

An endhost-centric approach to detect network performance problems

Good

Poor



Olivier Tilmans

olivier.tilmans@uclouvain.be

Maprg IETF 101, London

March 20, 2018

Joint work with Olivier Bonaventure

The performance of the transport layer is *indirectly monitored* in today's enterprise networks

- Netflow tracks bandwidth usage
- Stateful middleboxes infer TCP flow states

“ Authenticated and encrypted header and payload”


“Authenticated and encrypted header and payload”



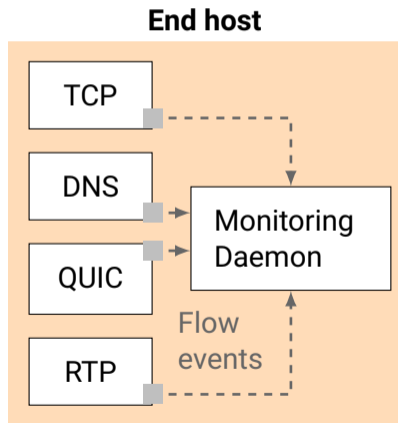
“ Authenticated and encrypted header and payload”

Back to an end-to-end approach to understand
transport layer performance

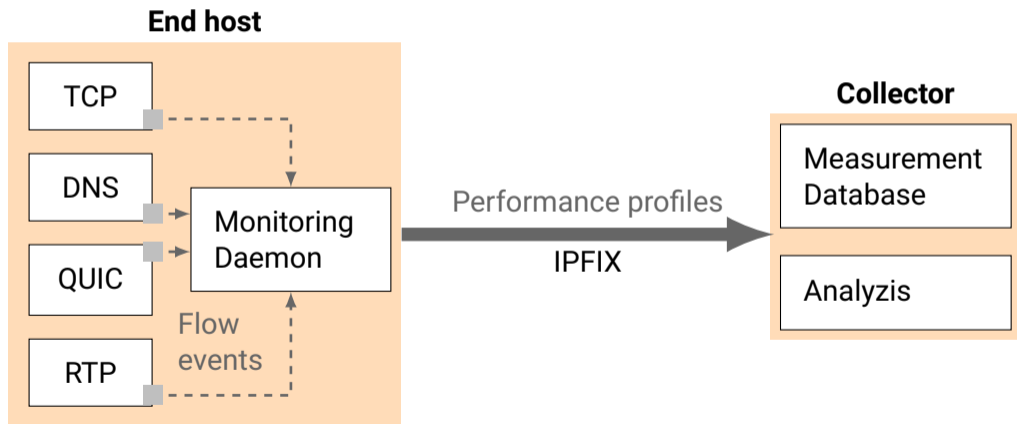
We instrument end-hosts transport stacks



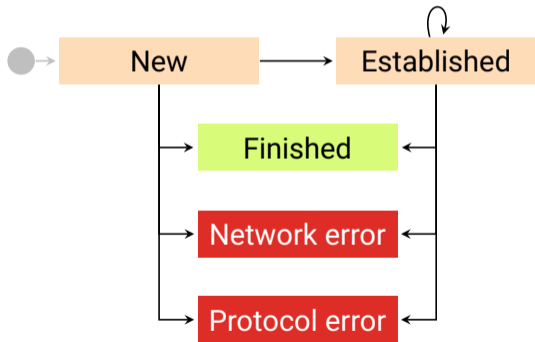
We instrument end-hosts transport stacks and export generic statistics



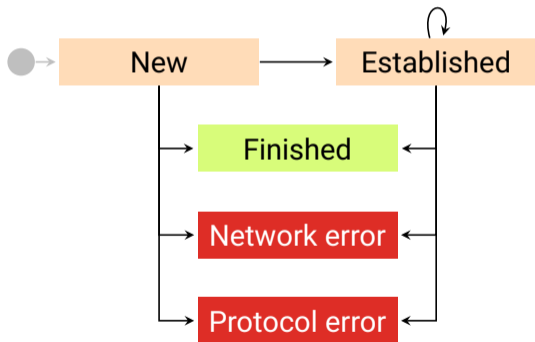
We instrument end-hosts transport stacks and export generic statistics



We abstract protocols using a generic FSM and compute performance profiles on state transitions



We abstract protocols using a generic FSM and compute performance profiles on state transitions




Each transition exports statistics, e.g.:

- Retransmissions
- Duplicated received segments
- RTT
- Jitter
- Time since the last transition
- IP nexthop/uplink ID

An endhost-centric approach to detect network performance problems

Good

Poor




- Instrumenting end-host transport stacks
- Observing network performance

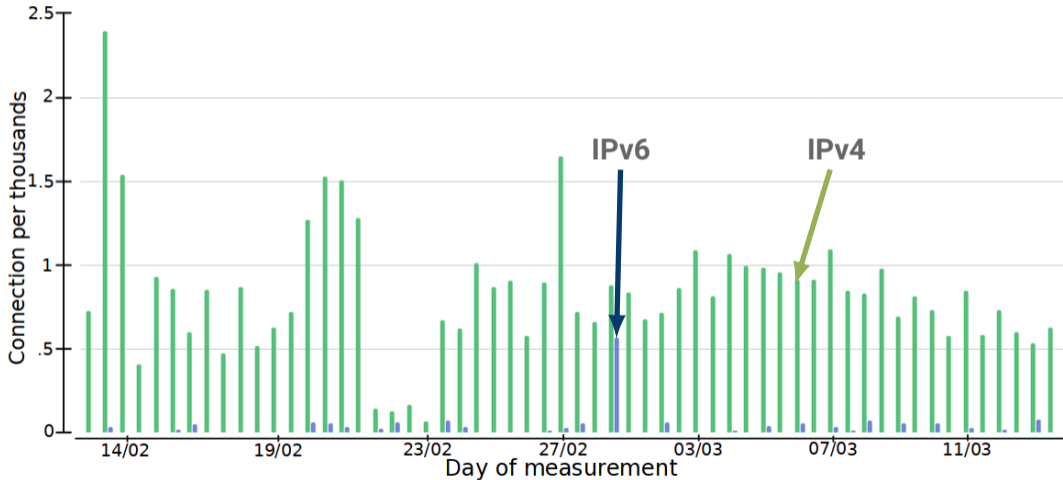
Prototype deployment on student lab computers

- Linux machines, with a non-intrusive instrumentation based on kprobes/eBPF
- TCP/DNS as examples to illustrate the methodology
- Data from the past month
- Few flows exhibit bufferbloat or reordering

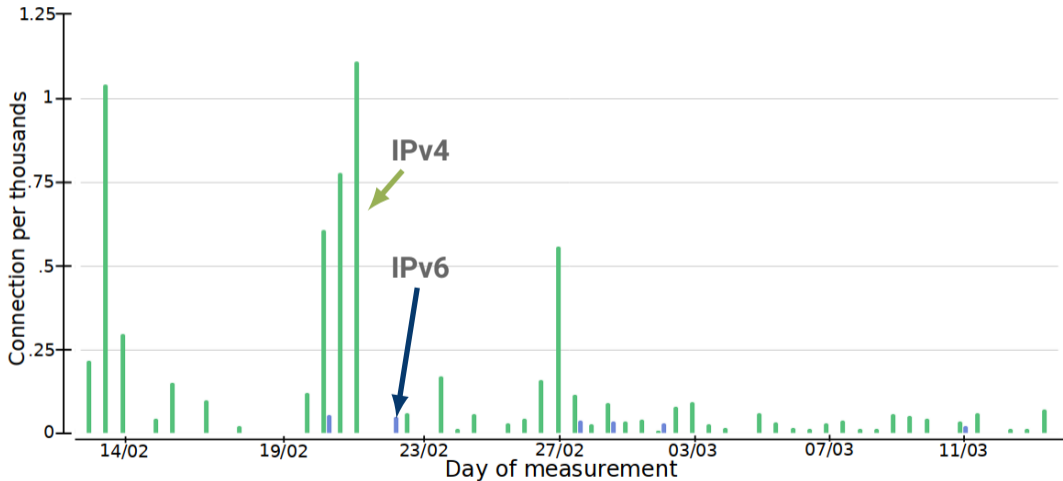
Most connections are made over IPv6



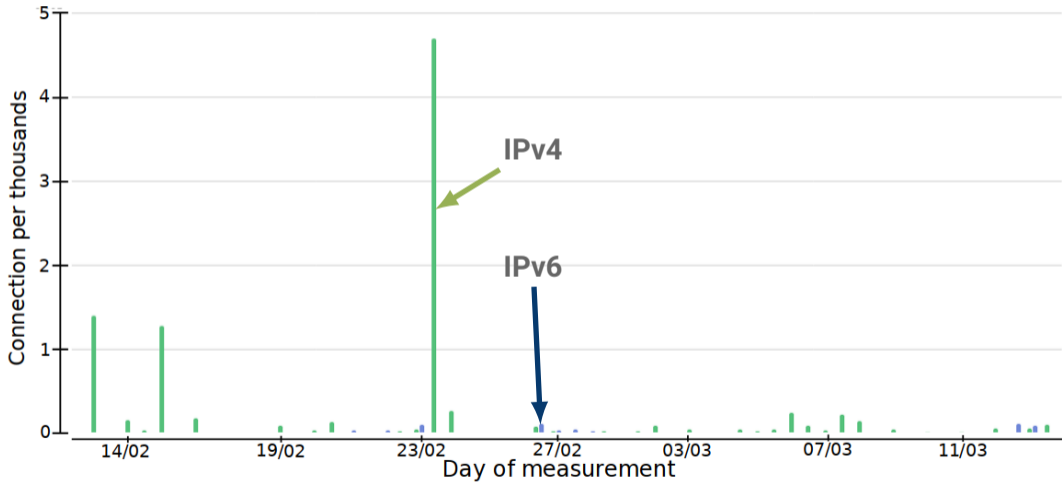
Querying TCP connections experiencing SYN retransmissions across IPv4/IPv6



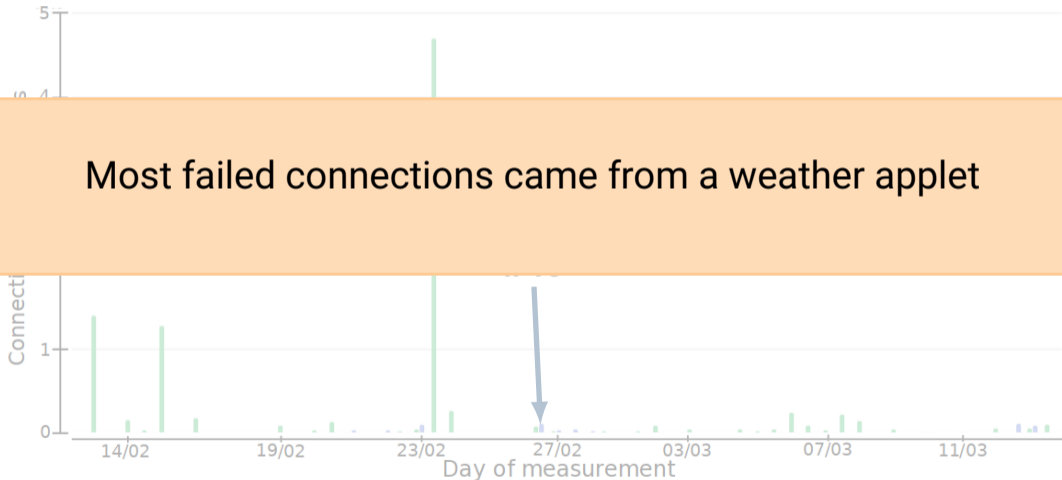
Comparing TCP connections experiencing more than one lost SYN



Querying TCP connections **that failed to establish**

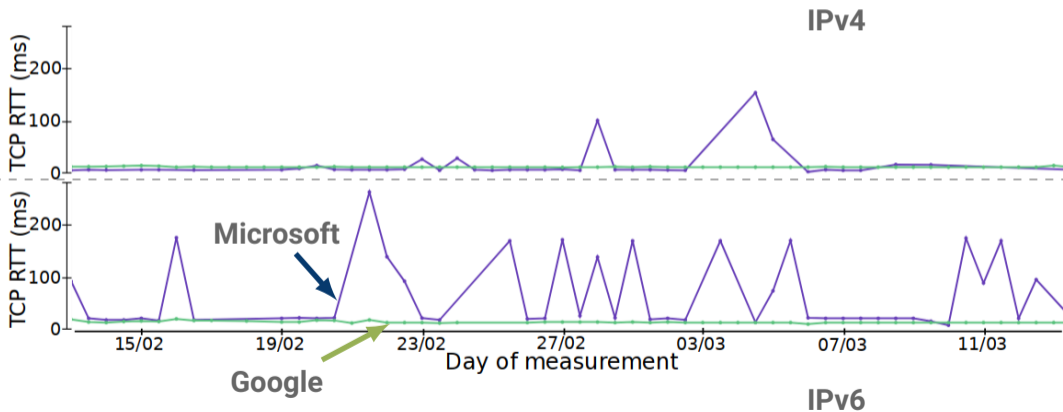


Querying TCP connections **that failed to establish**




Most failed connections came from a weather applet

Comparing median TCP RTTs across cloud providers and IPv4/IPv6



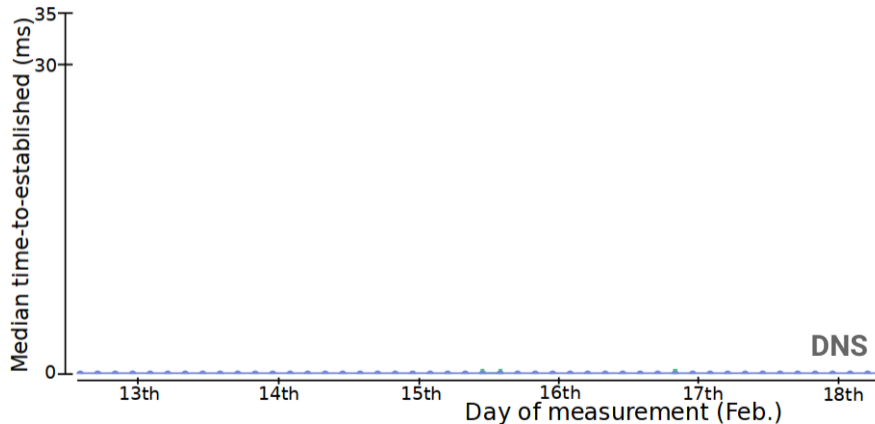
Comparing median TCP jitters across IPv4/IPv6 and similar destinations



“Moodle is slow”


“Moodle is slow”

Data-center is 200 m away from instrumented hosts



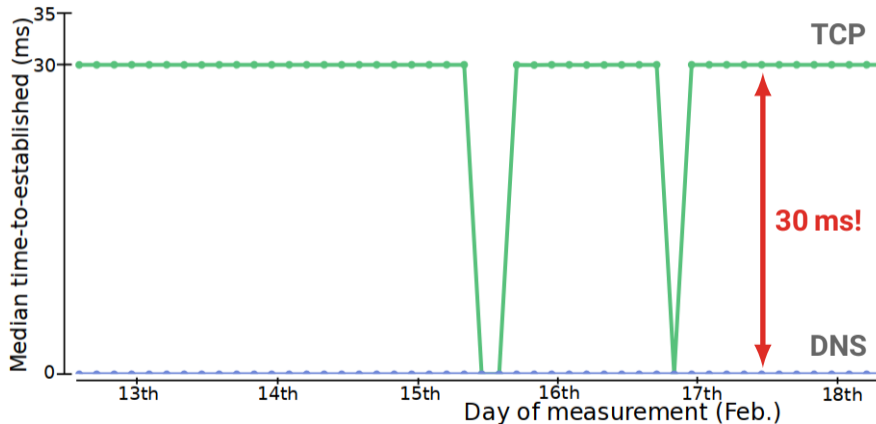
“Moodle is slow”

Data-center is 200 m away from instrumented hosts



“Moodle is slow”

Data-center is **200 m** away from instrumented hosts




“Moodle is slow”

Data-center is **200 m** away from instrumented hosts



The load-balancer was overloaded




An endhost-centric approach to detect network performance problems

olivier.tilmans@uclouvain.be

Good

Poor



- Instruments end-host transport stacks
- Supports encryption transparently
- Defines a generic instrumentation approach
- Leverages existing tooling around IPFIX