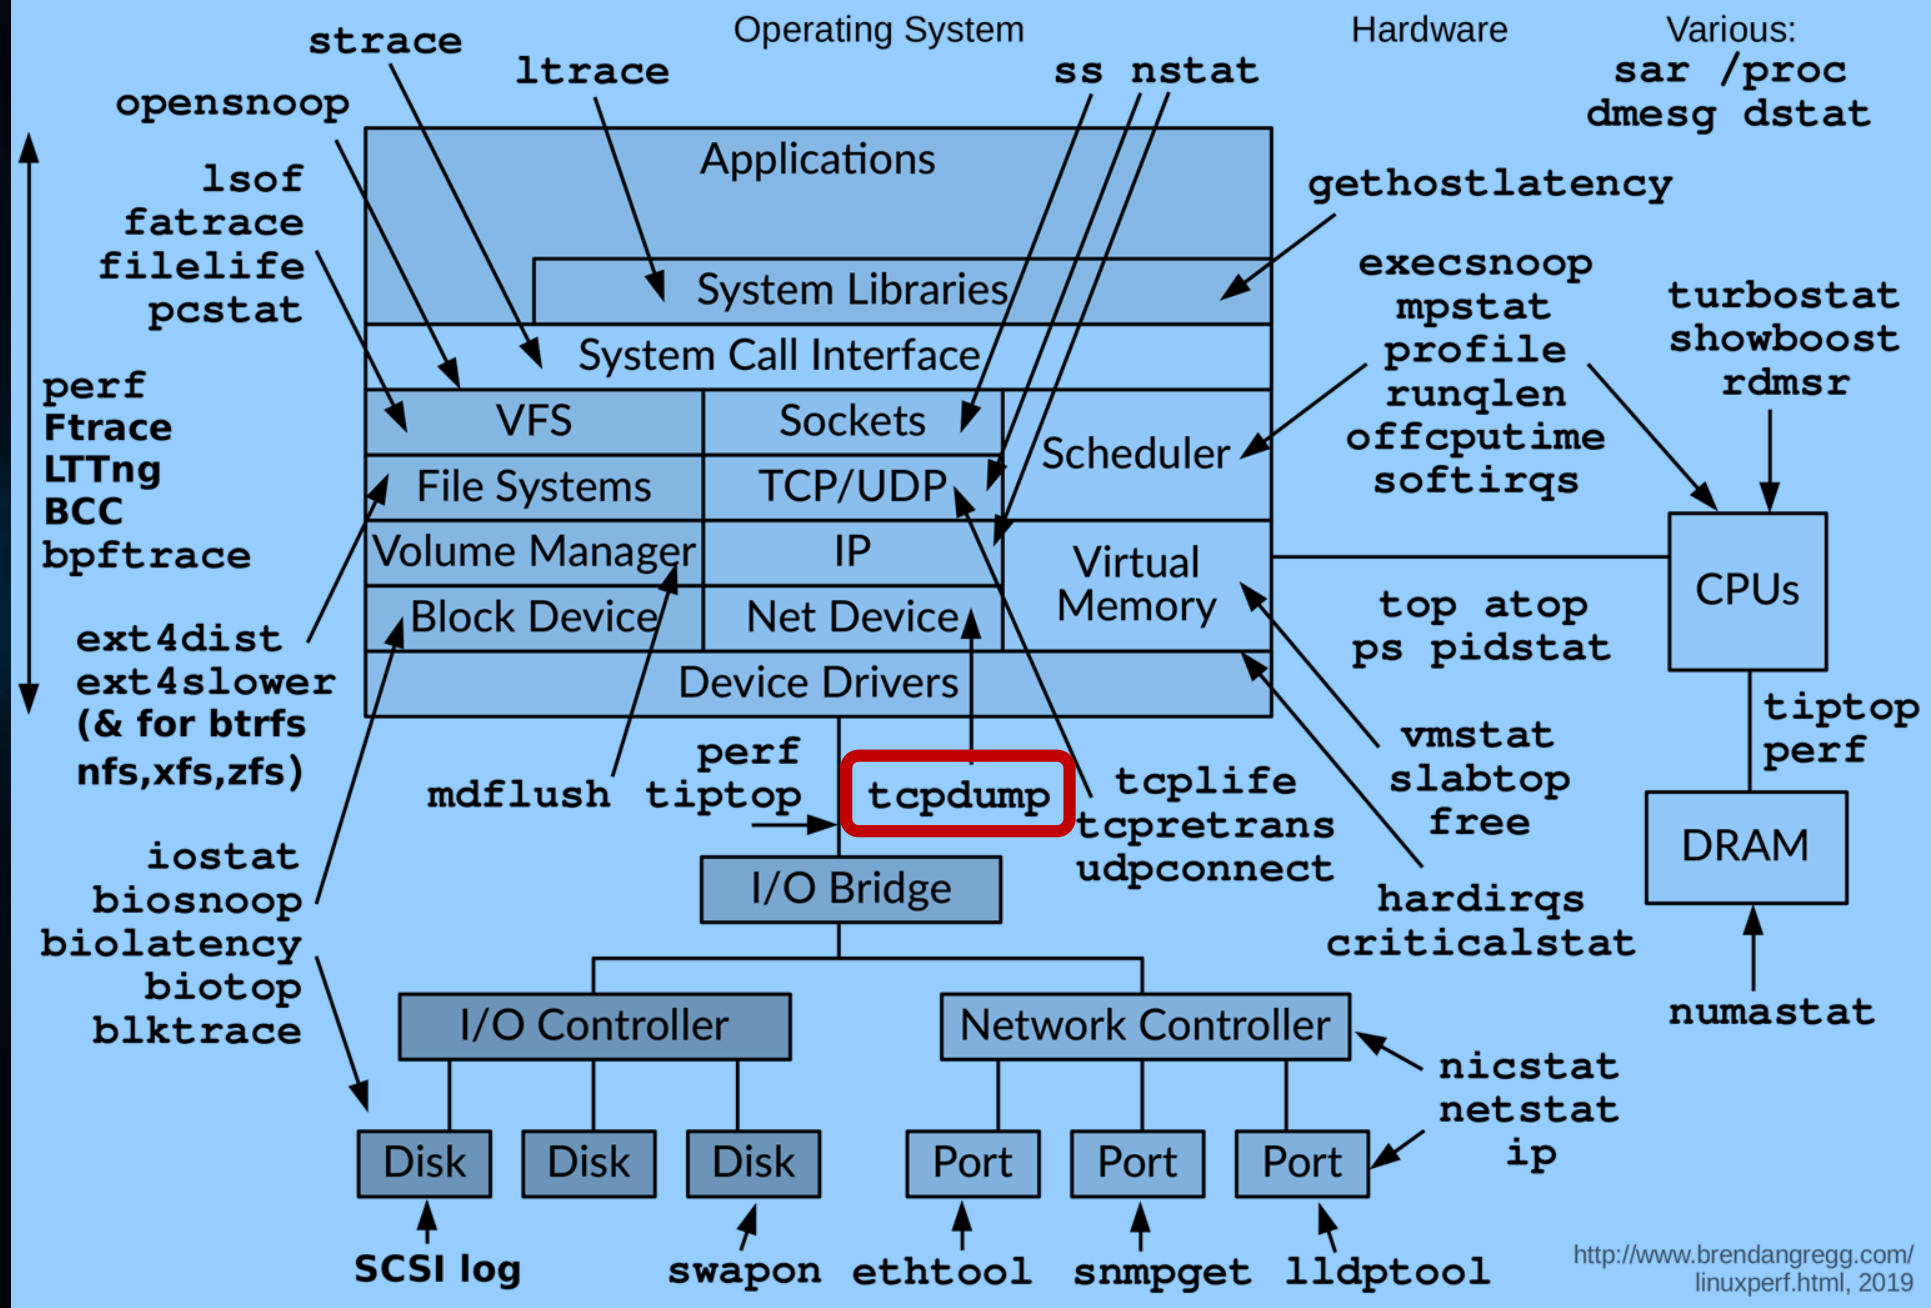

http://www.brendangregg.com/linuxperf.html, 2019

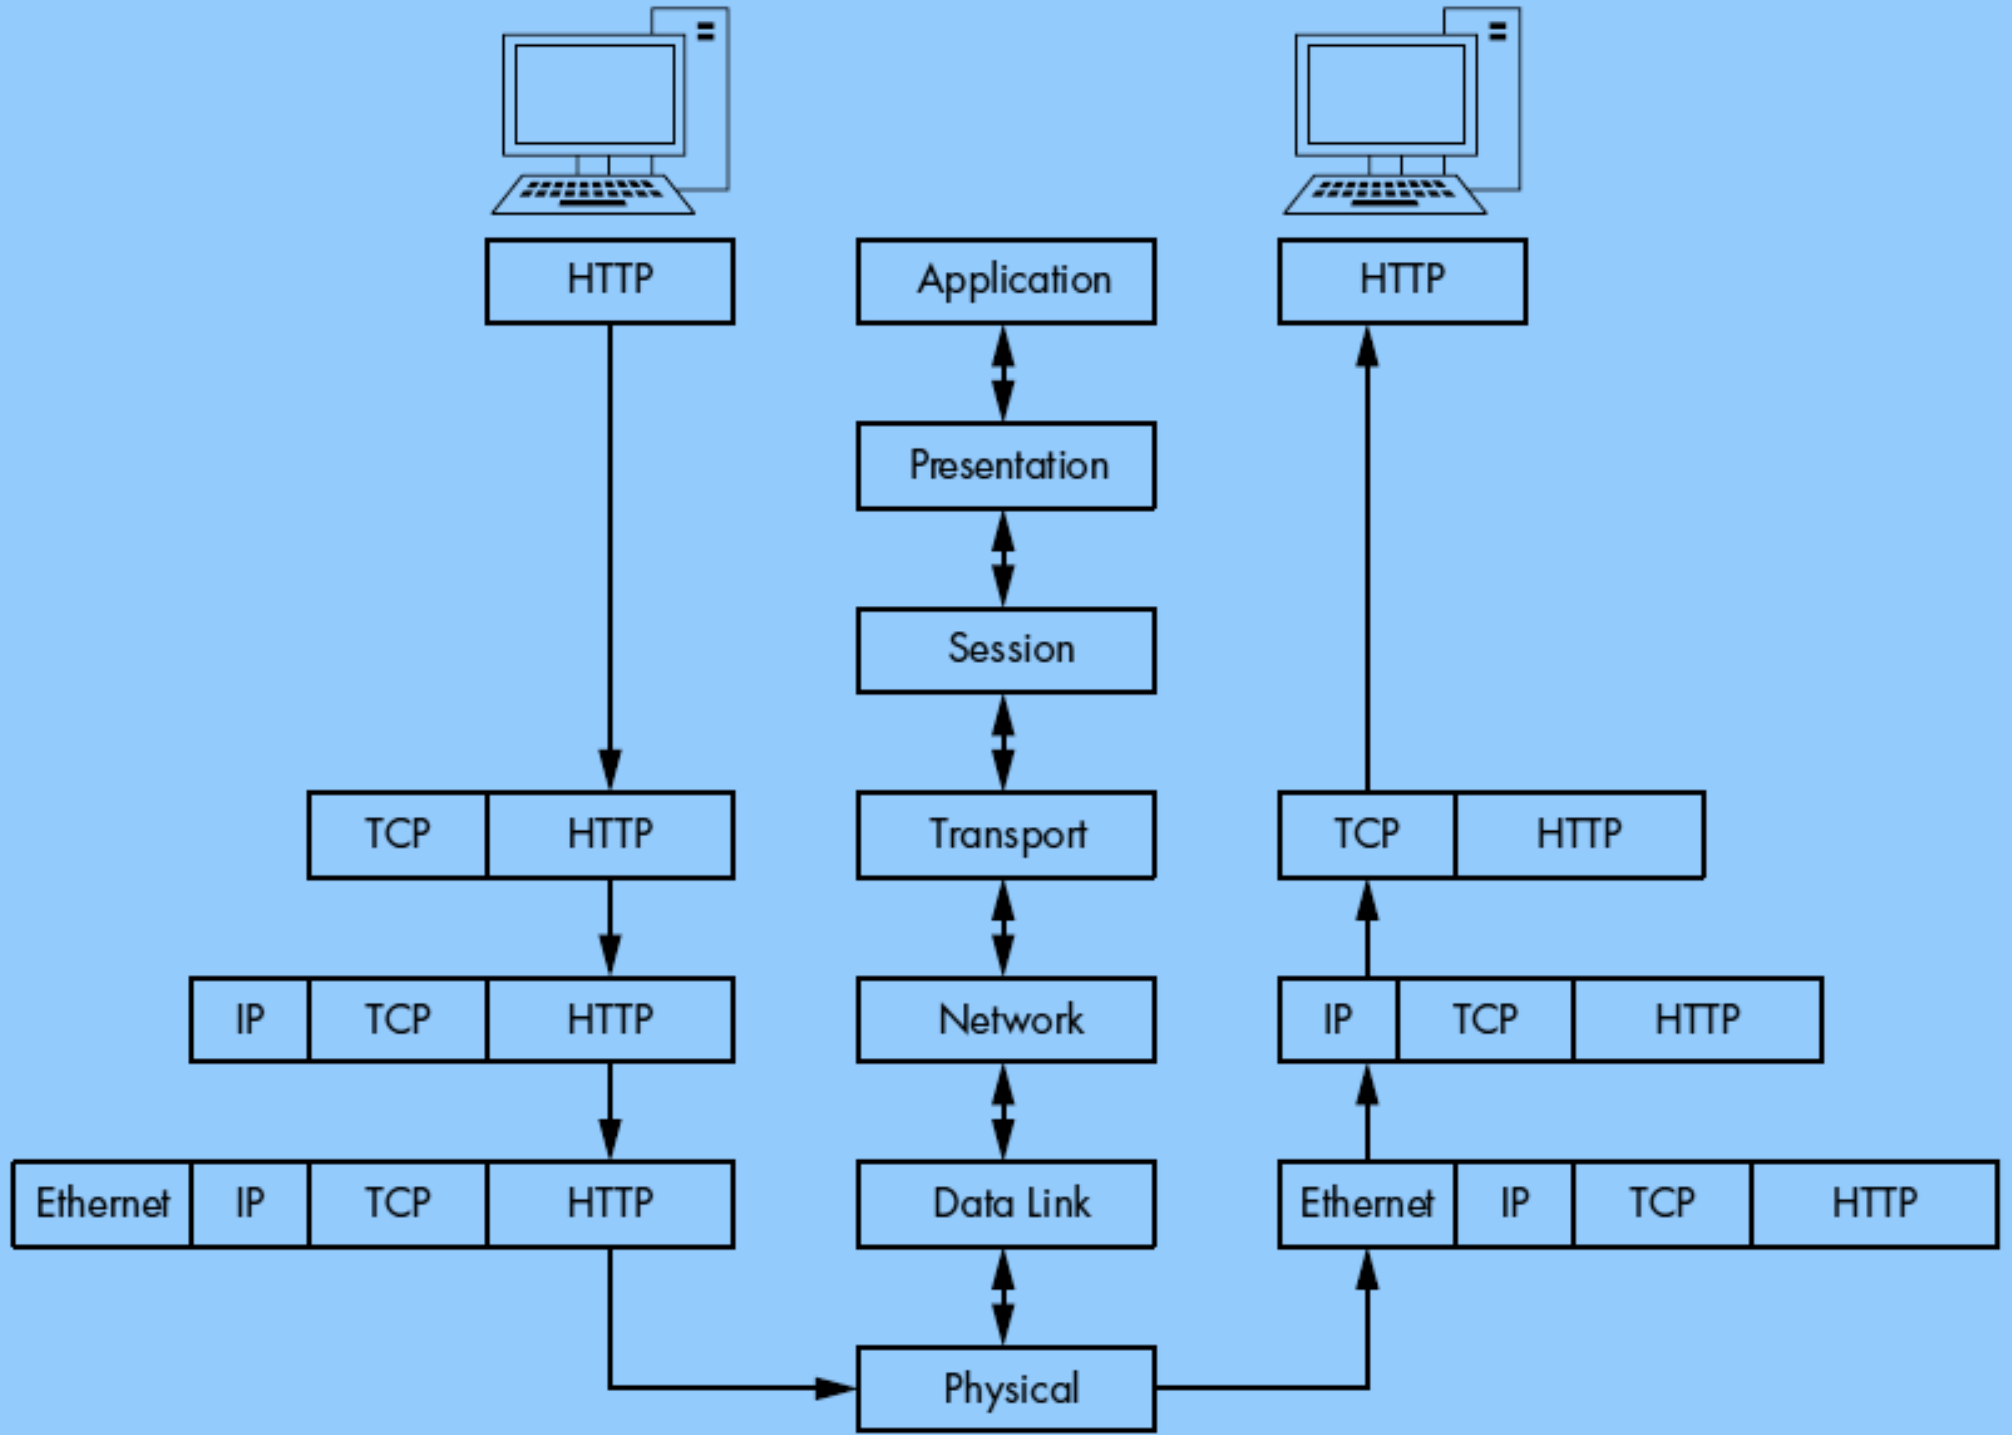
HTTP
Application
HTTP
Presentation
Session
TCP
HTTP
Transport
TCP
HTTP
IP
TCP
HTTP
Network
IP
TCP
HTTP
Ethernet
IP
TCP
HTTP
Data Link
Ethernet
IP
TCP
HTTP
Physical

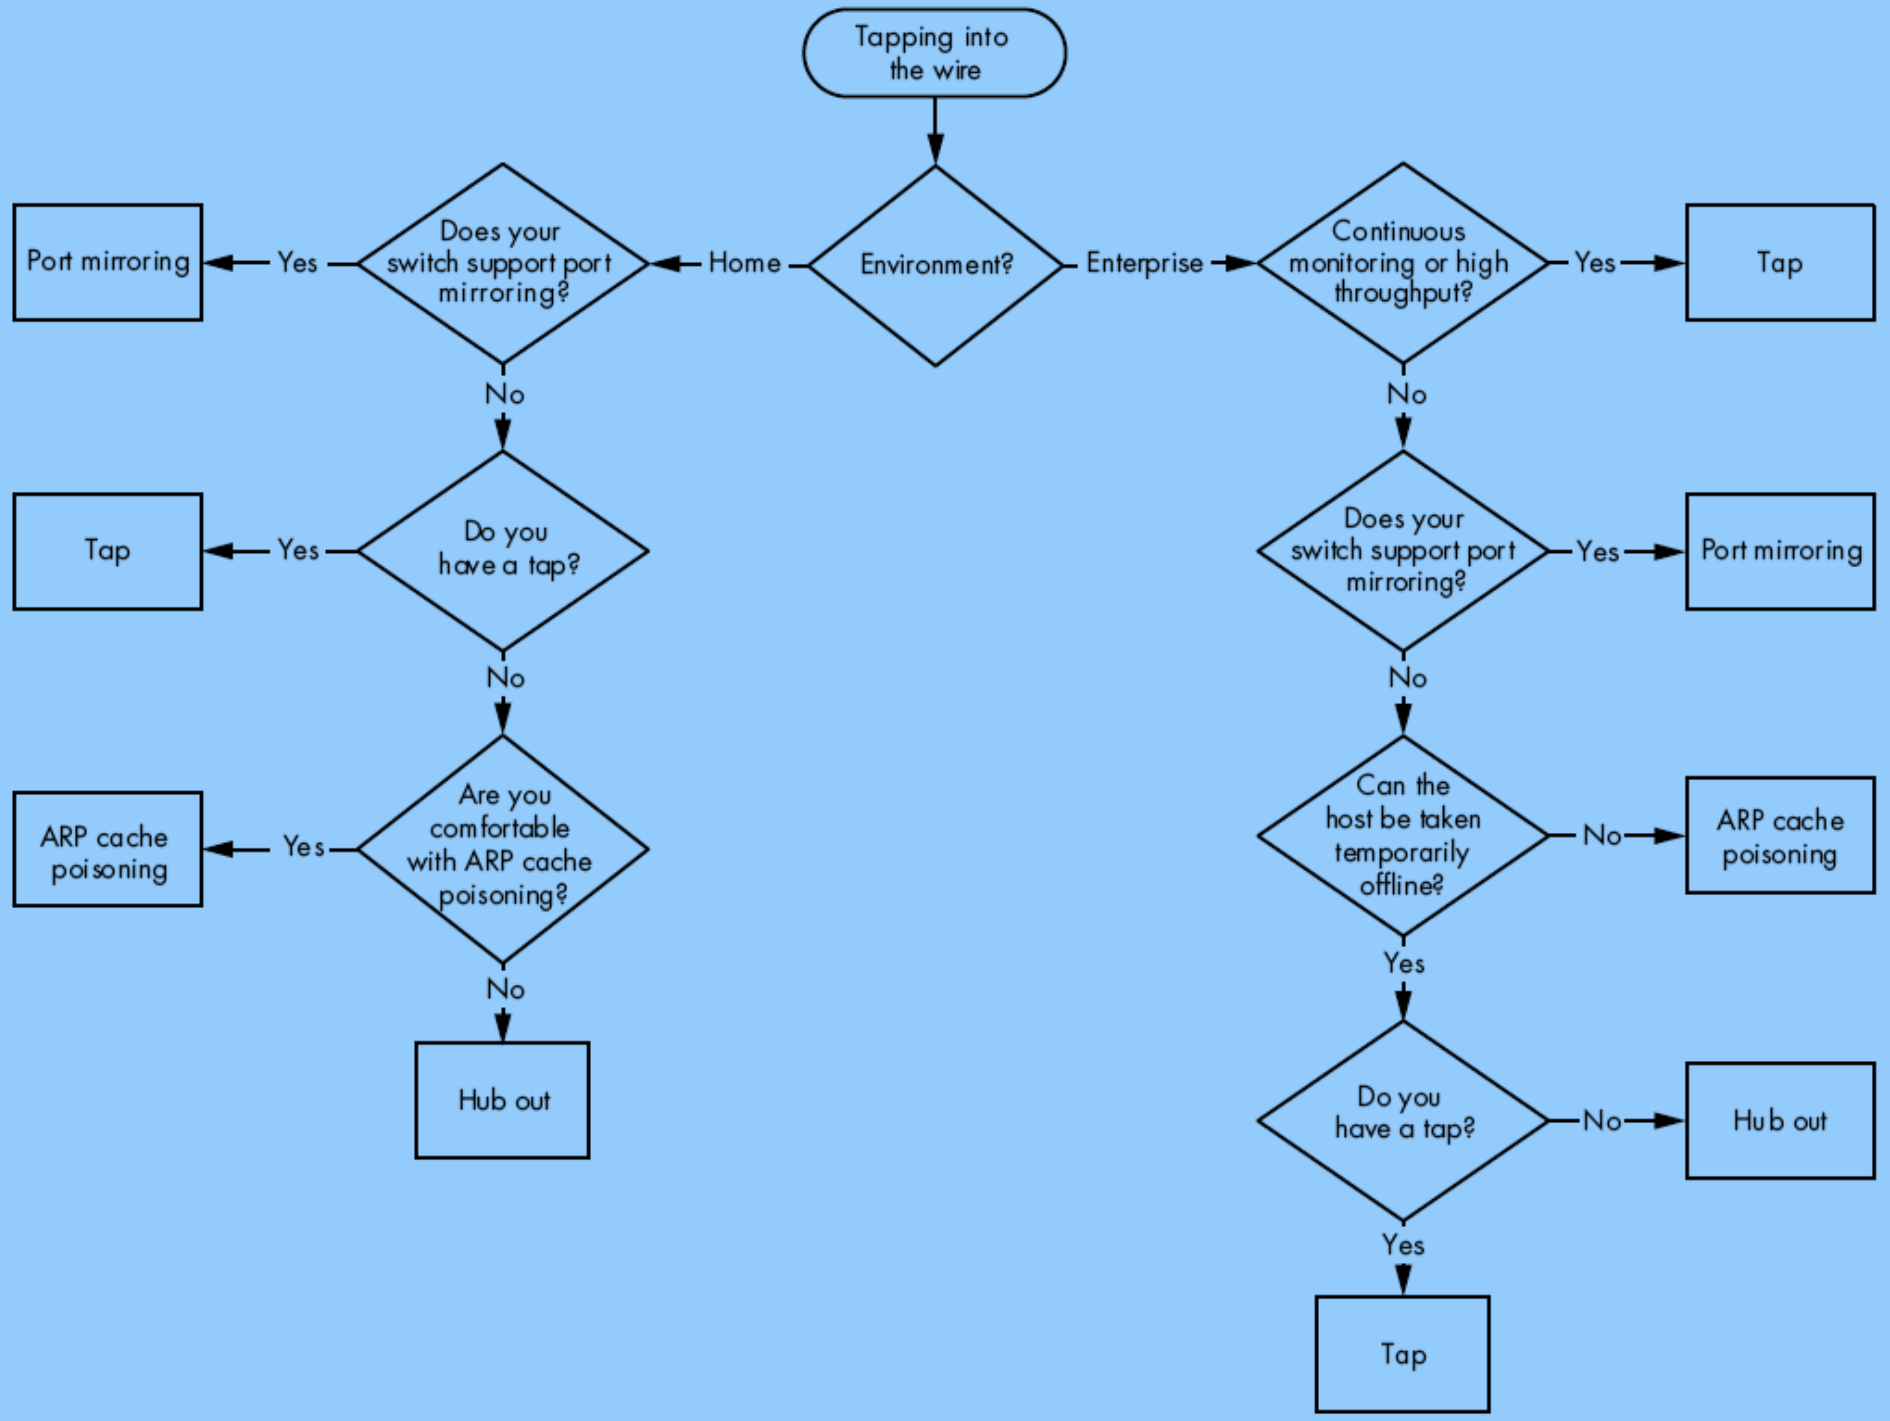
Tapping into the wire
Environment?
Home
Enterprise
Does your switch support port mirroring?
Yes
Port mirroring
No
Do you have a tap?
Yes
Tap
No
Are you comfortable with ARP cache poisoning?
Yes
ARP cache poisoning
No
Hub out
Continuous monitoring or high throughput?
Yes
Tap
No
Does your switch support port mirroring?
Yes
Port mirroring
No
Can the host be taken temporarily offline?
No
ARP cache poisoning
Yes
Do you have a tap?
No
Hub out
Yes
Tap

嗅探点


tcpdump
p
Server A
tcpdump
p
tcpdump
p
Server B
tcpdump
p
L2/L3 交换机
核心交换机
L2交换机
Client
网络A区
网络B区
网络C区
网络D区
网络E区

步骤主线

tcpdump

- tcpdump 是一个运行在命令行的嗅探工具
- 支持BPF过滤规则
  - 借助对网络层、协议、主机、网络或端口的过滤，并提供and、or、not等逻辑语句来去掉无用的信息

tcpdump

```
[ycflash@kvm7 tmp]$ sudo tcpdump -i ens6f0 -s0
tcpdump: listening on ens6f0, link-type EN10MB
^C411 packets captured
412 packets received by filter
0 packets dropped by kernel
[ycflash@kvm7 tmp]$
```

- 常用选项：
  - -i 指定网卡接口，any 表示所有接口
  - -s 指定嗅探报文长度
  - -n,-nn 禁止将字段解析为主机名、协议端口名
  - -e 打印二层MAC地址
  - -w 将捕获协议报文记录至文件
  - -c 捕获N个报文即停止
  - -v,-vv,-vvv 显示详细协议报文解析输出

例：tcpdump -i eth0 -s0 -w /tmp/1.pcap udp port 67 or udp port 68

BPF 过滤规则

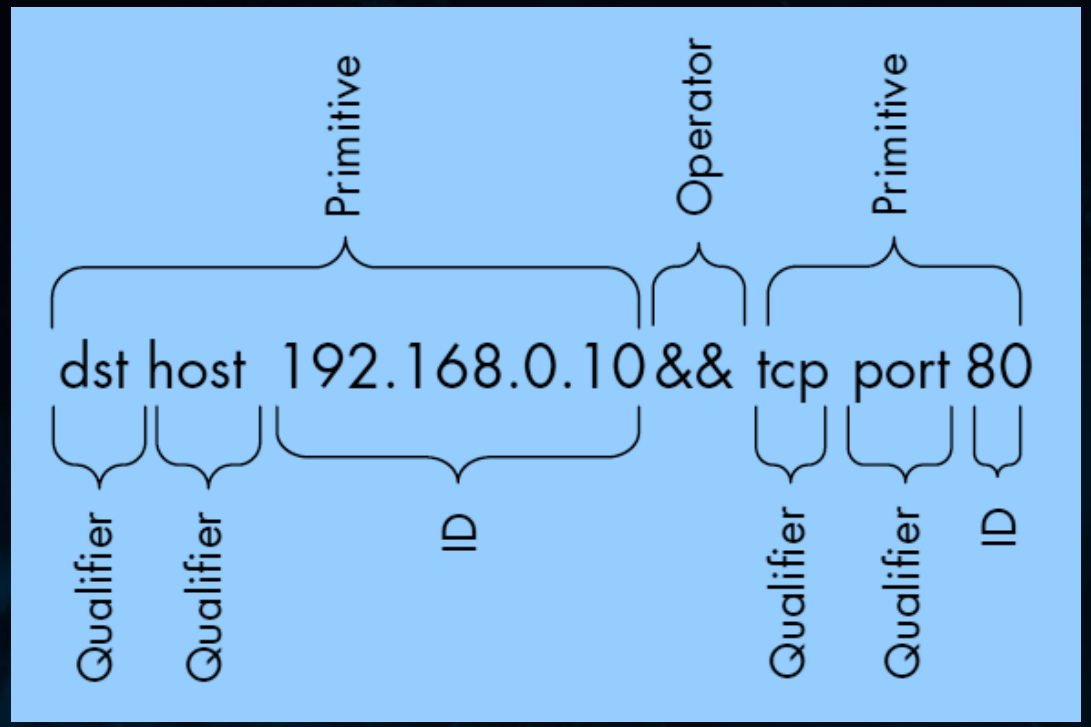

- 过滤规则语法
  - 原语
    - 类型
      - host, net, port, portrange
    - 方向
      - src, dst, src or dst, src and dst
    - 协议
      - ether, fddi, tr, wlan, ip, ip6, arp, rarp, decnet, tcp, udp
  - 逻辑操作
    - and, or, not 或 &&, ||, !
    - (,)

https://www.tcpdump.org/manpages/pcap-filter.7.html

BPF 过滤规则

- 示例
  - dst host <host>
  - src host <host>
  - dst net <net>
  - dst port <port>
  - udp port 67 or udp port 68
  - tcp port 23 and host 10.0.0.5

BPF 过滤规则

- 复杂表达式
  expr relop expr
  - 算术支持
    >, <, >=, <=, =, !=
  - 位操作
    +, -, *, /, %, &, |, ^, <<, >>
- proto 语法
  proto [*expr* : *size*]
  - proto：
    - ether, fddi, tr, wlan, ppp, slip, link, ip, arp, rarp, tcp, udp, icmp, ip6 or radio
  - expr：
    - 相对 proto 指定协议起始位置的偏移量offset, 从0起始
  - size：
    - 取值大小，1、2、4字节，默认1字节

BPF 过滤规则

- 示例
  - ether[0] & 1 != 0
  - ip[0] & 0xf != 5
  - ip[6:2] & 0x1fff = 0
  - tcp[tcpflags] & (tcp-rst|tcp-ack) == (tcp-rst|tcp-ack)

  - ((port 67 or port 68) and (udp[38:4] = 0x3e0ccf08))
  - ((port 67 or port 68) and (udp[32:4] = 0x0ACF0001))

wireshark

- Wireshark前身为Ethereal
- 免费开源
- 提供tcpdump没有的GUI
- 兼容PCAP记录文件格式

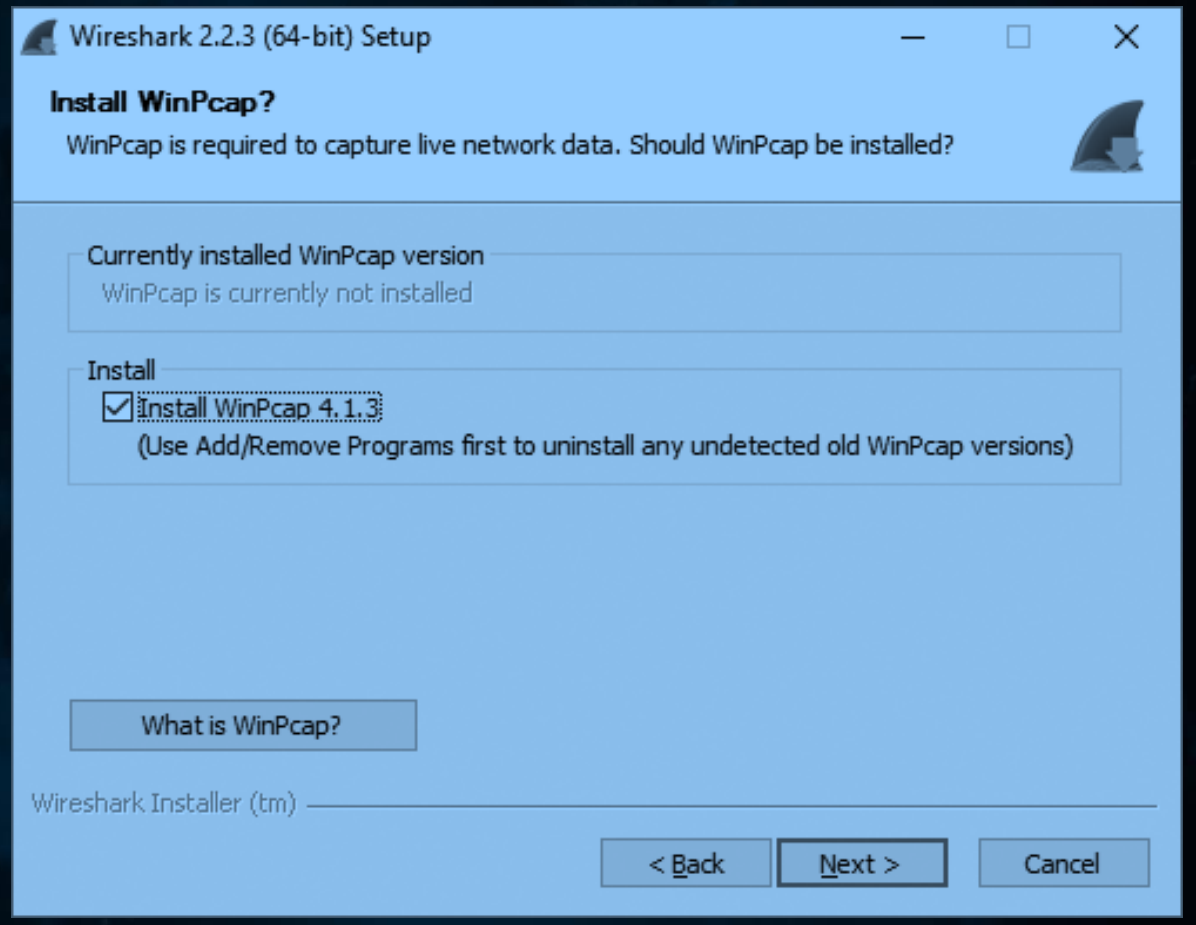

https://www.wireshark.org/#download

Interface	AIX	FreeBSD	HP-UX	Irix	Linux	macOS	NetBSD	OpenBSD	Solaris	Tru64 UNIX	Windows
ATM	?	?	?	?	✓	✗	?	?	✓	?	?
Bluetooth	✗	✗	✗	✗	✓[1]	✗	✗	✗	✗	✗	✗
CiscoHDLC	?	✓	?	?	✓	?	✓	✓	?	?	?
Ethernet	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
FDDI	?	?	?	?	✓	✗	?	?	✓	?	?
FrameRelay	?	?	✗	✗	✓	✗	?	?	✗	✗	✗
IrDA	✗	✗	✗	✗	✓	✗	✗	✗	✗	✗	✗
PPP[2]	?	?	?	?	✓	✓	?	?	✗	?	✓
TokenRing	✓	✓	?	✗	✓	✗	✓	✓	✓	?	✓
USB	✗	✗	✗	✗	✓[3]	✗	✗	✗	✗	✗	✗
WLAN[4]	?	✓	?	?	✓	✓	✓	✓	?	?	✓
Loopback (virtual)	?	✓	✗	?	✓	✓	✓	✓	✗	✓	N/A[5]
VLAN Tags (virtual)	✓	✓	✓	?	✓	✓	✓	✓	✓	✓	✓

https://gitlab.com/wireshark/wireshark/-/wikis/CaptureSetup/NetworkMedia

主界面

Ctrl+. 移动到TCP、UDP 会话下一包

Ctrl+, 移动到 TCP、UDP 会话上一包

报文时间戳显示调整：
View→Time Display format

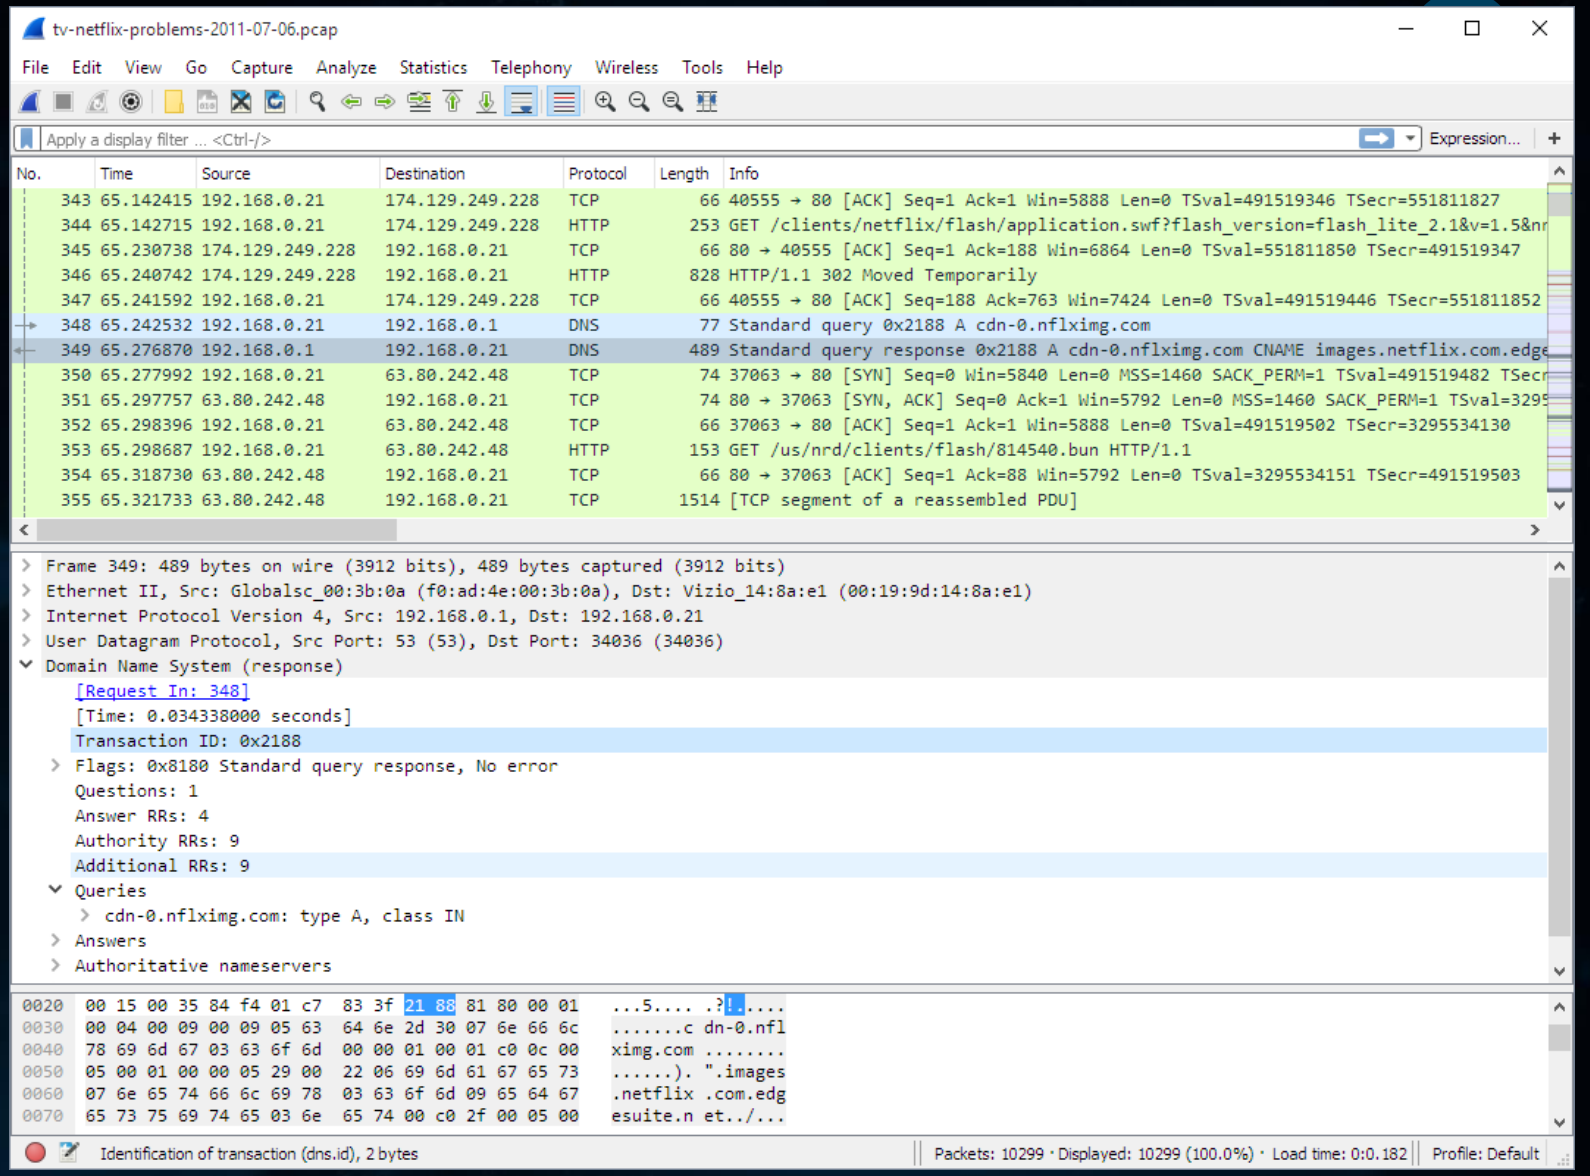

报文详情窗格

- 中括号标注的是分析用字段，非实际报文中的字段
- 如有链接，点击可跳转到关联报文

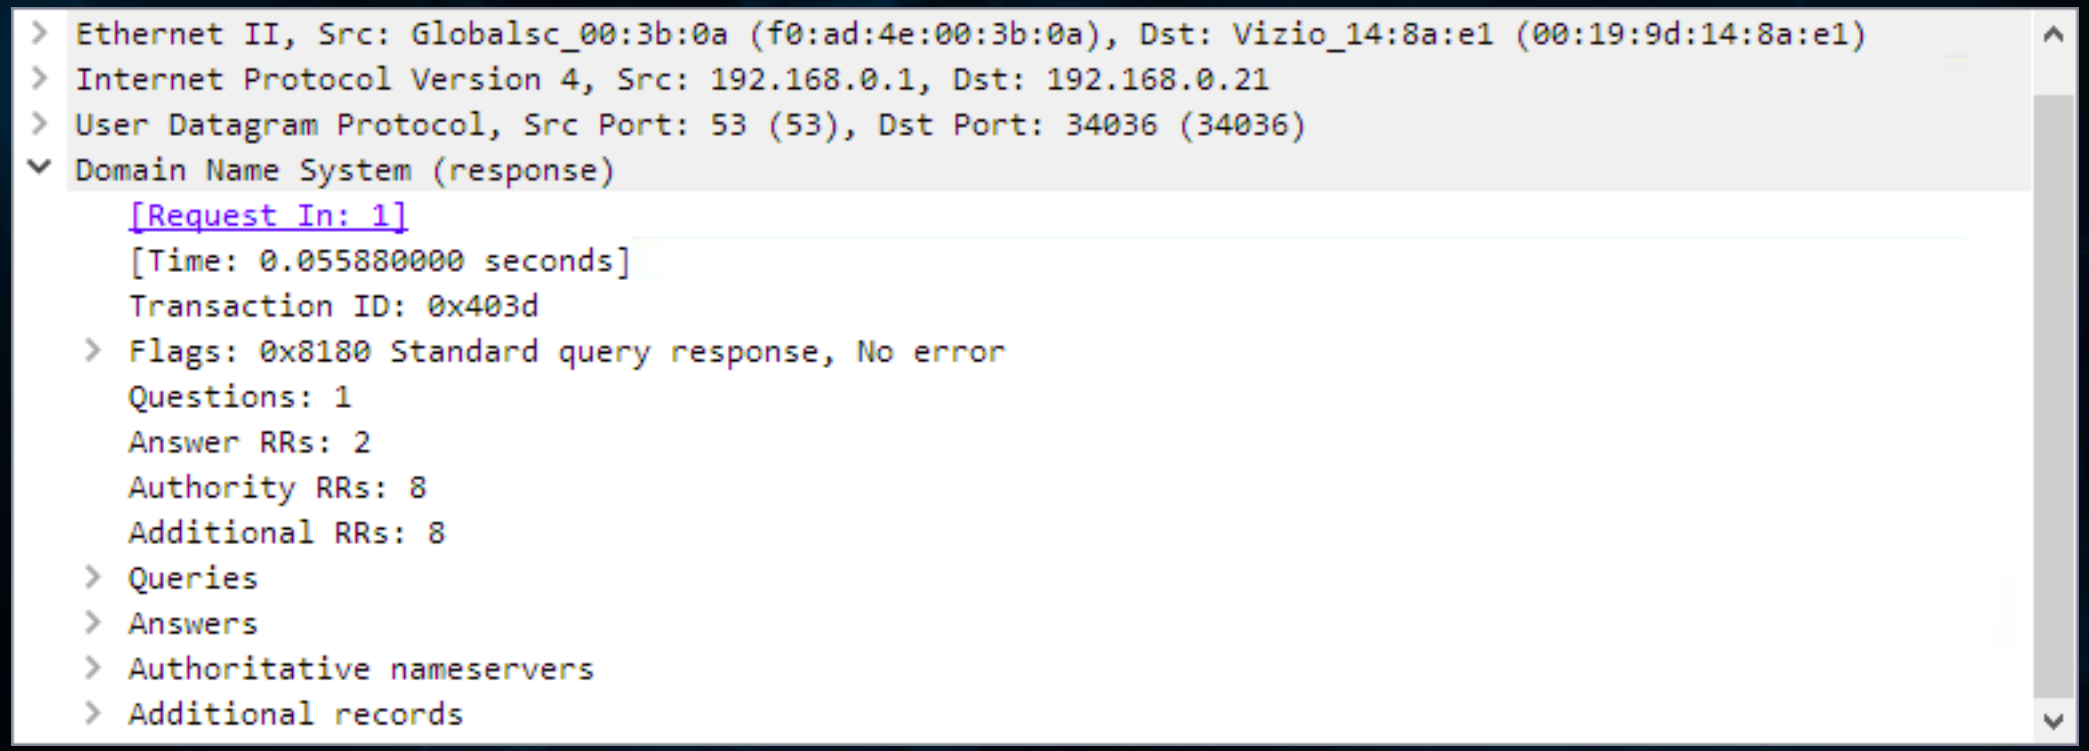

Capture Filter

- 捕包
  - Capture→Options→Start
  - Capture filter
    - 使用tcpdump/pcap相同的 BPF过滤规则语法

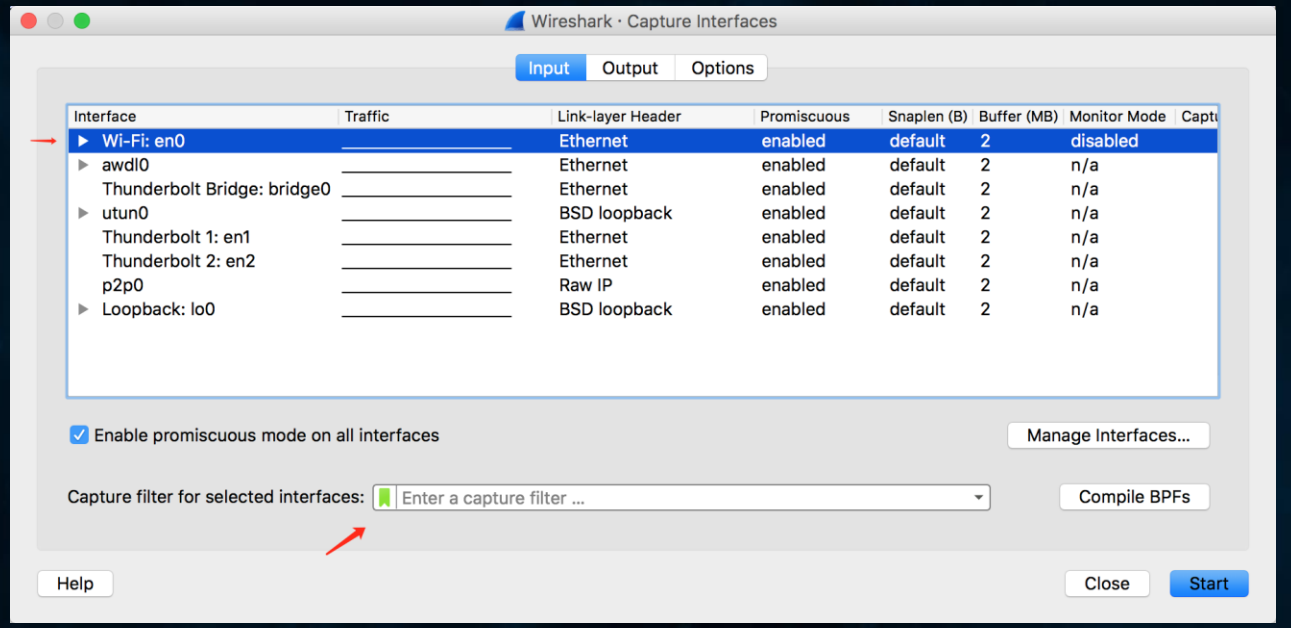

Display Filter

- 分析
  - display filter
    - 手写表达式
      - bootp
      - bootp.hw.mac_addr == 00:50:56:80:bd:c2
      - bootp.ip.relay == 10.207.0.1
    - Display Filter Expression向导
      - Analyze→Display Filter Expression
    - 选择需要过滤的字段
      - 右键→Apply as Filter

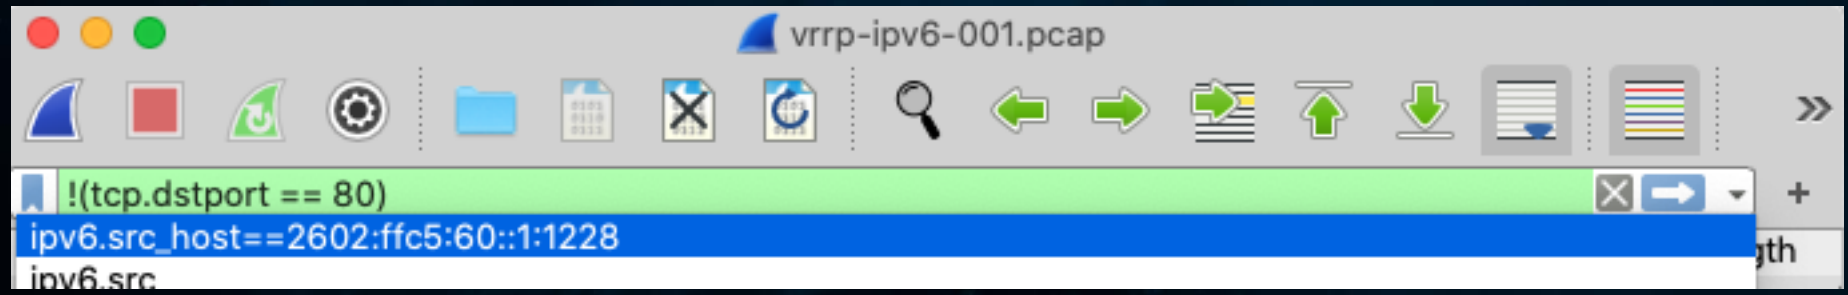

Display Filter

协议字段语法

- expr relop expr
  - **表达式支持 圆括号**
- **协议字段引用**
  - protocol.feature.subfeature
- **逻辑操作符**
  - and, &&
    - Logical AND
  - or, ||
    - Logical OR not,
  - !
    - Logical NOT
- **比较操作符**
  - eq, ==
    - Equal
  - ne, !=
    - Not Equal
  - gt, >
    - Greater Than
  - lt, <
    - Less Than
  - ge, >=
    - Greater than or Equal to
  - le, <=
    - Less than or Equal to
  - **例**
    - ip.addr==192.168.0.1
    - frame.len<=128
  - **比较**
    - ip.addr!=192.168.0.1
    - !(ip.addr==192.168.0.1)

Display Filter

- 匹配操作符
  - contains
    - 字符串匹配
  - matches, ~
    - 正则匹配（Perl 语法）
- 例
  - http contains "wireshark.org”
  - wsp.user_agent matches "cldc"

Display Filter

- 函数
  - upper
    - 转换大写
  - lower
    - 转换小写
  - len
    - 字符串长度
  - count
    - 计算次数
  - string
    - 转换为字符串

- 例
  - lower(http.server) contains "apache”
  - len(http.request.uri) > 100
  - count(ip.addr) > 2
  - string(frame.number) matches "[13579]$”
  - string(ip.dst) matches "^172\.(1[6-9]|2[0-9]|3[0-1])\..{1,3}\.255"

Display Filter

- **slice操作符**
  - proto[offset:len]
    - 提取 proto 协议/字段从 offset 开始的 len 长度字节
    - [i:j]
      - i = start_offset, j = length
    - [i-j]
      - i = start_offset, j = end_offset, inclusive.
    - [i]
      - i = start_offset, length = 1
    - [:j]
      - start_offset = 0, length = j
    - [i:]
      - start_offset = i, end_offset = end_of_field
    - offset 为负时，表示从字段结尾倒数计算
    - offset 支持半角逗号组合

- 例
  - eth.src[0:3] == 00:00:83
  - http.content_type[0:4] == "text”
  - ftp[1,3-5,9:] == 01:03:04:05:09:0a:0b

Display Filter

- in操作符
  - `in { }`
  - `..`
    - 指定连续范围

- 例
  - tcp.port in {80 443 8080}
  - http.request.method in {"HEAD" "GET"}
  - tcp.port in {443 4430..4434}

https://www.wireshark.org/docs/man-pages/wireshark-filter.html
https://www.wireshark.org/docs/dfref/
https://perldoc.perl.org/perlre.html

Display Filter

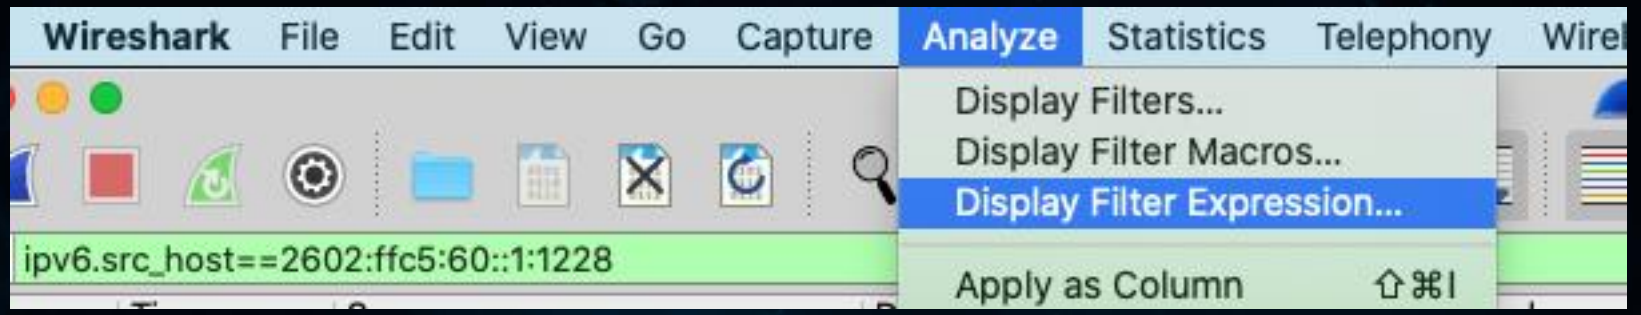

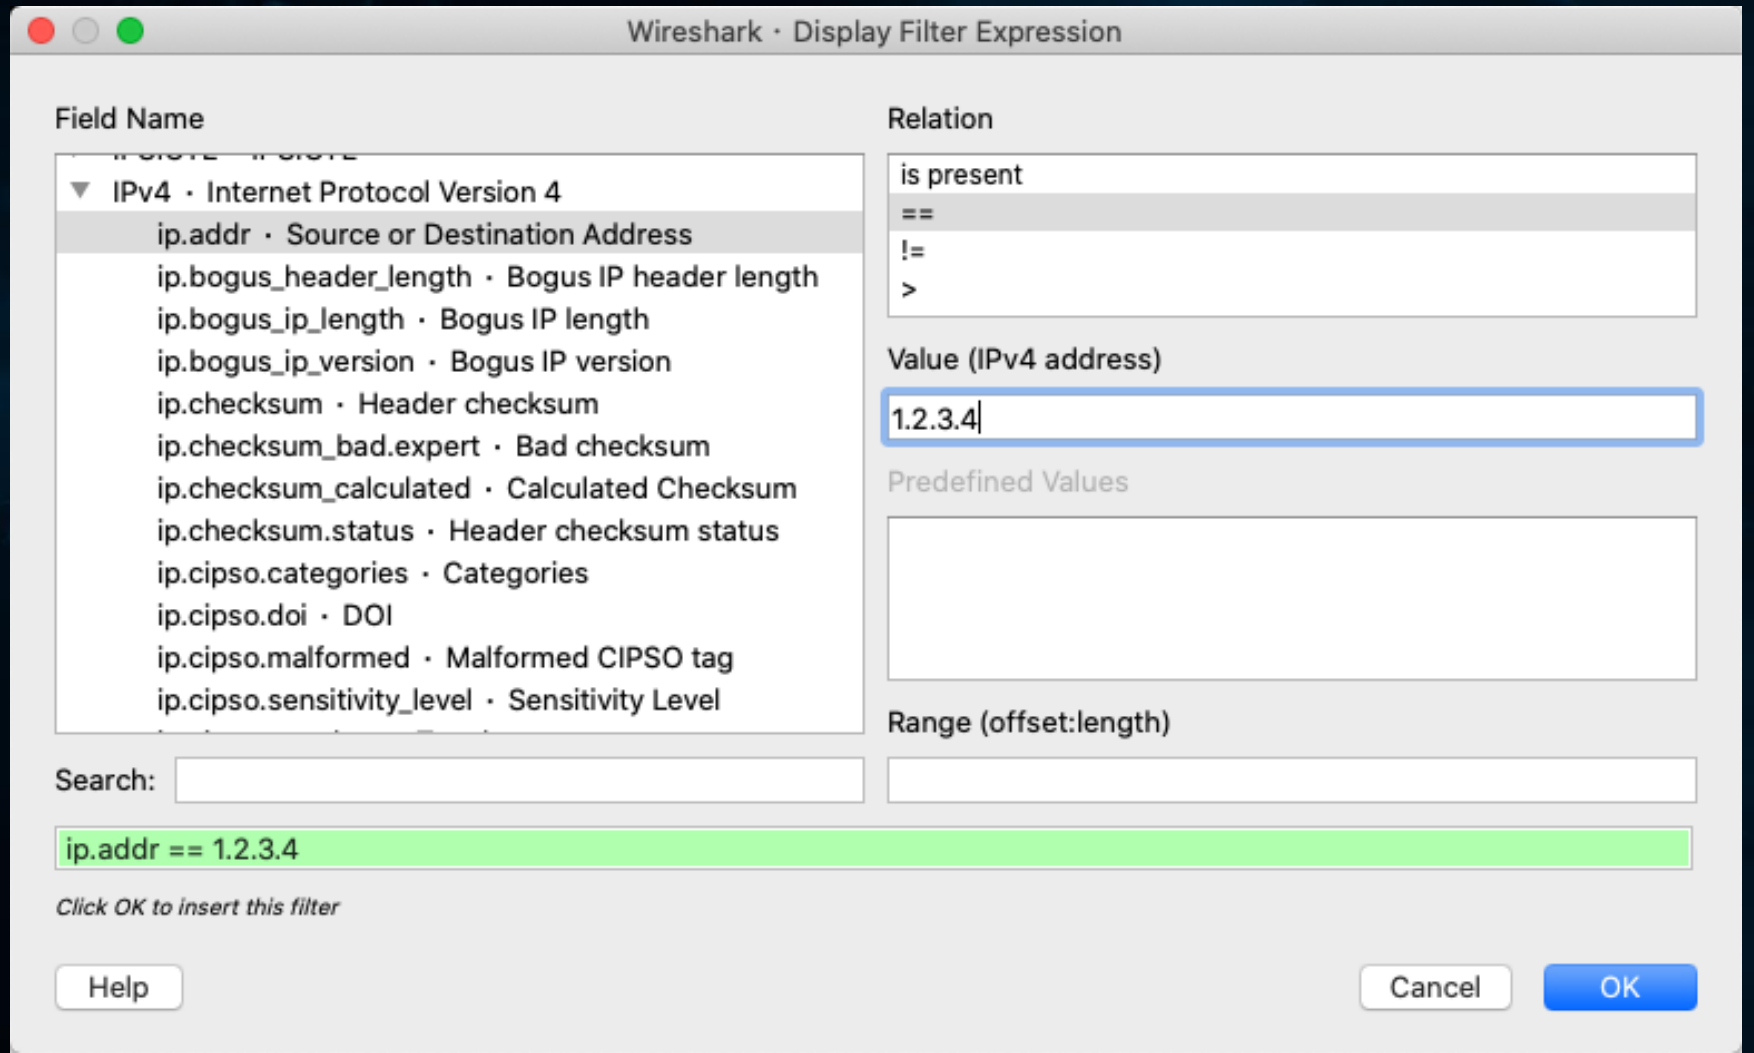

Display Filter

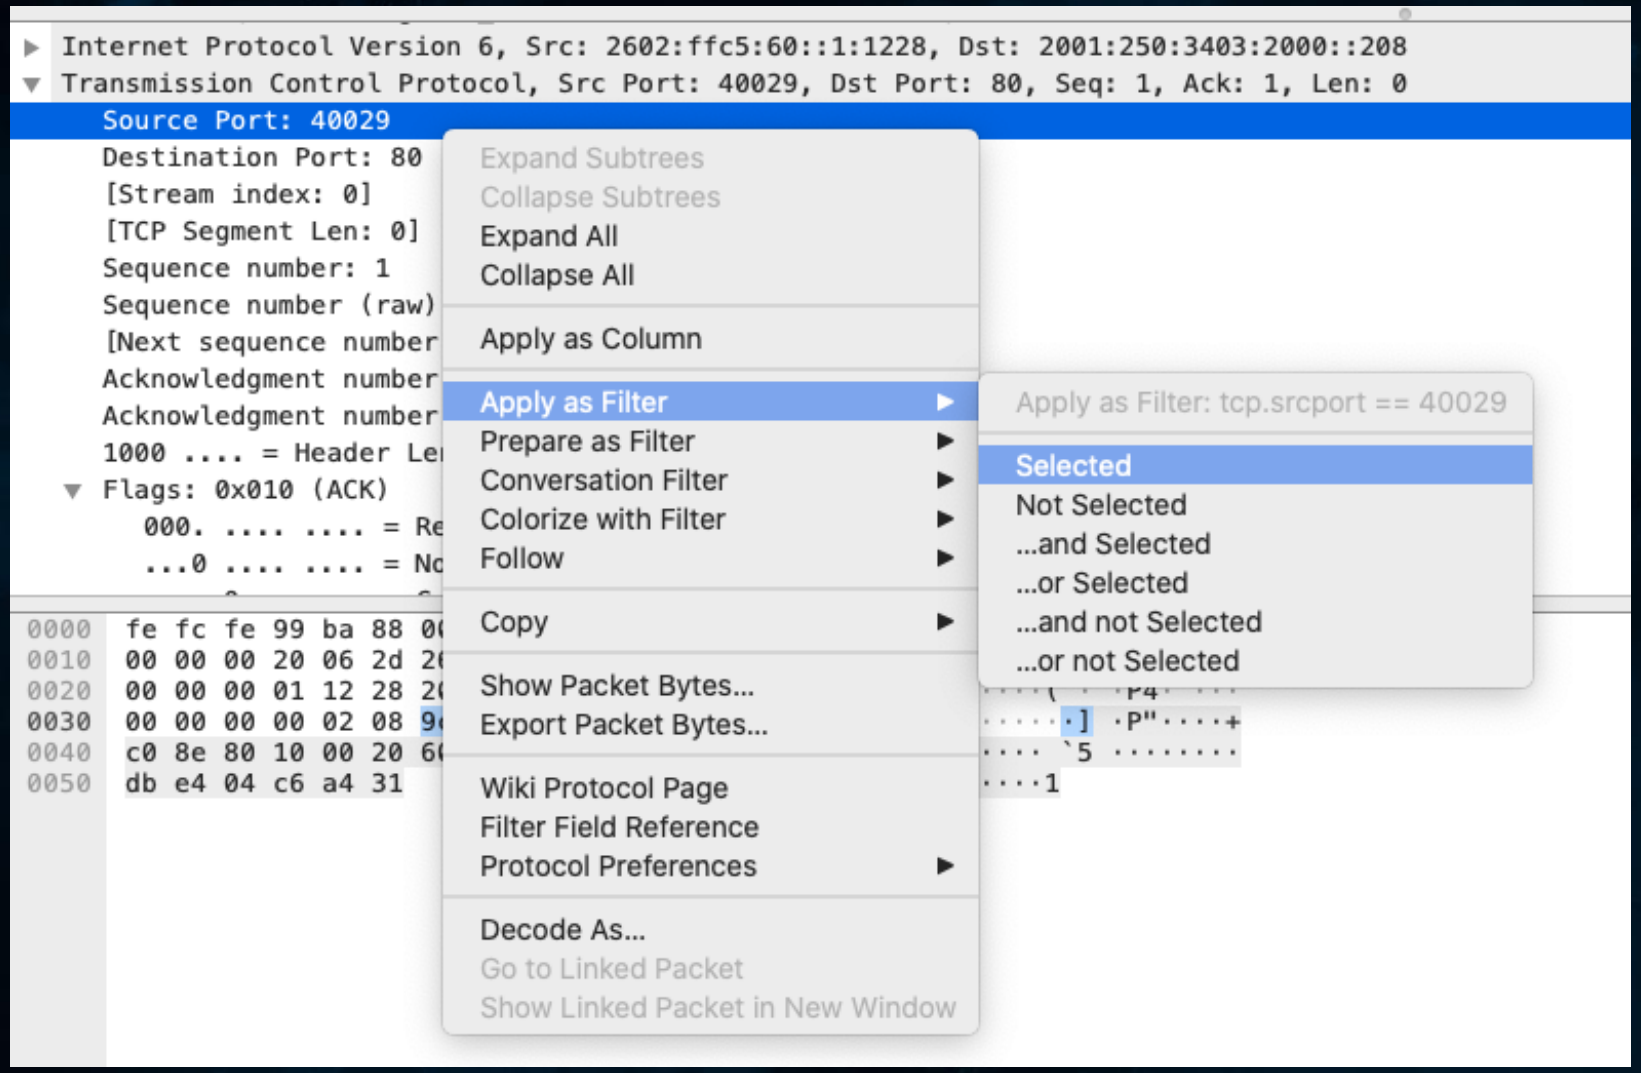

Statistics

- Statistics→Conversations
- Statistics→Protocol Hierarchy

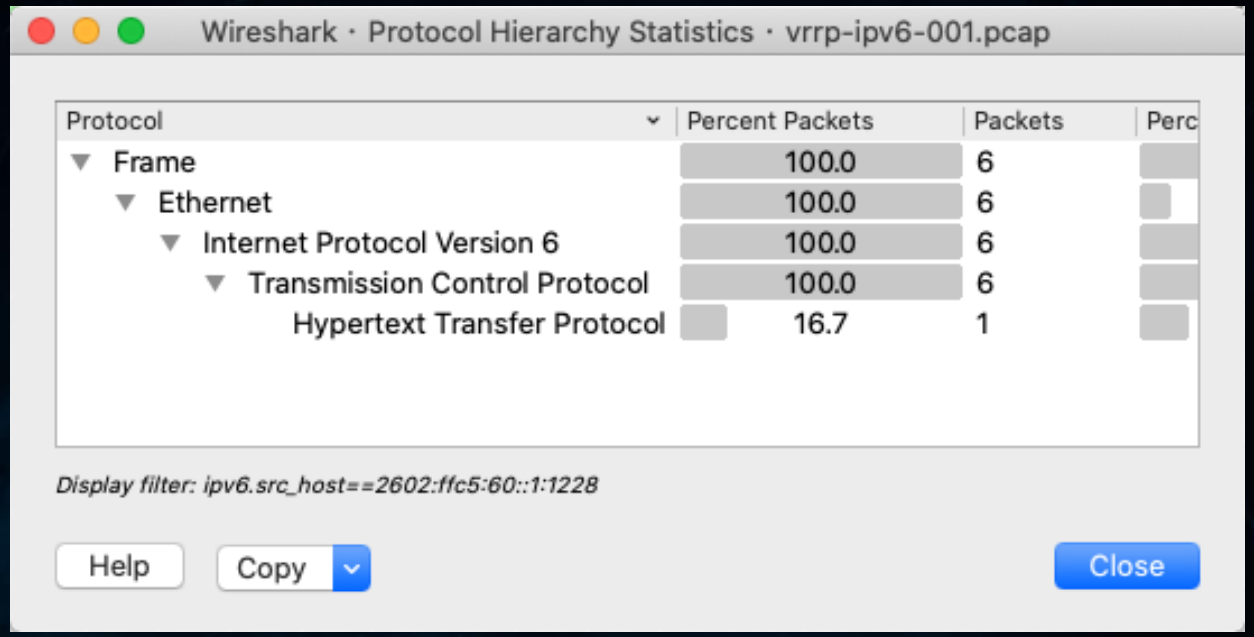

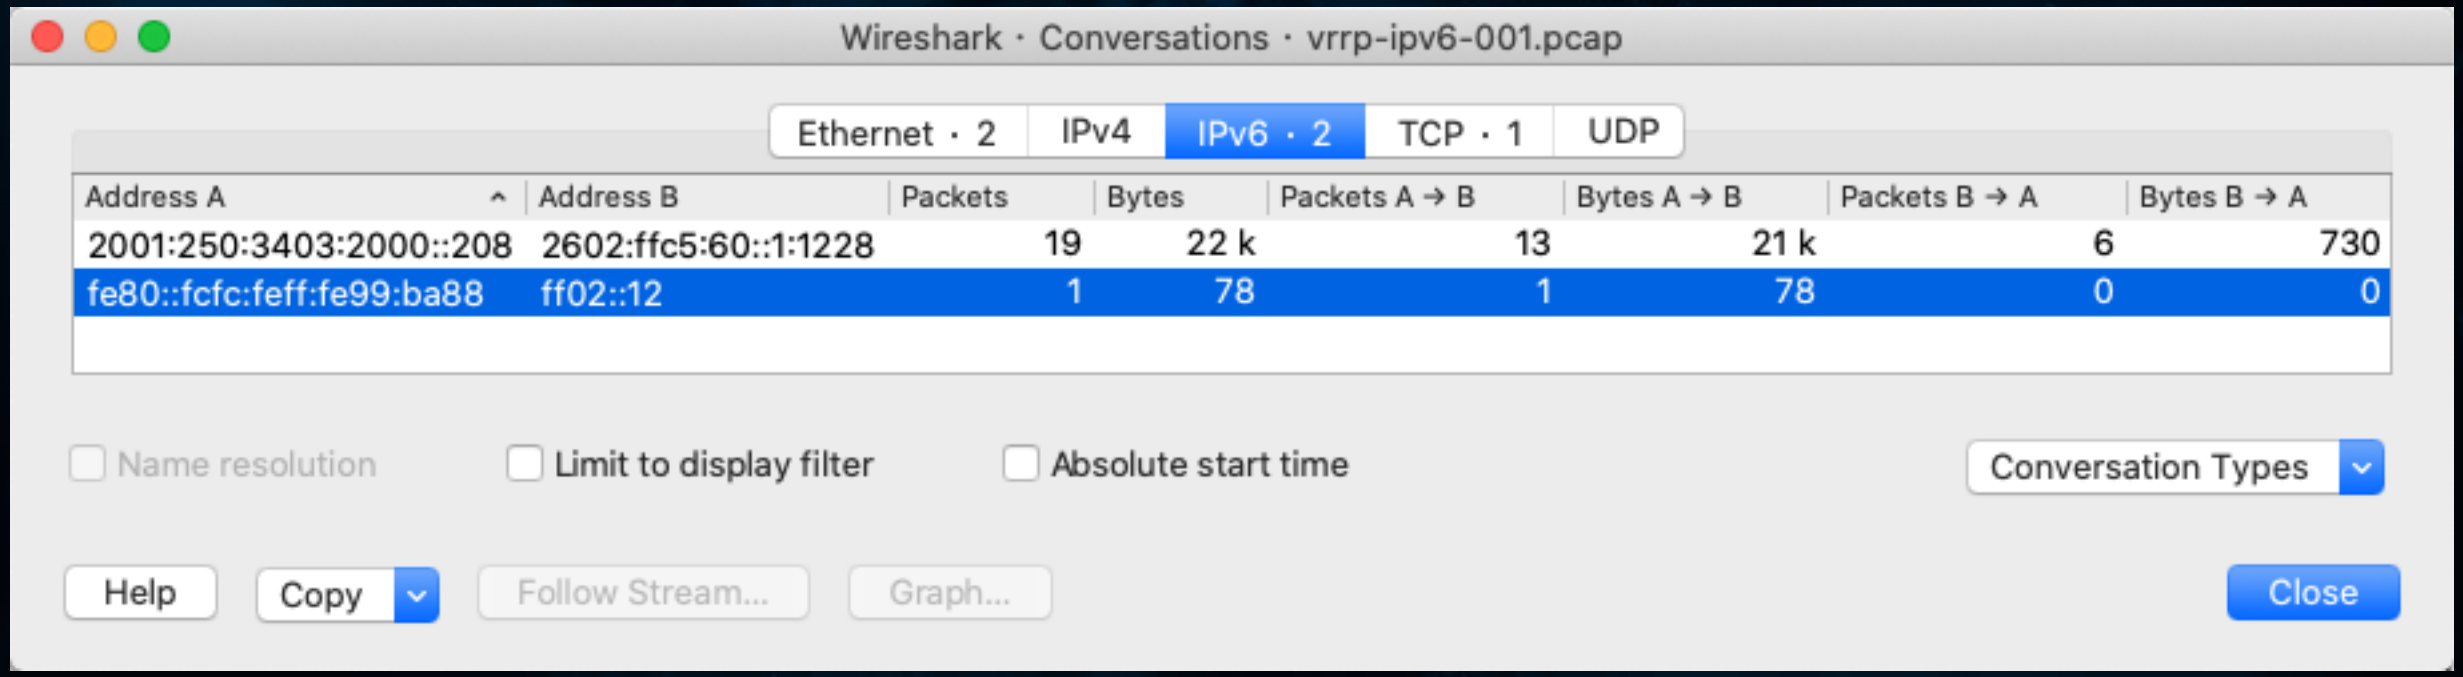

Statistics

- Statistics→Flow Graph

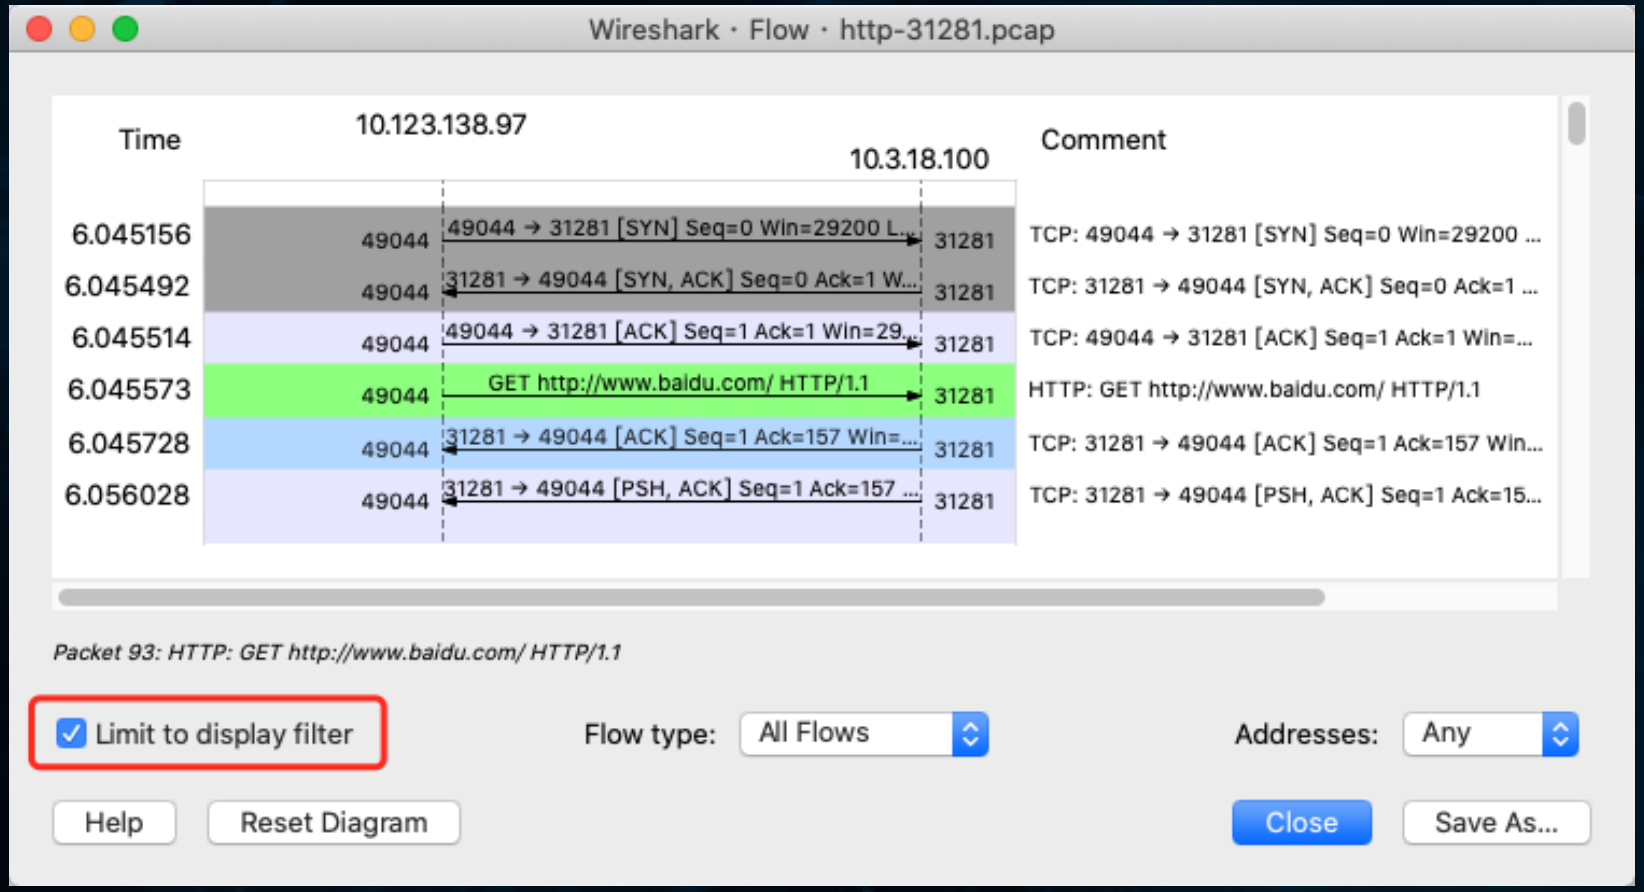

Decode As

- 用 Decode As 选择解析协议

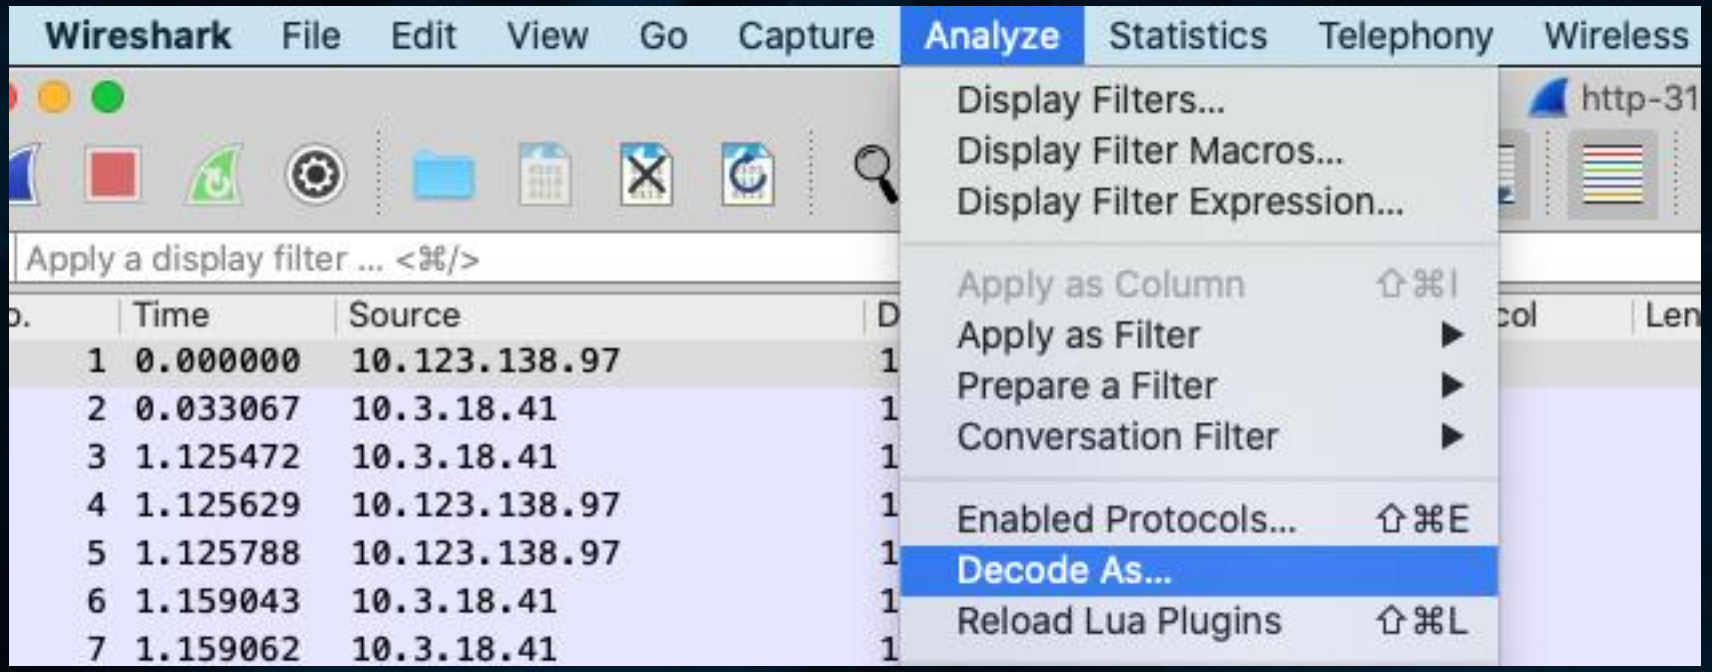

Follow protocol streams

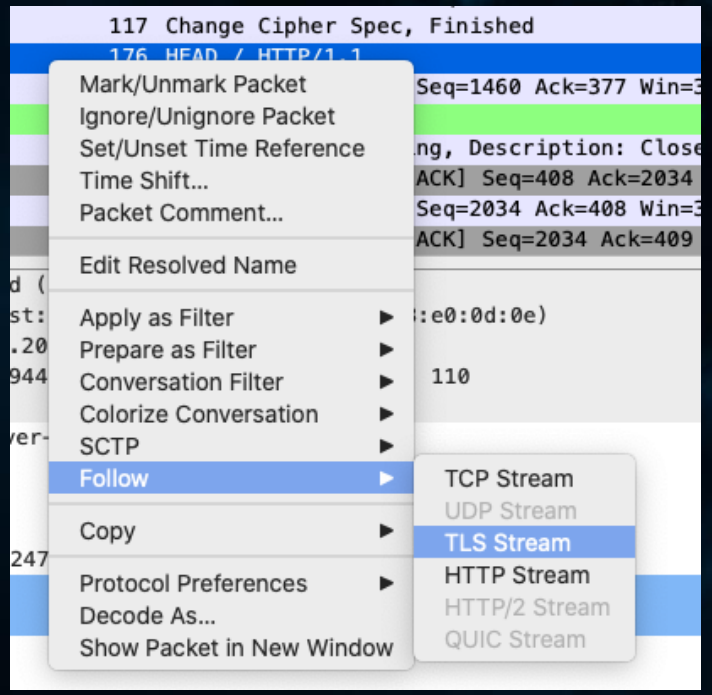

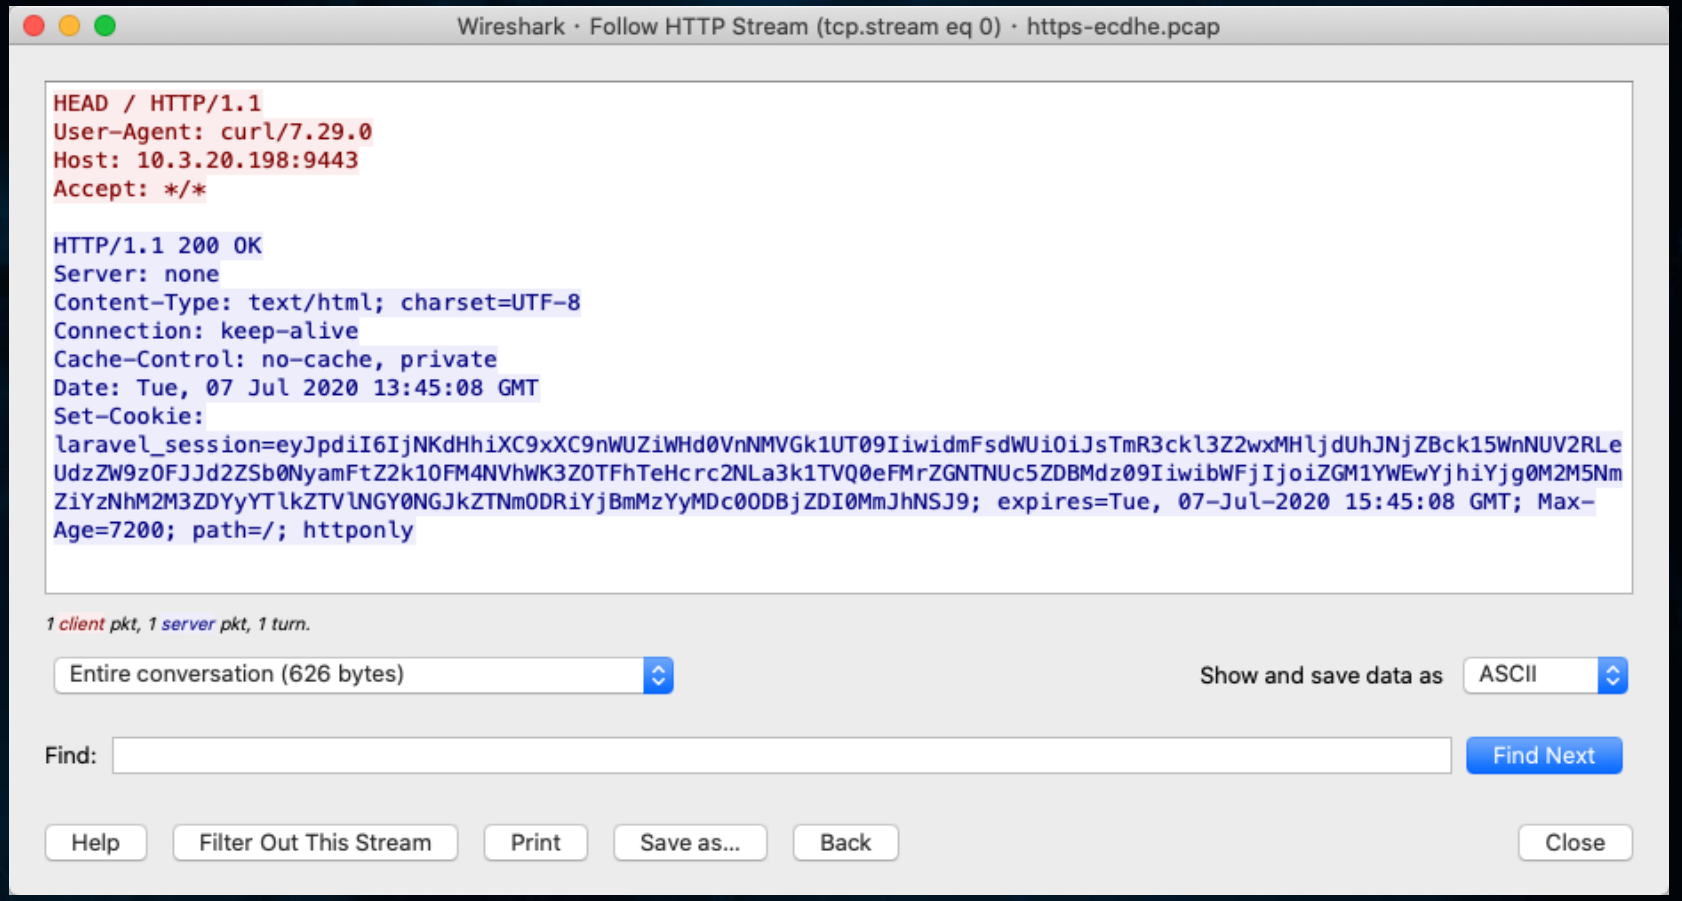

防火墙规则工具

- Tools→Firewall ACL Rules

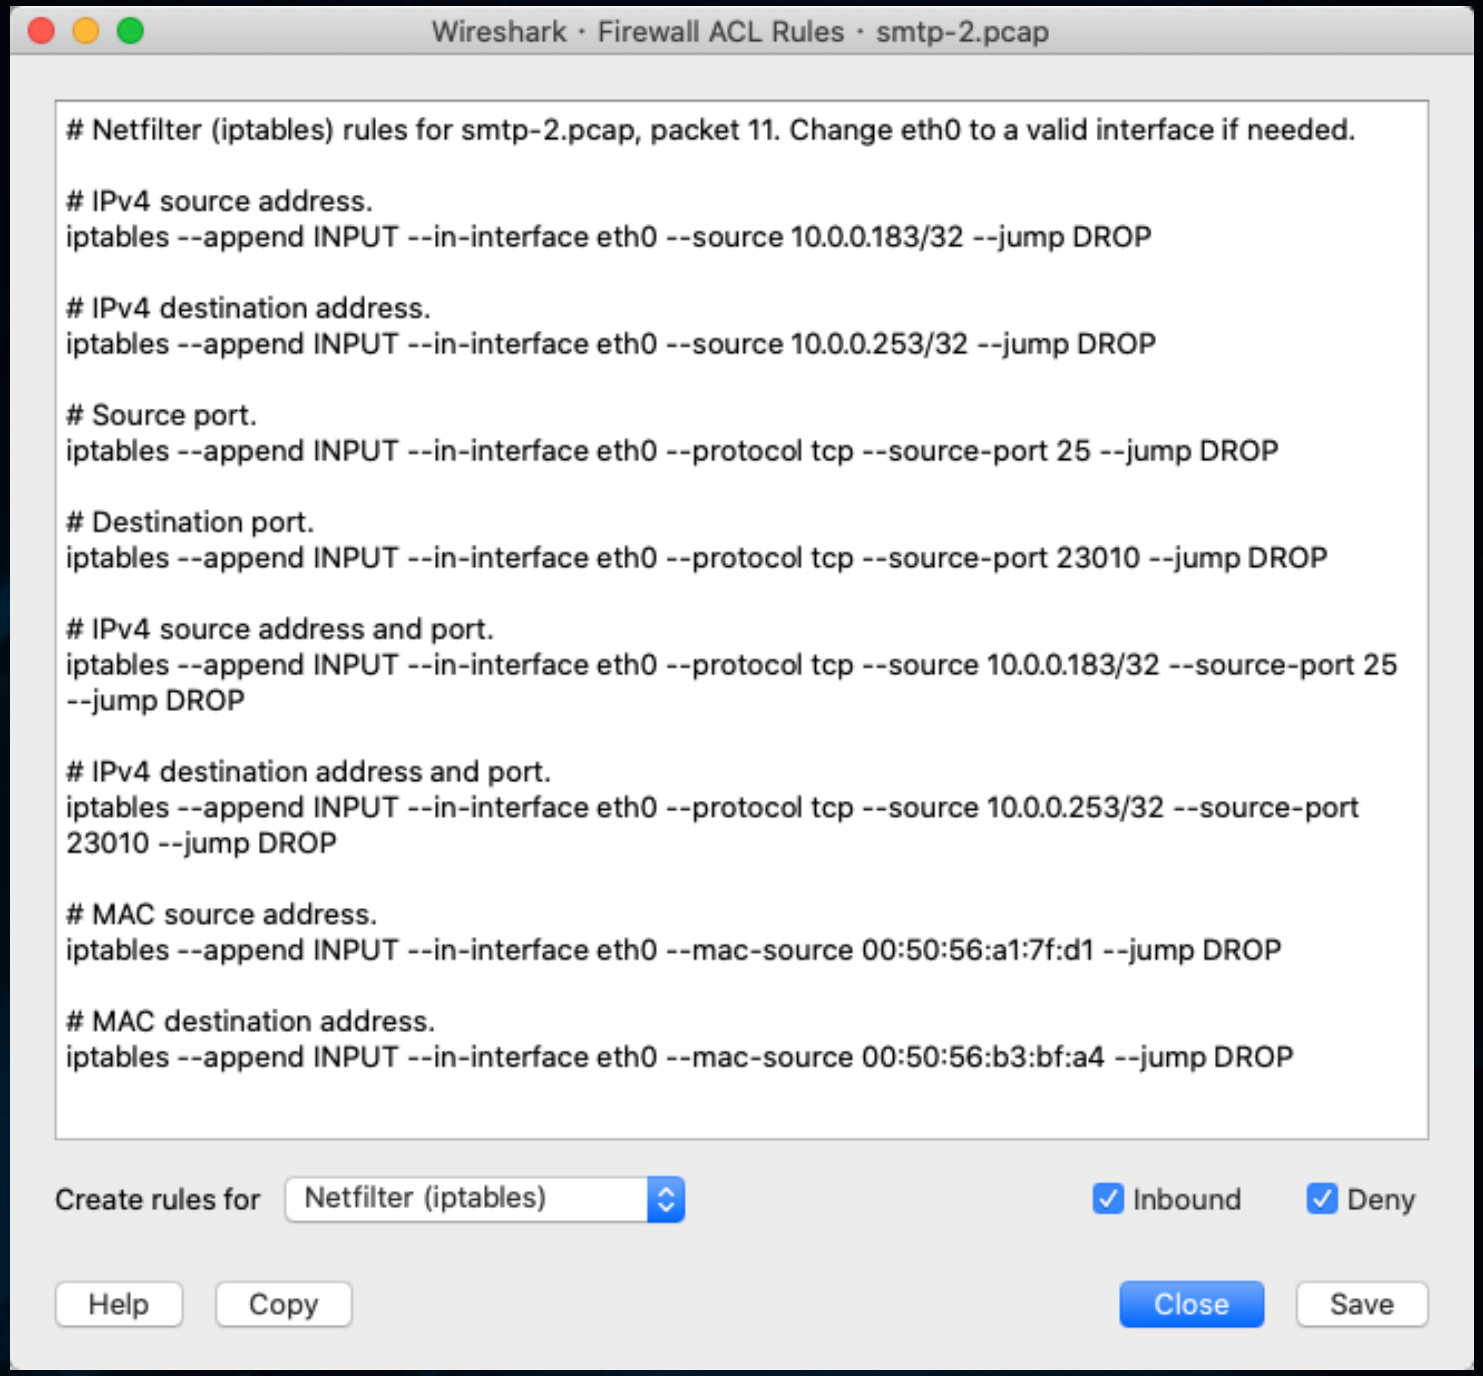

Offload

- xOffload
  - Linux
    - for i in rx tx sg tso ufo gso gro lro; \
      do ethtool -K eth0 $i off; done
  - Windows
    - 网卡属性→网络→配置→高级
      校验和卸载设置.
- TCP Chimney
  - Windows
    - netsh int ip set chimney disabled

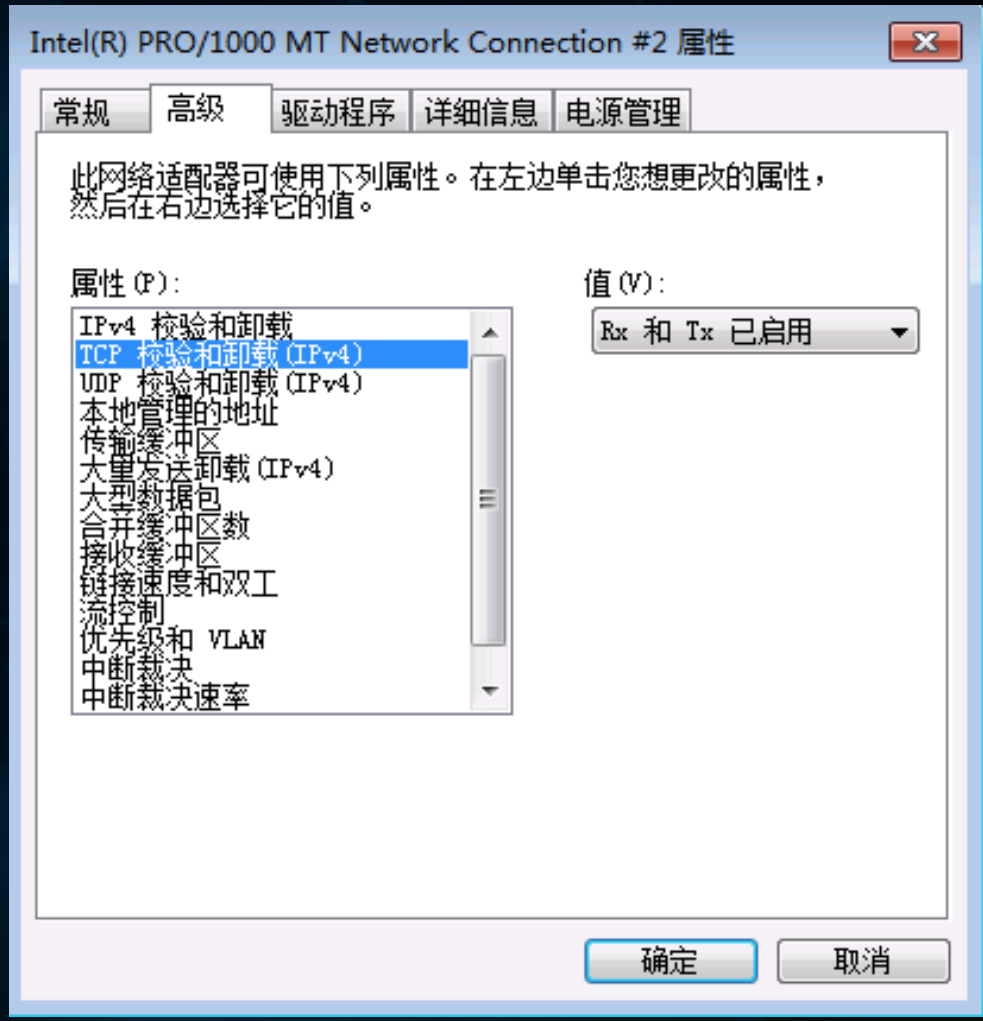

https://gitlab.com/wireshark/wireshark/-/wikis/CaptureSetup/Offloading

Offload

Transmission Control Protocol

- [x] Show TCP summary in protocol tree
- [ ] Validate the TCP checksum if possible
- [x] Allow subdissector to reassemble TCP streams
- [ ] Reassemble out-of-order segments
- [x] Analyze TCP sequence numbers
- [x] Relative sequence numbers (Requires "Analyze TCP sequence numbers")

- 或者关闭 wireshark 校验和检查
  - 属性→协议→TCP/UDP
  - 取消「Validate the TCP/UDP checksum if possible」

```
Sequence number (raw): 2821446576
[Next sequence number: 8187    (relative sequence number)]
Acknowledgment number: 127    (relative ack number)
Acknowledgment number (raw): 582686517
1000 .... = Header Length: 32 bytes (8)
▶ Flags: 0x018 (PSH, ACK)
Window size value: 32
[Calculated window size: 32]
[Window size scaling factor: -1 (unknown)]
▶ Checksum: 0xb5ab incorrect, should be 0xd61b(maybe caused by "TCP checksum offload"?)
[Checksum Status: Bad]
[Calculated Checksum: 0xd61b]
```

https

- SSLKEYLOGFILE=/tmp/a.key
  - 记录 SSL/TLS 密钥

```
[ycflash@kvm7 tmp]$ sudo tcpdump -i ens6f0 -s0 -w http-31281.pc
tcpdump: listening on ens6f0, link-type EN10MB (Ethernet), capt
^C50 packets captured
51 packets received by filter
0 packets dropped by kernel
[ycflash@kvm7 tmp]$
```

```
[ycflash@kvm7 tmp]$ export SSLKEYLOGFILE=/tmp/https-31281.key
[ycflash@kvm7 tmp]$ curl -I https://www.baidu.com
HTTP/1.1 200 Connection established

HTTP/1.1 200 OK
Accept-Ranges: bytes
Cache-Control: private, no-cache, no-store, proxy-revalidate, r
Connection: keep-alive
Content-Length: 277
Content-Type: text/html
Date: Mon, 12 Oct 2020 08:28:23 GMT
Etag: "575e1f5d-115"
Last-Modified: Mon, 13 Jun 2016 02:50:05 GMT
Pragma: no-cache
Server: bfe/1.0.8.18

[ycflash@kvm7 tmp]$ curl -I https://www.baidu.com
```

https

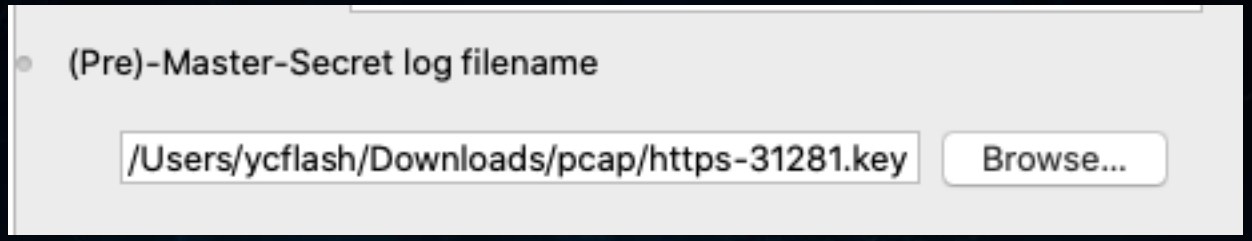

- 属性→协议→TLS→(Pre)-Master-Secret log filename

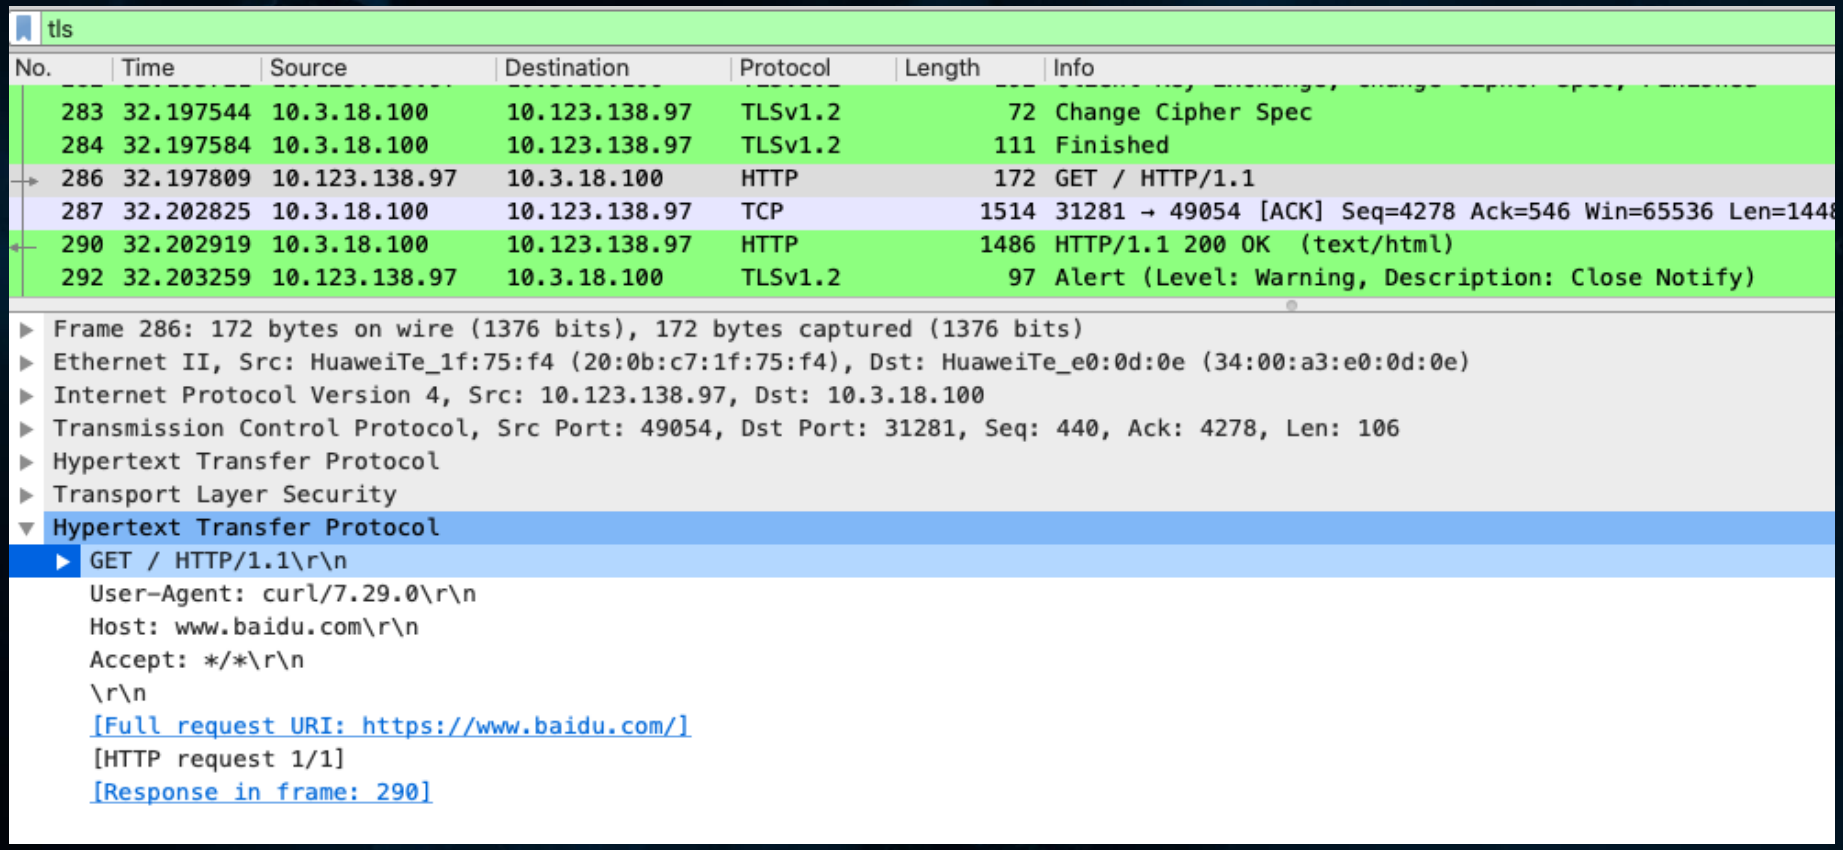

https://wiki.wireshark.org/TLS

SMTP黑名单

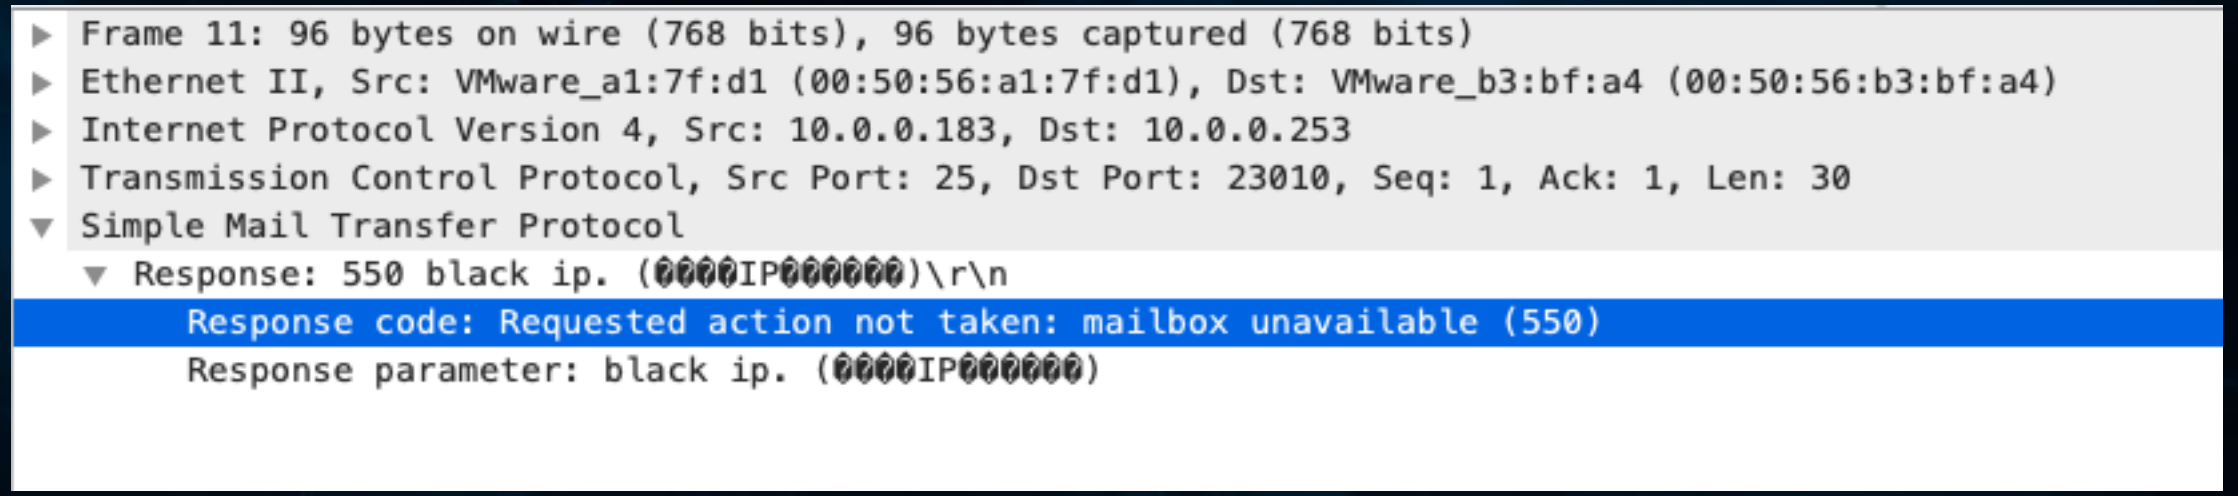
Frame 11: 96 bytes on wire (768 bits), 96 bytes captured (768 bits)
Ethernet II, Src: VMware_a1:7f:d1 (00:50:56:a1:7f:d1), Dst: VMware_b3:bf:a4 (00:50:56:b3:bf:a4)
Internet Protocol Version 4, Src: 10.0.0.183, Dst: 10.0.0.253
Transmission Control Protocol, Src Port: 25, Dst Port: 23010, Seq: 1, Ack: 1, Len: 30
Simple Mail Transfer Protocol
Response: 550 black ip. (????IP??????)\r\n
Response code: Requested action not taken: mailbox unavailable (550)
Response parameter: black ip. (????IP??????)

武汉某高校实例分析

- 现象：
  - DHCP部署及上线测试正常
  - 开学后，大量用户反映无法获取IP地址
  - 检查服务器报文收发正常
- 分析：
  - DHCP服务器报文收发正常，前期部署测试验证正常
  - 抓包分析中继上联口确认中继对服务器收发功能正常
  - 抓包分析中继下联口确认中继对客户端收发功能正常

武汉某高校实例分析

14:59:50.902954 IP (tos 0x0, ttl 62, id 34117, offset 0, flags [none], proto UDP (17), length 328)
    10.183.255.254.bootps > 172.31.1.7.bootps: [udp sum ok] BOOTP/DHCP, Request from 34:36:3b:67:c5:7a (oui Unknown), length 300, hops 1, xid 0xbbb03120, secs 44, Flags [none] (0x0000)
14:59:50.903135 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto UDP (17), length 328)
    172.31.1.7.bootps > 10.183.255.254.bootps: [bad udp cksum 8c20!] BOOTP/DHCP, Reply, length 300, hops 1, xid 0xbbb03120, secs 44, Flags [none] (0x0000)
14:59:51.910102 IP (tos 0x0, ttl 62, id 36909, offset 0, flags [none], proto UDP (17), length 328)
    10.183.255.254.bootps > 172.31.1.7.bootps: [udp sum ok] BOOTP/DHCP, Request from 34:36:3b:67:c5:7a (oui Unknown), length 300, hops 1, xid 0xbbb03120, secs 45, Flags [none] (0x0000)
14:59:51.910255 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto UDP (17), length 328)
    172.31.1.7.bootps > 10.183.255.254.bootps: [bad udp cksum 8b21!] BOOTP/DHCP, Reply, length 300, hops 1, xid 0xbbb03120, secs 45, Flags [none] (0x0000)
14:59:53.973644 IP (tos 0x0, ttl 62, id 42475, offset 0, flags [none], proto UDP (17), length 328)
    10.183.255.254.bootps > 172.31.1.7.bootps: [udp sum ok] BOOTP/DHCP, Request from 34:36:3b:67:c5:7a (oui Unknown), length 300, hops 1, xid 0xbbb03120, secs 47, Flags [none] (0x0000)
14:59:53.973816 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto UDP (17), length 328)
    172.31.1.7.bootps > 10.183.255.254.bootps: [bad udp cksum 8923!] BOOTP/DHCP, Reply, length 300, hops 1, xid 0xbbb03120, secs 47, Flags [none] (0x0000)
      Your-IP 10.183.3.251
      Gateway-IP 10.183.255.254
      Client-Ethernet-Address 34:36:3b:67:c5:7a (oui Unknown)
      Vendor-rfc1048 Extensions
        Magic Cookie 0x63825363
        DHCP-Message Option 53, length 1: ACK
        Server-ID Option 54, length 4: 172.31.1.7
        Lease-Time Option 51, length 4: 2244
        Subnet-Mask Option 1, length 4: 255.255.0.0
        Default-Gateway Option 3, length 4: 10.183.255.254
        Domain-Name-Server Option 6, length 8: 10.182.100.7,10.182.100.6
        END Option 255, length 0
        PAD Option 0, length 0, occurs 22

三次Request 三次ACK

武汉某高校实例分析

14:59:45.625002 IP (tos 0x0, ttl 255, id 33130, offset 0, flags [none], proto UDP (17), length 328)
    0.0.0.0.bootpc > broadcasthost.bootps: [udp sum ok] BOOTP/DHCP, Request from 34:36:3b:67:c5:7a (oui Unknown), length 300, xid 0xbbb03120, secs 44, Flags [none] (0x0000)
14:59:46.632226 IP (tos 0x0, ttl 255, id 33131, offset 0, flags [none], proto UDP (17), length 328)
    0.0.0.0.bootpc > broadcasthost.bootps: [udp sum ok] BOOTP/DHCP, Request from 34:36:3b:67:c5:7a (oui Unknown), length 300, xid 0xbbb03120, secs 45, Flags [none] (0x0000)
14:59:48.695767 IP (tos 0x0, ttl 255, id 33132, offset 0, flags [none], proto UDP (17), length 328)
    0.0.0.0.bootpc > broadcasthost.bootps: [udp sum ok] BOOTP/DHCP, Request from 34:36:3b:67:c5:7a (oui Unknown), length 300, xid 0xbbb03120, secs 47, Flags [none] (0x0000)
14:59:48.699753 IP (tos 0x0, ttl 64, id 42477, offset 0, flags [none], proto UDP (17), length 328)
    10.183.255.254.bootps > 10.183.3.251.bootpc: [udp sum ok] BOOTP/DHCP, Reply, length 300, hops 1, xid 0xbbb03120, secs 47, Flags [none] (0x0000)
                Your-IP 10.183.3.251
                Gateway-IP 10.183.255.254
                Client-Ethernet-Address 34:36:3b:67:c5:7a (oui Unknown)
                Vendor-rfc1048 Extensions
                  Magic Cookie 0x63825363
                  DHCP-Message Option 53, length 1: ACK
                  Server-ID Option 54, length 4: 172.31.1.7
                  Lease-Time Option 51, length 4: 2244
                  Subnet-Mask Option 1, length 4: 255.255.0.0
                  Default-Gateway Option 3, length 4: 10.183.255.254
                  Domain-Name-Server Option 6, length 8: 10.182.100.7,10.182.100.6
                  END Option 255, length 0
                  PAD Option 0, length 0, occurs 22

三次Request 一次ACK

武汉某高校实例分析

- 通过对中继上下行抓包分析，确认：
  - 服务器所有应答报文均送达DHCP中继
  - DHCP中继转发给客户端的应答报文发生丢失
- 解决：
  - 核查中继设备配置（N18K），发现其默认开启nfpp功能中有dhcp-guard选项，限制了同一MAC的DHCP报文速率为5条
  - 关闭nfpp中的dhcp-guard选项：

```
config t
nfpp
 no dhcp-guard enable
end
```

TCP timestamps 案例

No.	Time	Source	Destination	Protocol	Length	Info
11	0.055033	117.83.194.89	106.120.213.245	TLSv1.2	194	Client Key Exchange, Change Cipher Spec,
12	0.055258	106.120.213.245	117.83.194.89	TLSv1.2	310	New Session Ticket, Change Cipher Spec,
13	0.081926	117.83.194.89	106.120.213.245	TLSv1.2	634	Application Data
14	0.107303	106.120.213.245	117.83.194.89	TLSv1.2	244	Application Data
15	0.107370	106.120.213.245	117.83.194.89	TCP	68	8443 → 31095 [FIN, ACK] Seq=3767 Ack=951
16	0.132677	117.83.194.89	106.120.213.245	TLSv1.2	99	Encrypted Alert
17	0.132789	117.83.194.89	106.120.213.245	TCP	68	31095 → 8443 [FIN, ACK] Seq=982 Ack=3767
18	0.132795	106.120.213.245	117.83.194.89	TCP	68	8443 → 31095 [ACK] Seq=3768 Ack=983 Win=
19	0.132840	117.83.194.89	106.120.213.245	TCP	68	31095 → 8443 [ACK] Seq=983 Ack=3768 Win=
20	10.924512	117.83.194.89	106.120.213.245	TCP	76	10362 → 8443 [SYN] Seq=0 Win=28040 Len=0
21	11.326894	117.83.194.89	106.120.213.245	TCP	76	17305 → 8443 [SYN] Seq=0 Win=28040 Len=0
22	11.934599	117.83.194.89	106.120.213.245	TCP	76	[TCP Retransmission] 10362 → 8443 [SYN]
23	12.330044	117.83.194.89	106.120.213.245	TCP	76	[TCP Retransmission] 17305 → 8443 [SYN]
24	13.931373	117.83.194.89	106.120.213.245	TCP	76	[TCP Retransmission] 10362 → 8443 [SYN]
25	14.333213	117.83.194.89	106.120.213.245	TCP	76	[TCP Retransmission] 17305 → 8443 [SYN]
26	17.937810	117.83.194.89	106.120.213.245	TCP	76	[TCP Retransmission] 10362 → 8443 [SYN]
27	18.345099	117.83.194.89	106.120.213.245	TCP	76	[TCP Retransmission] 17305 → 8443 [SYN]

Wireshark · Conversations · 反代到电信云.pcap

Ethernet | IPv4 · 1 | IPv6 | TCP · 3 | UDP

Address A	Port A	Address B	Port B	Packets	Bytes	Packets A → B	Bytes A → B	Packets B → A	Bytes B → A
117.83.194.89	31095	106.120.213.245	8443	19	6055	11	1737	8	4318
117.83.194.89	10362	106.120.213.245	8443	4	304	4	304	0	0
117.83.194.89	17305	106.120.213.245	8443	4	304	4	304	0	0

Linux cooked capture?

- 抓包没有完整的二层帧头

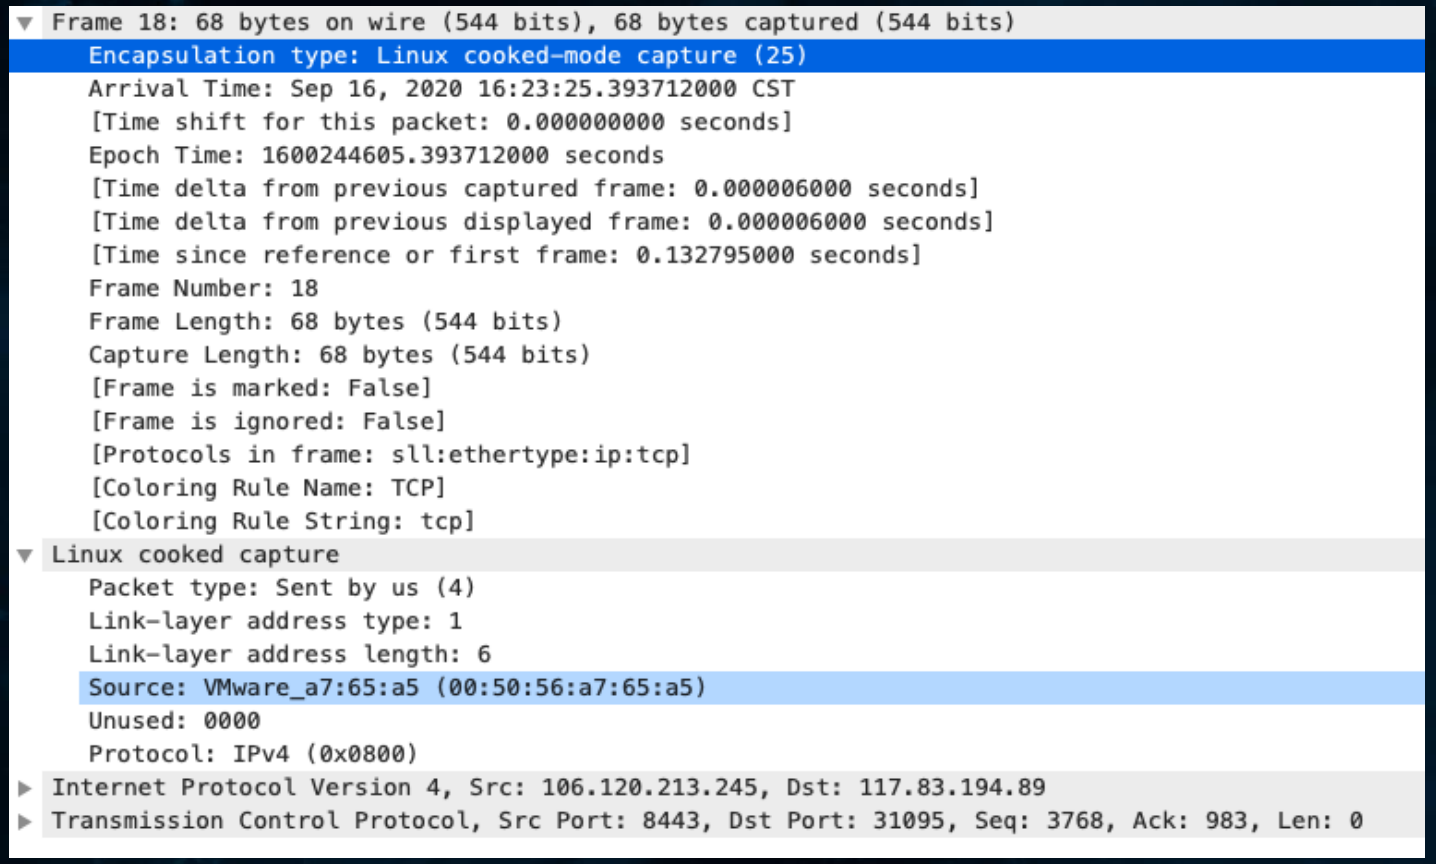

10G网络线速下抓包？

- 商业硬件方案
  - 使用采集卡如 Endace DAG 10X2-S 捕获全部报文
  - 使用 TAP 设备过滤功能过滤报文
- 开源软件方案
  - gulp
    - https://github.com/jmakov/gulp
  - PF_RING（免费，不满足线速捕获）
    - https://www.ntop.org/get-started/download/#PF_RING
  - PF_RING ZC（收费）
    - https://www.ntop.org/products/packet-capture/pf_ring/pf_ring-zc-zero-copy/

北京网瑞达科技有限公司
Beijing WRD Technology Co., Ltd.

提问时间

微信公众号

网瑞达小助理

欢迎合作

电话：010-62285865
传真：010-62285165
网站：www.wrdtech.com
地址：北京市海淀区文慧园路35号

北京网瑞达科技有限公司
Beijing WRD Technology Co., Ltd.