

TSVarea Open Meeting Minutes

Thurs 9am-11:30am March 6, 2014

Martin started with Note Well, then briefed the latest progress in TSVarea

- 3 BOFs this week:
 - Virtualized Network Function Pool (VNFPOOL)
 - Tunneling Compressed Multiplexed Traffic Flows (TCMTF)
 - Transport Services (TAPS)
- 1 new WG
 - TURN Revised and Modernized (TRAM)

Spencer presented TSV AD job description to address recent Nomcoms having trouble finding TSV area directors.

Spencer also listed a few changes to the workload at the IESG and TSV area levels:

- The IESG is including document shepherds on telechats and relying on shepherds more during review processing.
- The TSV Ads have recruited a TSVarea secretary, and plan to start a TSV review team.
- Need to encourage more nominees who are willing to be considered.

Stewart Bryan: I find it more than a full time job. WG can help

Joel (Op AD): the amount of effort that AD comes up to speed varies. Even with my many years working in IETF, I still find that I have to spend a lot of time when elected as AD. It gets easier after a year. It is less than a full time job now. I haven't heard any complain yet.

Lars: You should nominate whoever is qualified, even though there is incumbent there. We need more candidates.

Sharp: a lot of works in Transport Area are research oriented. Doing a lot of research is good. But we need to do more engineering work, to make it relevant to the industry.

Spencer: had a talk with Dave Mayer on TSD AD, Dave said that finding TSV ADs is more difficult because TSV deals with a lot of mathematic analysis – most other areas don't have the same requirement.

Lady (in pink shirt): I like the work in research area.

Lars: TSV AD needs to help more on research area.

Michelle Cotton (IANA): We need 2-3 people help out reviewing requests for port assignments. It is not large work load. We want to ease the work for the existing reviewers.

Technical presentations:

1. **Stochastic Forecasts Achieve High Throughput and low delay over cellular networks – Keith Winstein MIT**

Observation: videoconferences perform poorly over cellular network.

Cellular different from traditional network: link speed change very much.

ICMP-ping is very reliable. Every packet sent on the network does come back. Some takes 10 seconds, which is a big problem for video conference.

Interactive apps work poorly:

- We measured cellular network while driving: Verizon LTE, ATT's, Sprint's
- Skype, Google Hangout, Apple FaceTime

Skype performance:

- Every time Skype sends out a packet, it gets queued. Some packets get queued for 9 seconds.
- Average delay is more than a second.

What is wrong?

- Existing schemes react to congestion signals; packet loss, increase in round-trip time
- Feedback comes too late
- The killer: self-inflicted queuing delay

Sprout: new model:

- Motivated by "Data from cellular network are not steady"

Stuart Cheshire (Apple): I don't think it is a problem on the simulated link.

Bob Briscoe (BT): It is not a fair comparison. It will go up and then go down. The rate cycle will go down. The background will take away some.

Keith: it is possible that Sprout is not so good in real audio conference.

Gorry: Is Reno same as Cubic?

Keith: My guess is same as Cubic

Michael Scharf: the problem with you presenting this work is that large carriers will go to their LTE vendor, ask for a middle box to do the same thing.

Sprout is end to end, but comparable to in-net control.

The real fix is to convince provider to install some queuing management. Then this problem can be fixed.

If Cubic can be installed in Verizon, then the throughput and delay can be greatly improved.

Michael Scharf: if you ask operators to install it, then should go to convince them.

Keith: it is right that we can't convince operator to install anything. We can only do from application ends.

If application knows the network situation, the applications can improve its throughput/delay.

Limitation:

- Sprout wants to control all of the traffic on a queue
- What if cell link isn't the bottleneck
- Assumption: application always has data to send

It might be helpful to send garbage data to make the rate steady when the applications are not sending any data.

Summary: we should model the network link, infer current link

Ericsson participant : I see a lot of problems, you need to be more careful with this model. We also observe that you don't have bandwidth to send large frames. There are some improvements. The limitations you listed are the important issues to be included.

Keith: I'd love to see how Sprout with video encoder.

Dan Frost (Cisco): have you done Sprout on other networks?

Keith: yes, but mostly the on shaking links.

Dan Frost: do you have this tool in the Open source it would be very beneficial for other to experiment.

Keith: yes.

Jana Iyengar (Google): Curious what happen

2. Unicast UDP Usage Guidelines For Application Designers - Lars

RFC 5405/BCP 145 was written over a two-year period while Lars was serving as AD.

General guideline:

- Must tolerate a wide range of internet path conditions

Congestion control:

- Should control rate of transmission
- Should perform congestion control over all traffic

Not only need to monitor each flow, but also need to monitor how many parallel flows.

Bulk transfer applications:

- For bulk transfers;
 - o Should consider implementing TFRC
- For non-bulk transfers
should measure RTT and transmit Max.1 datagram/RTT
- Else should send at most 1 datagram every 3 seconds.

Stewart Bryan: it is common to send 3 in a row

UDP tunnels:

- For tunnels carrying IP traffic, should not perform congestion control
- For non-IP tunnels or rate not determined by traffic, should perform congestion control.

Gorry: “should not perform congestion control” means should not do TCP-type mechanism.

Message size guidelines:

- Should not send datagrams that exceed the PMTU
- Should discover PMTU

Reliability guidelines:

Middlebox traversal guidelines:

- Should not always send middlebox keep-alive
- May use keepalives when needed

ICMP guidelines:

Thomas Narten: This is like butter and apple-pie. How closely are applications in meeting those guideline? Is this something that IETF recommend, but no one follows? There are a lot of tunnels being built, but no one wants to touch them.

Lars: I don't know.

Ed Crabbe (Google): BCP document is not mandate. Do you think IETF should standardize inter-domain protocols?

Lars: guideline is more: if you don't follow this, you can hurt others. For security, you hurt yourself if you don't follow guidelines.

Joel Halpern: The UDP tunnel has been changed dramatically. This is not reality. You can't assume that IP has congestion control

Lars: There are some assumption built in the applications

Thomas Narten: My view is this doc is good. On the other hand, we as a community are so afraid of breaking internet. But today, whenever there is something wrong with internet, there are big service providers jumping in to fix the problem.

Lars: there is a problem of the term "congestion control"

Thomas Narten: in the past, there is risk

Stuart Cheshire (Apple): If you are successful, other people will use protocols in ways you never imagined. This is what the IAB says in RFC 5218 on successful protocols.

Erik Normark: we try to do the right thing for people to use our protocols. A lot of Tunnel protocols carry Ethernet header

Lars: When there is no middle box, there is no difference.

Curtis: there are situations, every time you tunnel UDP, you create some issue. It is not a good idea to tunnel in UDP.

Carsten Borman:

Ross Callon (Juniper): we had discussions at various places. If we want our protocols to be used, we have to pay attention to what operators want to do. There are billions of billions dollar, they are capable of doing load splitting, TCP/UDP. Most routers can't do load splitting on GRE tunnels. Many years ago, there are many types of links that need congestion control. Today, many of those link types don't exist anymore.

Ed Crabbe: if use "Should use" in the document, is it better?

Tim: why need to mention this document is over 6 years old

Martin: The time of the document is not important.

Tim: I assume that mentioning “6 years old” indicates it is very out of date. Time has changed, the recommendation was good many years ago.

Jana (Google): what is important is what is on top of UDP.

3. **Experiments with ECN – Stuart Cheshire (Apple)**

- Apparent State of ECN: seems a good way to reduce delay/loss in the network.
- Turned off – present but actually used.
- Turning on has no effect: pretty much same as TCP

TCP for streaming video:

- TCP handles packet loss extremely well
- But.. causes irregular data delivery to client. No problem for file transfer. But problem for streaming video over TCP: Netflix, YouTube, etc.

Conclusion:

- CODEL helps, ECN helps a little, AQM+ECN really helps a lot

Why don't we have ECN: doesn't help much. Sometimes worse because bugs in the code.

- Clients:
- Routers: no one ask, enabling ECN might expose code bug (risk)
- All risk; no reward

Dave Oran (Cisco): ECN alone doesn't help very much. ECN coupled with AQM helps.

Jana (Google): I don't see how you can do ECN without weight marking.

Stuart: smart AQM will have the marking.

Michael Tuexen: I want to implement protocol on top of UDP

4. **ECN the identifier of a new service model – Bob Briscoe BT**

- We don't deploy anything if it is not coupled with a service.
- Classic ECN wasn't useful enough to deploy
- Opportunity to identify a new service model using ECN
 - o Consistent low delay service for all.

-
- Incremental deployment path
- ECN is not about drop

Classic ECN: just avoiding loss has not been good enough.

- Classic ECN = “drop equivalence”
- Cannot justify the deployment pain
- For a small performance gain.

ECN can do so much more.

Immediate ECN + Smooth TCP response.

- Inspired by data center TCP (DCTCP)

Lady (in pink sweater): question on your Round Trip delay: TCP slow start.

Stuart: it is not slow start. TCP consistent probes of the network to determine the network speed.

Lady (in pink sweater): if you have short round trip flow and long round trip flow together

Dave (Cisco): stress that we can barely get AQM deployed.

Bob: I am not saying not deploy AQM, but don't hard code the “drop” parameters.

Jana: have you considered ECN on top of ECN.