

Generic Robust Low Latency Tunneling

A proposal for a performance enhancing tunnel

Marie-José Montpetit (marie@mjmontpetit.com)

Brandon Williams, Akamai Technologies (brandon.williams@akamai.com)

July 20 2017

Some New Use Cases

- Wireless and streaming have traditionally been the focus of FEC and performance enhancing proxies and proxy based FEC
- But other use cases are on the rise:
 - Enterprise (overlay) networks
 - Cloud computing (in and out of the cloud)
 - A rebirth of heterogeneous mobile/fixed networks (airplane p2p, multi-satellite, trains etc.)

What is needed are

flexible, dynamic, application and policy-based mechanisms

to enhance the performance of rising services beyond traditional QoS

A Robust Low Latency Tunnel

- Tunnel with *dynamic negotiated performance enhancement*
 - A protocol not a de-facto tunnel
- *End-to-end semantics*
- *In user space*
 - Will not be impacted by middleboxes
 - Will not interact with crypto
 - No transport termination hop by hop when not needed

What Needs to Be Negotiated?

- FEC algorithm/implementation
 - Depends on both technical and legal constraints and requirements *
- Reliability requirements
 - Based on packet error rates and profiles
- In-order delivery (y/n)
- Delay tolerance vs. reliability
- Packet pacing to maintain constant delivery rate
- Congestion control / fairness?
 - For example: do we need the protocol to be TCP aware
- Micro-flow support vs. muxing
 - Micro-flow independence from tunnel protocol evolution

What Exists/Can Be Built Upon

- A growing number of FEC codes/protocols
- QUIC tunneling codes
- TAPS concepts to decouple application from transport
- New Path Aware Networking Research
- Other (ICCRG, SDNRG etc.)

Next Steps

- Architecture draft
- Implementation draft

Is this is important work for the future of NC as part of the network performance toolkit?

Acknowledgements: Discussions Leading to this Presentation

- Vincent Roca (INRIA)
- Ian Swett (Google)