FUNET Measurements - Lessons Learned

Markus Peuhkuri

2005-02-17

Content

- FUNET measurements
- What we are collecting
- What we plan to do

FUNET measurements

- FUNET backbone (2.5 Gbit/s)
- On average, 10 Mbit/s of compressed trace
  $\Rightarrow \approx 1$ TiB/week
- System properly running since May 2004
- Stales statistics calculated
- Samples stored: 7.5 TiB
  - 4 complete weeks
  - 71 complete days, 24 partial
  - 2^{-28} s resolution (3.7 ns)

Stateless statistics calculated for all data

- IP protocol (TCP, UDP, ...) counts for every second
  - packet length histogram (40, 64, 128, 256, 512, 1024)
- OSPF packet timestamps and lengths
- For every 10-second file
  - TCP, UDP port counts and packet length histogram
  - TTL histogram by IP protocol
  - TCP retransmissions/reorders in 32k byte window

Capture machine

- Dual Xenon 2.4 GHz
- 2 GB of memory
- 120 GB system disk
- 4*160 GB IDE disks for data
- 2*1000BaseT NIC
- Endace DAG 4.23 OC48 capture cards

- Linux 2.4.20
- Performance:

  Disk I/O write 77 MB/s = 617 Mbit/s
  (currently as 2 RAID-0 stripes)

  Compression / Anonymization initial tests:

  single-thread 2.5 Gbit/s (disk-disk)

  double-thread 7 Gbit/s (estimated based on CPU usage)

  Compression ratio about 12 bytes / packet, 1:40 reduction for wire speed
- No single packet lost!
- Unfortunatly, no conditional full capture

Storage and analysis

- Daily volume
  - 23rd December: 55 GiB
  - 26th September: 124 GiB
- Weekly volume: 600 – 750+ GiB
  ⇒ to do a week-long analysis more than 1 TiB disk capacity needed
- Capture machine keeps disks 90 % full
- Problems with CSC tape archive
  - maintenence periods
  - results holes if data ages
    ⇒ need for temporary storage as buffer
- Analysis needs 1 GiB memory to start with
  - needed to buffer IP address db
  - more stateful analysis
  - basic flow analysis about real-time

Findings: protocols used: week 2004/48 and 2005/5

- TCP protocols
  - top 10 ports use 25 – 30 % of bandwidth
  - http (16 – 20 %), nntp (20 –30 %)
  - p2p traffic halved
  - email traffic 0.5 – 0.7 %
- UDP protocols
  - game traffic
  - dns
  - malicious traffic

Findings: TCP retransmissions and hop count

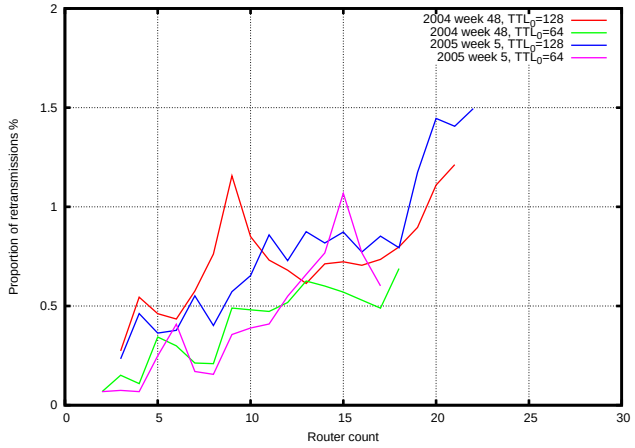

Findings: TCP retransmission distance

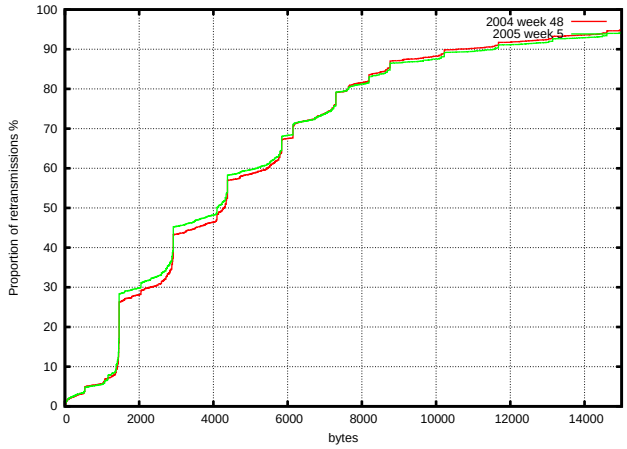

Findings: OSPF non-Hello IAT

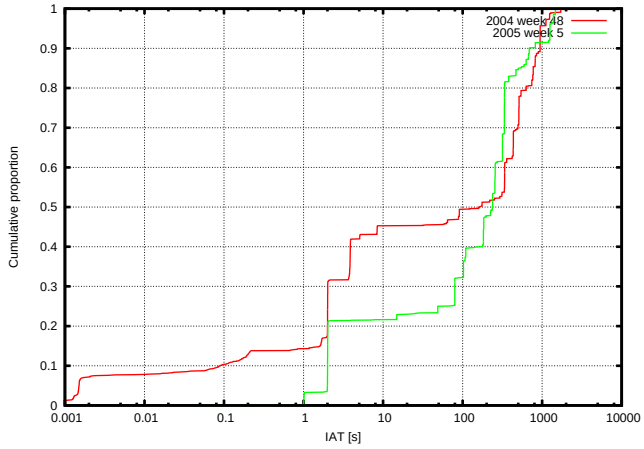

Future plans

- FUNET probe as long as possible
  - Endace DAG for 10G is $\approx 2 \times 25\,\text{k€}$
  - CSC owns data and rights to distribute
    - data potentially identifying organisations also sensitive

- International co-operation
  - different locations
  - synchronisation
  - data access rights
- If there are policy, technical and scientific issues, try to solve parallel
- Capturing ever faster networks
  - sampling strategies
- Deploying low-cost PC-based captures for residential broadband
  - silent capture unit $\approx 600\,€$
  - wireless edge?