PcapParser:
DNS pcap Made Easy

Shane Kerr / BII Labs / shane@biigroup.cn
Runxia Wan · 万润夏 / BII Labs / rxwan@biigroup.cn

2016-11-13 / Seoul · 서울 / IEPG

pcap for recording DNS traffic

- DNS software historically has little or no facilities to record traffic
  - Server slow? Turn query logging off, you fool!
- Traffic recording happens external to DNS
  - On the box, via mirrored ports, and so on
- UDP pattern quite simple
  - 1 packet query, 1 packet answer

Bii 天地互连

Problems with pcap

- Modern DNS is *usually* a simple 2-packet UDP exchange
- Modern DNS *can* be TCP at any time
  - TCP was always used for zone transfers
- Large messages in responses less clean
  - UDP fragments at IP layer
  - TCP may fragment, is always a stream

Bii 天地互连

PcapParser

- Go program
- Takes pcap input, writes pcap output
- Defragments IP packets
- Reads TCP streams
- Writes unfragmented UDP/IP packets
- DNS-specific
  - Needed when reading TCP streams
  - DNS messages always fit in a UDP packet

Bii
天地互连

Defragmenting IP

- Conceptually simple: collect fragments and re-build the original
- IPv4 and IPv6 are basically the same
  - Kind of a crappy algorithm ☹
- gopacket library has IPv4 support!
- Runxia created IPv6 support
  - Pushing upstream, non-trivial IPR stuff ☹

Bii 天地互连

Reassembling TCP streams

- Potentially really hard
  - Need to implement a lot of TCP
- gopacket library has TCP reassembly! 😄
- Handle DNS TCP stream built from pcap
  - 2-byte length, then data
- Build UDP/IP packet with DNS payload

Bii
天地互连

Final Thoughts on DNS Logging

- pcap is not a great match
  - Lots of duplicated & unneeded information
  - Missing information (query/answer pairs, server state)
- dnstap aims to be ”real” DNS logging
  - Although does not define a file format?
- CBOR seems to be popular
- Some day encryption will make out-of-server logging worthless (hopefully)

Bii 天地互连

Links

- http://dnsv6lab.net/2016/09/06/DNS-pcap-fragments/
- https://github.com/RunxiaWan/PcapParser

Bii
天地互连