Some Perspectives on the Performance Impact of Link-Speed Switching Outages

Vern Paxson

International Computer Science Institute and
Lawrence Berkeley National Laboratory
Berkeley, California USA

vern@icsi.berkeley.edu

July 18, 2007

Where I’m Coming From

- Networking researcher:
  - Measurement/performance
  - (also security)
- Lots of work in the past on transport protocols in general (IETF co-Area Director for Transport) …
- … and TCP in particular
  - Co-author of RFCs on TCP congestion control, retransmission timeout management, implementation problems
- But: no “skin in the game”
  - I don’t really care what you decide
  - Just would like to provide some useful input
- (Though new project w/ K. Christensen, B. Nordman on how energy efficiency affects network architecture)

Some Assumptions I Make

- Link speed transitions are rare
  - If the point is energy efficiency, this only pays off for lengthy periods of low-speed use (lots of idle time)
    - Sub-assumption: transitions have a cost
    - (If transitions very cheap, could consider making them all the time)
- Logic for deciding transitions can be compartmentalized
  - Including whether or not management is involved
- Logic is aware of explicit real-time mechanisms like AVB
- Dampening is used to avoid oscillations
  - Can be as simple as “once high, stay high for quite a while”
- If decision made to switch to lower speed, *very unlikely* a transition outage coincides with link use

Some Questions I Have

- Just how much minimal buffer will future switches have?
- When the rate changes, is it bidirectional or unidirectional?

Outages and TCP

- For TCP performance, can consider three types of packet loss:
  - SYNs
  - Acks
  - Data

- SYNs: if lost, sender endures lengthy timeout
  - But: they're not going fast yet anyway
  - And: they often incur such delays already due to DNS latency or congestive loss
  - And: if transitions are rare, then impact here very low

Outages and TCP, con't

- Acks: loss of individual packets has very minor impact on TCP throughput …
  - … unless entire window lost …
  - … which requires an outage of an RTT duration
- Data:
  - Important to keep in mind that most connections are small and slow
    - Thus, unlikely to have data in flight during transition
    - And even if so, they're already inconsequential
  - Only real concern is high speed connections
  - These have key property: <u>an entire RTT's worth of data is in flight</u>

High-Speed TCP Connections

- A modern TCP that has plenty of data to send <u>does not time out upon packet loss</u> …
- … unless all data packets are lost.
- I.e.: requires transition outage $\geq$ RTT

- In addition: recovery time from loss scales inversely with RTT
  - Window expands (slow start or cong. avoid.) per RTT
- What if RTT is small (e.g., 1-10 msec)?
  - Then maybe risk timeout, but recovery from it is quick

High-Speed TCP Connections, con't

- What if no timeout?
- Then rate halved, with linear recovery …
  - … per RTT

- Summary for TCP:
  - Impairment will be truly rare
  - Requires (rare) large/high-speed connection, and …
  - … timeout will only occur for those with small RTTs
    - Best able to recover, especially intra-LAN
  - … rate-halving only problematic for those with high RTTs (100s of msec)
    - And these are particularly rare for large + high-speed

What About (Streaming) Non-TCP?

Important to distinguish:

- … one-way streaming
  - For which there will already be a significant play-out buffer
- … Internet multimedia (soft/adaptive real-time)
  - Which will already be written to adapt to varying conditions and loss
- … two-way interactive
  - This is the potentially hard case
  - If it uses LAN features (e.g., AVB) then transition logic can monitor its use
  - Also note: cell phones show that users are fine with occasional glitches, even of 100s msec or higher

Summary

- My two cents:
  - $\leq$ 1 msec: just can't be any real problem
  - $\leq$ 10 msec: this is going to work fine, no one will ever notice
  - $\leq$ 100 msec: you can get away with this too, but this is the time scale where you start to want to explore just what you're trading off